

Università Degli Studi Di Pavia

Dipartimenti di Scienze Politiche e Sociali, Giurisprudenza, Ingegneria Industriale e
dell'Informazione, Scienze Economiche e Aziendali, Studi Umanistici

Corso Di Laurea Magistrale Interdipartimentale in Comunicazione Digitale

AI-Driven Transformation in Pharmaceutical Digital Advertising: Balancing
Innovation, Regulatory Constraints, and Ethical Risks

Relatore:

Prof. Alfredo Sassi

Correlatore:

Alessandro Ferrari

Tesi di laurea di

Basim Ahmed Hanin

ANNO ACCADEMICO 2025/2026

Table of Contents

Chapter 1: Introduction	1
1.1 The Digital Transformation of Life Sciences	1
1.2 Problem Statement: Innovation vs. the Regulatory Wall.....	2
1.3 Generation Z and the Trust Deficit.....	3
1.4 Research Questions	4
Chapter 2: The Digital Pharmaceutical Paradigm: Personalization, Skepticism, and Regulatory Friction.....	6
2.1 From Mass Marketing to Hyper-Personalization: The Algorithmic Shift	6
2.1.1 The Decline of the One-Size-Fits-All Paradigm	7
2.1.2 Precision Profiling and the Next Best Action.....	8
2.1.3 The Precision Ecosystem and the Autonomy Risk	10
2.2 Digital Skepticism and Ad Avoidance Behaviors: The Role of Generation Z	12
2.2.1 The Algorithmic Intuition and the Trust Deficit	13
2.2.2 The Platform Paradox: Visual Validation vs. Scientific Evidence.....	15
2.2.3 The Human Element Requirement in the AI Era.....	17
2.3 International Regulatory Landscapes and the Regulatory Lag.....	19
2.3.1 The Jurisdictional Gray Areas of Social Media.....	21
2.3.2 Data Privacy vs. Marketing Liberty: The Battle over Tracking	22
2.3.3 Borderline Contexts and the Failure of Self-Regulation	24
Chapter 3: AI Technologies in the Creative and Strategic Process	27
3.1 Generative AI for Content Automation and Medical Compliance	27
3.1.1 Enterprise LLMs and Retrieval-Augmented Generation (RAG) in Pharma Copywriting.....	28
3.1.2 Modular Content Creation and Hyper-Scalability.....	30
3.1.3 Automated Pre-Screening and Compliance by Design.....	32
3.2 Predictive Architectures and Next Best Action Engines	36
3.2.1 Predictive Behavioral Modeling and Data Ingestion Pipelines	37
3.2.2 The Next Best Action Orchestration Framework	38
3.2.3 Dynamic Creative Optimization and Algorithmic Visual Strategy	40
3.3 The Personalization Paradox: Balancing Relevance and User Privacy.....	42
3.3.1 Algorithmic Reactance and the Privacy Paradox in Healthcare	42
3.3.2 The Privacy Utility Tradeoff and Data Minimization Architectures	44
3.3.3 Resolving the Paradox: Risk-Proportionate Functional Asymmetry and the Data + Empathy Framework	46

Chapter 4: The Regulatory Wall and Borderline Contexts.....	48
4.1 Surrogate and Hidden Advertising: Regulatory Gaps in High-Exposure Contexts.....	49
4.2 The Role of Digital Platforms: Algorithmic Self-Regulation vs. National Regulatory Authorities.....	51
4.2.1 Case Study: The Semaglutide Communication Crisis.....	53
4.3 Compliance by Design: Leveraging AI as a Pre-Screening Tool for Regulatory Alignment.....	56
4.4 Ethical Dilemmas: Algorithmic Transparency and the Risk of AI-Generated Misinformation.....	59
4.5 Toward Responsible AI Governance in Pharmaceutical Digital Communication.....	61
Chapter 5: Empirical Analysis: Case Studies and Field Research.....	63
5.1 Comparative Analysis of Selected AI-Driven Initiatives	65
5.1.1 Case 1: Eurofarma - Scrolling Therapy (Functional AI as a Clinical Rehabilitation Utility)	65
5.1.2 Case 2: ASCO - The Most Beautiful Sound (Translational Acoustic Data as Emotional Oncology Support)	70
5.1.3 Case 3: Dove - The Code (Critical Benchmark: Ethical Resistance to Algorithmic Aesthetic Homogenization).....	74
5.1.4 Cross-Case Synthesis: Key Strategic and Operational Findings	81
5.2 Applying the Data + Empathy Framework to AI-Driven Pharma Communication	103
5.2.1 Architectural Configuration: Risk-Proportionate Functional Asymmetry.....	105
5.2.2 Cognitive and Behavioral Considerations: Addressing Aversion, Reactance, and Authenticity	110
5.2.3 Strategic and Operational Implementation Matrix	113
5.3 Field Research: Generation Z Perceptions of AI-Mediated Healthcare Communication	116
5.3.1 Research Design and Data Collection	117
5.3.2 Survey Instrument and Analytical Strategy	118
5.3.3 Quantitative Findings: Technological Role Allocation, Trust, and Privacy Dynamics	120
5.3.4 Qualitative Findings: Human Responsibility in AI-Mediated Healthcare Communication	126
5.3.5 Integrated Interpretation: Empirical Triangulation through the Data + Empathy Framework.....	130
Chapter 6: Conclusions and Future Outlook	136
6.1 Synthesis of Findings and Response to the Research Questions.....	137
6.2 Theoretical Contribution: The Data + Empathy Framework.....	140
6.3 Managerial and Governance Implications for Pharmaceutical Digital Communication.....	141
6.4 Research Limitations.....	143
6.5 Future Research and Concluding Remarks	145
References	147

Abstract (English)

The pharmaceutical industry is undergoing a profound digital transformation driven by Artificial Intelligence (AI) and advanced data analytics, progressively shifting from traditional mass-broadcast models toward more personalized, context-aware, and omnichannel forms of communication. However, this technological acceleration interacts with significant institutional and organizational constraints, including the “pacing problem” of regulatory lag, complex data-privacy requirements, ethical concerns surrounding algorithmic transparency, and evolving expectations regarding trust, accountability, and human oversight. This thesis investigates how AI enables the transition from mass marketing to precision communication, how pharmaceutical organizations navigate regulatory gray areas, how Generation Z responds to AI-mediated health communication, and how data-driven precision can be integrated with human-centered communication.

Employing a multidimensional research design, the study combines a technical and regulatory analysis of advanced AI architectures—including Retrieval-Augmented Generation (RAG), Knowledge Graphs, Next Best Action (NBA) engines, and Compliance-by-Design pre-screening—with a comparative analysis of three international initiatives: *Eurofarma’s Scrolling Therapy*, *ASCO’s The Most Beautiful Sound*, and *Dove’s The Code*. The research further incorporates an exploratory field study involving 147 respondents aged 18–29, combining quantitative rating scales with qualitative thematic analysis.

The empirical findings reveal a pattern of *conditional acceptance*. High levels of everyday AI familiarity coexist with clear boundaries concerning technological authority: 82.3% of respondents preferred direct interaction with a healthcare professional in serious health scenarios, while 82.3% assigned high importance to knowing whether health-related content had been generated or modified through AI. Respondents also displayed substantial discomfort with behavioral personalization, while attitudes toward synthetic empathy were predominantly neutral or negative. These findings suggest that technological familiarity does not automatically translate into willingness to delegate sensitive or consequential healthcare functions to Artificial Intelligence.

To interpret these tensions, the thesis develops the **Data + Empathy framework**, grounded in the principle of *risk-proportionate functional asymmetry*. Rather than imposing a rigid separation between an invisible AI backend and a human front-end, the framework proposes that the appropriate degree of AI autonomy, visibility, and human oversight should vary according to functional complexity, communicative sensitivity, evidentiary requirements, and potential consequences. Computational systems may therefore assume greater autonomy in standardized and lower-risk informational functions, while human involvement becomes increasingly important as communication approaches clinical interpretation, emotional vulnerability, uncertainty, or consequential decision-making.

The findings ultimately suggest that responsible pharmaceutical digital communication requires a transition from capacity-driven automation maximization toward utility-oriented, transparent, and context-sensitive technological role allocation, supported by proportionate human governance and accountability.

Keywords: *Artificial Intelligence, Pharmaceutical Digital Advertising, Data + Empathy Framework, Regulatory Lag, Compliance-by-Design, Generation Z, Algorithmic Trust, Omnichannel Communication.*

Abstract (Italiano)

L'industria farmaceutica sta attraversando una profonda trasformazione digitale guidata dall'Intelligenza Artificiale (IA) e da sistemi avanzati di analisi dei dati, passando progressivamente da tradizionali modelli di comunicazione di massa a forme di comunicazione più personalizzate, contestualizzate e omnichannel. Tuttavia, questa accelerazione tecnologica si confronta con rilevanti vincoli istituzionali e organizzativi, tra cui il “problema del ritmo” (*pacing problem*) legato al ritardo normativo, complessi requisiti in materia di protezione dei dati, questioni etiche relative alla trasparenza algoritmica e aspettative crescenti in termini di fiducia, responsabilità e supervisione umana. La tesi analizza come l'IA abiliti la transizione dal marketing di massa alla comunicazione di precisione, come le organizzazioni farmaceutiche affrontino le aree grigie normative, come la Generazione Z reagisca alla comunicazione sanitaria mediata dall'IA e come la precisione basata sui dati possa essere integrata con una comunicazione human-centered.

Adottando un disegno di ricerca multidimensionale, lo studio combina un'analisi tecnica e normativa di architetture avanzate di IA — tra cui Retrieval-Augmented Generation (RAG), Knowledge Graph, motori Next Best Action (NBA) e sistemi di pre-screening Compliance-by-Design — con un'analisi comparativa di tre iniziative internazionali: *Scrolling Therapy* di Eurofarma, *The Most Beautiful Sound* di ASCO e *The Code* di Dove. La ricerca comprende inoltre uno studio empirico esplorativo condotto su 147 rispondenti di età compresa tra 18 e 29 anni, integrando scale quantitative di valutazione e analisi tematica qualitativa.

I risultati empirici evidenziano un modello di *accettazione condizionale*. Un'elevata familiarità quotidiana con l'IA coesiste con confini ben definiti rispetto all'autorità attribuita alla tecnologia: l'82,3% dei rispondenti ha espresso una preferenza per l'interazione diretta con un professionista sanitario in scenari di salute seri, mentre l'82,3% ha attribuito elevata importanza alla possibilità di sapere se un contenuto sanitario sia stato generato o modificato attraverso l'IA. I partecipanti hanno inoltre mostrato un significativo disagio rispetto alla personalizzazione basata sul comportamento online, mentre gli atteggiamenti nei confronti dell'empatia sintetica sono risultati prevalentemente neutrali o negativi. Tali risultati suggeriscono che la familiarità tecnologica non si traduce automaticamente nella disponibilità a delegare all'Intelligenza Artificiale funzioni sanitarie sensibili o ad alto impatto.

Per interpretare queste tensioni, la tesi sviluppa il **framework Data + Empathy**, fondato sul principio di *asimmetria funzionale proporzionata al rischio*. Anziché proporre una separazione rigida tra un'IA confinata a un backend invisibile e un'interazione front-end esclusivamente umana, il framework sostiene che il grado appropriato di autonomia, visibilità e supervisione dell'IA debba variare in funzione della complessità del compito, della sensibilità comunicativa, dei requisiti evidenziatori e delle potenziali conseguenze dell'interazione. I sistemi computazionali possono quindi assumere maggiore autonomia nelle funzioni standardizzate, informative e a rischio relativamente basso, mentre il coinvolgimento umano diventa progressivamente più importante quando la comunicazione riguarda interpretazione clinica, vulnerabilità emotiva, incertezza o decisioni ad alto impatto.

Nel complesso, i risultati suggeriscono che una comunicazione digitale farmaceutica responsabile richieda il passaggio dalla massimizzazione dell'automazione basata sulla capacità tecnologica a un'allocazione dei ruoli tecnologici orientata all'utilità, trasparente e sensibile al contesto, sostenuta da una governance umana e da forme di responsabilità proporzionate al rischio.

Parole chiave: *Intelligenza Artificiale, Pharmaceutical Digital Advertising, Data + Empathy Framework, Regulatory Lag, Compliance-by-Design, Generazione Z, Algorithmic Trust, Omnichannel Communication.*

Chapter 1: Introduction

1.1 The Digital Transformation of Life Sciences

The pharmaceutical industry is undergoing a significant digital transformation driven by the convergence of advanced data infrastructures, Artificial Intelligence, increasingly fragmented media environments, and evolving stakeholder expectations. Traditional communication models based primarily on broad audience segmentation and relatively standardized outreach are progressively being complemented by more data-informed, personalized, and digitally mediated approaches to healthcare communication (Jones et al., 2024; Sabapathy, 2025).

Within this transformation, Artificial Intelligence has acquired growing strategic relevance across multiple stages of pharmaceutical communication. Machine learning and generative technologies can support information retrieval, audience analysis, content adaptation, predictive decision-making, workflow automation, and regulatory pre-screening. These capabilities potentially allow life sciences organizations to move from generalized communication toward more contextual forms of interaction in which informational content, timing, channel, and level of personalization can be adapted to different stakeholder needs (Jones et al., 2024; Sabapathy, 2025).

However, the strategic relevance of Artificial Intelligence in pharmaceutical communication cannot be assessed exclusively in terms of computational efficiency. Healthcare communication operates within an environment characterized by scientific uncertainty, legal constraints, informational asymmetries, privacy concerns, and varying levels of stakeholder vulnerability. A communication system that increases technological precision may simultaneously introduce new forms of opacity, intrusive personalization, misinformation, or inappropriate delegation of responsibility. The technological transformation of the sector therefore raises a broader communication question: not simply how Artificial Intelligence can be used, but **which functions should be delegated to computational systems, under what conditions, and with what degree of human oversight.**

This question becomes particularly significant as pharmaceutical organizations increasingly interact with stakeholders who are familiar with digital technologies while simultaneously exposed to algorithmic targeting, online misinformation, synthetic media, and extensive commercial data collection. In such an environment, technological sophistication alone cannot be assumed to generate trust. The perceived legitimacy of AI-mediated communication may also depend on factors such as scientific provenance, transparency, privacy, functional utility, and the continued presence of identifiable human accountability.

This thesis examines these tensions within pharmaceutical digital communication by integrating technological, regulatory, behavioral, and strategic perspectives. Its central objective is to investigate how the opportunities created by Artificial Intelligence can be reconciled with the requirements of scientific credibility, stakeholder autonomy, regulatory responsibility, and human-centered communication.

To address this challenge, the research progressively develops the **Data + Empathy framework**, a conceptual and applied model for evaluating the appropriate role of Artificial Intelligence within pharmaceutical communication. Rather than treating technological capability as sufficient

justification for deployment, the framework considers computational possibility alongside human context and organizational responsibility.

In its final formulation, Data + Empathy does not prescribe an absolute separation in which Artificial Intelligence must always remain invisible while humans exclusively occupy the visible communication interface. Instead, the framework proposes a principle of **risk-proportionate functional asymmetry**: the appropriate degree of AI autonomy, visibility, and human oversight should vary according to the function performed, the sensitivity of the communication context, the evidentiary requirements involved, and the potential consequences of technological failure.

Under this perspective, Artificial Intelligence may assume comparatively greater autonomy in standardized, informational, administrative, or easily reversible functions, while human involvement becomes increasingly important as communication approaches clinical interpretation, emotional vulnerability, uncertainty, or consequential decision-making. The objective is therefore neither maximum automation nor generalized technological restraint, but **appropriate technological deployment**.

1.2 Problem Statement: Innovation vs. the Regulatory Wall

The central problem addressed by this thesis emerges from the tension between the expanding capabilities of Artificial Intelligence and the unusually demanding regulatory, ethical, and communicative conditions of the pharmaceutical sector.

AI-mediated technologies can increase the speed and granularity with which organizations analyze stakeholder behavior, retrieve scientific information, adapt content, and coordinate digital communication. At the same time, their deployment introduces risks related to algorithmic opacity, inappropriate behavioral profiling, biased or misleading outputs, insufficient scientific substantiation, and the generation of communication that may appear authoritative despite lacking adequate evidentiary support (Amini Khiabani & Rezaeian, 2026; Jones et al., 2024; Sabapathy, 2025).

These risks become particularly significant in healthcare because errors in communication cannot be evaluated exclusively as failures of advertising effectiveness. Health-related information can influence perceptions of illness, therapeutic expectations, decisions to seek professional care, and broader relationships of trust between organizations, healthcare professionals, and patients. As a consequence, pharmaceutical digital communication must reconcile the speed and adaptability associated with contemporary computational systems with regulatory structures designed to prioritize scientific substantiation, patient safety, and accountable communication.

One dimension of this tension concerns what technology-governance literature describes as the **pacing problem**. Technological systems, platform architectures, and digital communication practices may evolve considerably faster than the legislative and administrative processes through which new regulatory expectations are established and enforced (Marchant et al., 2011). This temporal asymmetry does not imply the absence of regulation, but it can create periods of uncertainty in which existing legal principles must be interpreted and applied to technological configurations that were not explicitly anticipated when those rules were developed.

The resulting regulatory friction is particularly visible within decentralized digital environments. Pharmaceutical organizations operate across platforms in which corporate communication coexists

with user-generated content, influencer narratives, disease-awareness material, wellness communication, synthetic media, and promotional practices that may cross jurisdictional or categorical boundaries. Borderline products and indirect health communication further complicate the distinction between information, advertising, education, and commercial persuasion.

A second dimension of the problem concerns personalization. Data-driven systems can make communication increasingly contextual and potentially useful to specific stakeholders. Yet the data collection and predictive inference required to achieve this precision may themselves generate discomfort when they are perceived as opaque, disproportionate, or intrusive. This tension constitutes the **Personalization Paradox** examined throughout the thesis: greater relevance does not necessarily correspond to greater perceived legitimacy or greater stakeholder trust.

The problem is therefore not reducible to a choice between innovation and regulation. Nor can it be resolved simply by increasing automation or, conversely, by excluding Artificial Intelligence from stakeholder communication. What is required is a governance approach capable of distinguishing between technological capabilities that create meaningful stakeholder utility and those whose deployment may introduce disproportionate informational, regulatory, or relational risks.

The Data + Empathy framework is developed in response to this problem. Within the framework, **Data** refers to the analytical and computational capabilities through which organizations can retrieve, process, infer, predict, and adapt information. **Empathy** is not conceptualized as synthetic emotional language produced by an algorithm, but as a human-centered evaluative principle through which organizations consider stakeholder needs, vulnerability, experience, and autonomy. **Governance** subsequently determines whether, where, and under what conditions technological capability should be translated into organizational action.

The resulting conceptual logic can be summarized as follows:

Data → Computational Possibility

Empathy → Human-Centered Appropriateness

Governance → Organizational Permission, Limitation, and Accountability

This framework provides the analytical foundation through which the thesis examines technological infrastructures, regulatory constraints, real-world communication initiatives, and stakeholder perceptions of AI-mediated healthcare communication.

1.3 Generation Z and the Trust Deficit

A further dimension of the research concerns younger digital audiences and, specifically, the relationship between technological familiarity and trust in AI-mediated health communication.

Generation Z represents a particularly relevant population for the study of contemporary digital communication because its members have developed within media environments characterized by social platforms, recommendation systems, personalized digital services, online advertising, and increasingly widespread access to Artificial Intelligence. At the same time, extensive exposure to digital commercial environments does not necessarily imply uncritical acceptance of algorithmic mediation.

Research on digital advertising has documented various forms of advertising avoidance and resistance to intrusive commercial communication, while studies of younger populations emphasize the increasingly complex role played by online environments in the discovery and evaluation of health-related information (Redondo & Aznar, 2018; Mansfield et al., 2025). Digital platforms provide substantial access to information and peer experience, but they can also expose users to misleading health narratives, synthetic content, inconsistent source quality, and blurred distinctions between commercial and informational communication.

These conditions complicate simplistic descriptions of younger audiences as either inherently enthusiastic about Artificial Intelligence or categorically resistant to it. Technological familiarity, trust, and willingness to delegate specific functions to AI are conceptually distinct variables. An individual may regularly use generative AI while remaining reluctant to rely on the same technology for diagnostic interpretation, highly personalized health recommendations, emotionally sensitive communication, or therapeutic decision-making.

This distinction is particularly important to pharmaceutical communication. AI-mediated systems can potentially improve access to information, reduce administrative friction, support healthcare professionals, and adapt educational material. However, increasing technological involvement also raises questions concerning source verification, disclosure, behavioral tracking, error tolerance, and the appropriate boundary between computational assistance and human authority.

Rather than assuming a uniform generational attitude toward Artificial Intelligence, this thesis therefore investigates **conditional acceptance**: the possibility that younger stakeholders differentiate between AI applications according to the function performed and the perceived sensitivity of the context.

This perspective subsequently informs the exploratory field research developed in Chapter 5, which examines attitudes among respondents aged 18–29 toward AI-mediated healthcare information, scientific provenance, disclosure, behavioral personalization, synthetic empathy, technological error, and human responsibility.

The relevance of Generation Z to this thesis therefore extends beyond its status as a digitally familiar demographic cohort. It provides an analytically useful context in which to investigate a broader contradiction at the center of contemporary digital communication: the same audiences that routinely interact with advanced computational technologies may simultaneously establish clear boundaries concerning the degree of informational, emotional, and decisional authority they are willing to delegate to those systems.

Understanding these boundaries is essential for pharmaceutical organizations attempting to integrate Artificial Intelligence into communication strategies without assuming that technological familiarity automatically translates into technological trust.

1.4 Research Questions

Against this background, the thesis investigates the transformation of pharmaceutical digital communication through four interconnected research questions:

RQ1: How is Artificial Intelligence enabling the transition from mass marketing to precision communication in pharmaceutical digital advertising?

RQ2: How do pharmaceutical organizations navigate regulatory gray areas, including borderline contexts and platform-based self-regulation?

RQ3: What perceptions and trust dynamics characterize AI-mediated health communication among the Generation Z respondents examined in this research?

RQ4: How can human-centered evaluation and data-driven precision be integrated through the Data + Empathy framework to support responsible and credible AI-mediated pharmaceutical communication?

These questions are addressed through a multidimensional research design.

The first part of the thesis develops the theoretical and technological foundations of AI-mediated pharmaceutical communication, examining the transition from mass communication toward data-driven personalization, predictive architectures, Next Best Action systems, generative AI, Retrieval-Augmented Generation, modular content, Dynamic Creative Optimization, and automated compliance support.

The second analytical dimension examines the regulatory and ethical environment within which these technologies operate, with particular attention to regulatory lag, platform governance, indirect and borderline communication, algorithmic transparency, misinformation, Compliance by Design, and Human-in-the-Loop governance.

The third component applies these theoretical perspectives to three purposively selected communication initiatives: Eurofarma's *Scrolling Therapy*, ASCO's *The Most Beautiful Sound*, and Dove's *The Code*. These cases are not treated as empirical validation of the Data + Empathy framework, but as comparative examples through which different configurations of technological utility, scientific translation, and ethical boundary-setting can be examined.

Finally, the thesis incorporates exploratory field research involving younger adult respondents to investigate perceptions of AI-mediated healthcare communication. Quantitative descriptive analysis is complemented by an exploratory association analysis and a qualitative examination of open-ended responses. The empirical component is intended to refine the conceptual framework through stakeholder-level evidence rather than to establish statistically generalizable conclusions about Generation Z as a whole.

Taken together, these analytical levels enable the thesis to move from technological capability to communicative application, from regulatory constraint to organizational governance, and from theoretical assumptions concerning younger audiences to exploratory evidence concerning their stated expectations.

The resulting research perspective is based on a central proposition: *the strategic value of Artificial Intelligence in pharmaceutical communication depends not on the maximum extent to which communication can be automated, but on the extent to which computational capability is assigned an appropriate role relative to stakeholder need, contextual sensitivity, scientific evidence, and human responsibility.*

Chapter 2: The Digital Pharmaceutical Paradigm: Personalization, Skepticism, and Regulatory Friction

The introduction of Artificial Intelligence into pharmaceutical communication has not simply expanded the range of tools available to organizations. It has changed the conditions under which health information is produced, distributed, and evaluated. Pharmaceutical companies now operate in an environment where data can be used to adapt communication to specific audiences, digital platforms influence what information becomes visible, and users can compare institutional messages with a wide and uneven variety of online sources. Greater communicative precision has therefore developed alongside new concerns about privacy, autonomy, and the influence of algorithmic systems on individual exposure to health-related content.

This chapter examines the tensions created by this transformation. The transition from mass communication to hyper-personalized models illustrates how predictive profiling may increase relevance while making the boundary between useful adaptation and intrusion increasingly difficult to define. The attitudes associated with Generation Z further complicate this picture, since familiarity with digital platforms and Artificial Intelligence does not necessarily imply trust in their use within sensitive health contexts. These developments also take place within a fragmented regulatory environment that often evolves more slowly than the technologies it is expected to govern, leaving some forms of platform-mediated and borderline communication only partially addressed. This intersection of personalization, skepticism, and regulatory uncertainty provides the context for Chapter 3, which examines the technological architectures through which pharmaceutical communication is increasingly automated and personalized.

2.1 From Mass Marketing to Hyper-Personalization: The Algorithmic Shift

The historical architecture of pharmaceutical marketing was largely shaped by an environment characterized by substantial information asymmetry. Within this conventional framework, life sciences enterprises operated as centralized repositories of clinical and product-related knowledge, executing predominantly broadcast-oriented models of communication designed to disseminate generalized and highly standardized scientific or promotional messages to broad categories of Healthcare Professionals and patient populations alike (Jones et al., 2024). This paradigm was reinforced by the comparatively limited number of information-distribution channels available to medical practitioners, who historically relied heavily on physical detail aids, peer-reviewed journals, professional conferences, and periodic face-to-face interactions with pharmaceutical sales representatives to remain informed about therapeutic developments.

Consequently, patients generally occupied a less active position in the circulation of medical information than they do within contemporary digital environments, partly because access to specialized medical literature and patient-oriented information channels was considerably more limited. The legacy broadcast model, rooted in the corporate communication infrastructures of twentieth-century industrial organizations, relied on relatively centralized message production and mass-distribution mechanisms intended to influence professional awareness, prescribing behavior, and patient engagement through repeated exposure and controlled communication flows. The marketing process was therefore comparatively linear, moving primarily from pharmaceutical organizations toward physicians and patients, with more limited opportunities for real-time feedback concerning immediate cognitive reception or behavioral variation.

However, this top-down methodology has progressively transformed through its interaction with the contemporary fragmented digital ecosystem, where the expansion of web-based platforms has

contributed to the decentralization of healthcare information production and consumption (Verma et al., 2024). The contemporary transformation toward hyper personalization should therefore not be interpreted as a simple tactical optimization or a superficial update of existing advertising portfolios. Rather, it represents a broader structural evolution enabled by the convergence of increasingly sophisticated Big Data infrastructures and advanced Artificial Intelligence modeling.

This convergence can significantly alter the traditional dynamics of medical information consumption, reducing the extent to which communication departments independently determine the sequence and content of stakeholder exposure. Healthcare Professionals and patients increasingly interact with multiple information sources, while algorithmic systems can continuously ingest, categorize, and interpret diverse data signals. As a consequence, the relatively passive audience assumed by legacy communication models is progressively replaced by more active, digitally enabled stakeholders whose informational expectations may vary substantially according to context. This behavioral evolution encourages pharmaceutical organizations to move beyond exclusively linear broadcasting toward more dynamic, context-aware engagement frameworks capable of navigating the structural tension between increasing communicative precision and individual stakeholder autonomy (Sabapathy, 2025).

2.1.1 The Decline of the One-Size-Fits-All Paradigm

Traditional segmentation models within the pharmaceutical sector were historically based on relatively broad demographic criteria, clinical specializations, geographic boundaries, and other stable professional or patient classifications. These conventional classification mechanisms provided organizations with an operational basis for differentiating communication, but they could also implicitly assume a degree of homogeneity among Healthcare Professionals and patient cohorts, despite substantial differences in cognitive patterns, clinical challenges, informational requirements, and communication preferences. Such limitations have become increasingly visible when confronted with the multidimensional complexity of the contemporary Patient Journey (Verma et al., 2024).

In the current digital environment, the route from initial symptom recognition to clinical consultation, therapeutic decision-making, treatment management, and long-term health behavior is increasingly mediated by a continuous influx of online touchpoints, digital support networks, independent medical forums, search engines, and institutional resources (Sabapathy, 2025). This structural fragmentation means that two individuals sharing an identical diagnosis may nevertheless navigate substantially different informational trajectories, influenced by individualized concerns, health literacy, prior knowledge, digital habits, and localized interactions that conventional segmentation structures may only partially capture.

As discussed by Jones et al. (2024), this expansion of digital information has significantly reshaped the traditional dynamics of the healthcare information marketplace. Stakeholders, including both specialized medical practitioners and increasingly autonomous patient consumers, are no longer exclusively dependent on top-down corporate broadcasting. They now operate within a more plural information environment and may expect relevant, value-driven scientific content capable of addressing specific clinical or informational queries rather than generic promotional appeals that prioritize product positioning over identifiable utility.

The modern medical audience consequently has access to a broader range of tools for evaluating information quality and source credibility, meaning that corporate messaging increasingly benefits from moving beyond generic promotional vocabulary toward transparent scientific evidence, explicit safety information, and demonstrable informational value (Jones et al., 2024). When a pharmaceutical enterprise fails to deliver communication that corresponds to stakeholder context, its messaging may

be perceived as irrelevant or intrusive corporate noise, potentially weakening engagement between the organization and its digital audiences.

Legacy Broadcast Model:

[Static Demographics] → [Monolithic Messaging] → [Potential
Cognitive Rejection]

Modern Personalization Model:

[Behavioral Touchpoints] → [Context Aware Content] → [Potentially
More Relevant Engagement]

The growing interest in moving beyond exclusively mass-oriented communication models is further reinforced by the broader phenomenon of content fatigue, in which increasing volumes of digital information compete for limited human attention (Sabapathy, 2025). In a digital ecosystem saturated with corporate newsletters, banner advertisements, social media outreach, professional content, and user-generated information, stakeholders may experience substantial cognitive saturation and selective attention.

This fatigue can contribute to behavioral responses such as digital blindness, disengagement, or active avoidance, particularly when communication is perceived as irrelevant, repetitive, or overly intrusive (Verma et al., 2024). For the pharmaceutical industry, the commercial risk associated with generic and poorly contextualized messaging therefore extends beyond suboptimal returns on marketing investment. Repeated exposure to low-value communication may also reduce perceived relevance, encourage users to block or ignore particular communication channels, and decrease the likelihood that subsequent information receives meaningful attention.

At the same time, this does not imply that standardized communication has become obsolete or that all interactions should be individually personalized. Standardized communication remains appropriate in contexts where consistency, broad public-health information, or universal safety messaging is required. The strategic challenge lies instead in determining when greater contextual adaptation creates genuine stakeholder value and when increased personalization may introduce unnecessary complexity, privacy concerns, or perceptions of intrusiveness.

2.1.2 Precision Profiling and the Next Best Action

A principal technological catalyst driving the reconfiguration of pharmaceutical digital outreach is the transition from predominantly descriptive, historical data analysis toward increasingly predictive and near-real-time modeling. This operational transformation is not a sudden or detached phenomenon within digital marketing, but rather a process rooted in the historical evolution of prescription tracking architectures and commercial medical data analysis. As critically examined by Fugh-Berman (2008), commercial Health Information Organizations have long operated as data intermediaries within the healthcare sector, aggregating commercial pharmacy fulfillment records and other forms of healthcare-related information to support increasingly granular analysis of physician prescribing patterns.

While this legacy data analysis allowed commercial departments to evaluate past behavioral patterns and subsequently adjust regional sales-representative strategies, contemporary Artificial Intelligence expands this baseline practice by reducing the temporal distance between data collection, interpretation, and strategic response. Modern digital strategy frameworks increasingly transition from retrospective data observation toward predictive behavioral modeling supported by machine-learning architectures (Sabapathy, 2025). Instead of relying exclusively on static repositories or

periodic reporting, contemporary systems may incorporate continuously updated data streams derived from multiple authorized enterprise interactions.

These AI-driven platforms may integrate data from corporate Customer Relationship Management systems, virtual event attendance, medical website navigation, engagement with scientific newsletters, and professional downloads or educational resources, depending on the organizational context and applicable data-governance requirements. By processing heterogeneous behavioral signals, predictive systems can identify patterns that may indicate changing clinical interests or informational requirements, contributing to a more dynamic representation of stakeholder interaction than conventional static categorization models.

Legacy Descriptive Profiling:

[Historical Pharmacy Data] → [Static Repositories] → [Delayed
Strategic Adjustment]

Modern Predictive Modeling:

[Real Time Interaction Signals] → [Predictive AI Engine] → [Next
Best Action Decision Support]

Specifically, these predictive systems can support the Next Best Action framework, which functions as an algorithmic orchestration mechanism within omnichannel medical communication. The Next Best Action infrastructure moves beyond simple automated scheduling by assisting organizations in evaluating several operational variables, including the appropriate timing of outreach, the most suitable authorized communication channel, and the type of scientific content that may correspond most closely to the stakeholder's current informational context (Sabapathy, 2025).

Instead of deploying uncoordinated and generic promotional broadcast messages, the system can assist in identifying an appropriate moment for communication, for example following engagement with a professional resource concerning a specific therapeutic topic, and subsequently prioritize a suitable distribution channel, such as a secure portal message or personalized professional resource. The engine may also support the selection of a relevant scientific asset, such as a peer-reviewed comparative study, localized epidemiological data, or an invitation to a specialized scientific webinar. Importantly, these recommendations should be understood as probabilistic decision support rather than as deterministic predictions of prescribing behavior or stakeholder intention.

Comparable forms of personalization are increasingly being explored within patient-directed communication and support environments, potentially transforming the nature of therapeutic and educational services. By analyzing authorized interactions within digital health portals, patient-support programs, or monitoring applications, systems can contribute to the construction of more adaptive educational pathways that move beyond generic information delivery. Such personalization may respond to characteristics including expressed preferences, interaction history, stage within a support journey, or variations in health literacy, where appropriate and legally permissible.

As discussed by Verma et al. (2024), contemporary personalization systems can potentially use individual interaction patterns to support the identification of informational gaps or changes in engagement. However, such signals should not be interpreted as reliable indicators of hidden psychological states, treatment anxiety, or future non-compliance without additional evidence. Clinical adherence and patient behavior are influenced by multiple psychological, social, medical, and economic factors that cannot necessarily be inferred from digital activity alone.

The strategic value of predictive systems therefore lies primarily in their ability to support the identification of potentially relevant informational needs and to help organizations prioritize appropriate communication, rather than independently diagnosing stakeholder motivations or predicting future therapeutic behavior with certainty. This distinction is particularly important in pharmaceutical communication, where increasing predictive capability must remain proportional to the sensitivity of the underlying data and the consequences of automated inference.

2.1.3 The Precision Ecosystem and the Autonomy Risk

The computational realization of the Next Best Action framework does not operate within an information vacuum; rather, it forms part of a broader, continuously evolving precision ecosystem (Sabapathy, 2025). This digital infrastructure can achieve substantial analytical depth through the synthesis of heterogeneous data sources, potentially combining institutional Real World Evidence with professional engagement information and, where appropriate and legally permissible, insights derived from public digital conversations or social listening (Verma et al., 2024). Real World Evidence can provide longitudinal clinical insights concerning treatment use and outcomes across diverse populations, while social listening may help organizations identify emerging topics, concerns, misconceptions, or broader health-related discussions within public digital environments (Jones et al., 2024).

By integrating different information sources, the precision ecosystem can assist pharmaceutical organizations in understanding the relationship between clinical developments and stakeholder perception (Sabapathy, 2025). However, these data sources differ substantially in methodological quality, purpose, evidentiary status, and regulatory sensitivity. Their integration should therefore not be interpreted as the construction of a complete or objective representation of individual stakeholder reality. Increasing analytical granularity also introduces a structural tension that extends beyond marketing efficiency, raising questions concerning user autonomy, informational diversity, and the acceptable boundaries of predictive inference.

As life sciences corporations adopt increasingly sophisticated personalization systems, algorithms optimized to increase informational relevance may also create a risk of narrowing the range of content made salient to individual stakeholders (Jones et al., 2024). When predictive systems prioritize communication primarily on the basis of historical interaction, previous clinical interests, or past prescribing patterns, they may reinforce existing informational preferences rather than consistently exposing users to a broader range of relevant scientific perspectives (Verma et al., 2024).

If an Artificial Intelligence system identifies that a specific physician repeatedly engages with evidence supporting a particular therapeutic approach, the system may reasonably prioritize additional material within that area. However, if relevance optimization consistently suppresses alternative or clinically significant information because it is predicted to generate lower engagement, the resulting communication environment may contribute to a form of epistemic narrowing. The risk therefore lies not in personalization itself, but in personalization systems that prioritize behavioral optimization without adequate safeguards for informational diversity and scientific balance (Sabapathy, 2025).

Precision Data Synthesis:
[Real World Evidence] + [Social Listening] → [Predictive
Optimization Engine]

Ethical and Operational Tension Risks:
[Hyper Personalization Target] → Risk Vector → [Potential
Epistemic Narrowing]
[Algorithmic Profiling] → Risk Vector → [Reduced Perceived
Autonomy]

Furthermore, the increasing precision enabled by contemporary Artificial Intelligence architectures raises an important ethical dilemma concerning the relationship between helpful personalization and potentially exploitative psychological profiling. As discussed by Amini Khiabani and Rezaeian (2026), advanced behavioral analytics can create conditions in which organizations possess increasing capacity to infer preferences, interests, and potentially sensitive characteristics from digital behavior. The existence of such computational capability does not automatically establish the ethical legitimacy of using every available inference for persuasive purposes.

When enterprise platforms attempt to infer emotional vulnerability, health-related anxiety, professional stress, or other sensitive characteristics in order to optimize communication, the distinction between educational support and persuasive exploitation becomes particularly important. Even technically accurate inferences may be ethically problematic when their primary purpose is to capitalize on moments of vulnerability rather than to provide proportionate informational utility.

Optimizing promotional layouts or disease-awareness communication specifically to exploit moments of emotional vulnerability or psychological fatigue would therefore raise significant ethical concerns, potentially transforming data-driven agility from a mechanism of contextual support into a form of behavioral influence that may undermine the independent judgment of consumers or medical professionals (Amini Khiabani & Rezaeian, 2026).

This thesis contends that while the technological capability to construct increasingly precise communication ecosystems represents a significant computational development, their long-term strategic validity and ethical legitimacy depend on establishing a transparent and verifiable equilibrium between analytical capability and stakeholder autonomy. The operational success of digital pharmaceutical marketing cannot therefore be measured solely through algorithmic efficiency, click frequencies, engagement levels, or behavioral conversion metrics. It must also take into account scientific balance, informational integrity, proportionality of data use, and the stakeholder's capacity to understand and retain meaningful agency within the interaction.

Ultimately, this ethical tension contributes to the foundational justification for the Data + Empathy framework developed throughout this research. However, the framework should not be interpreted as requiring Artificial Intelligence to remain permanently confined to an invisible background infrastructure. Instead, later chapters refine this initial distinction through the principle of risk-proportionate functional asymmetry.

Within this approach, computational systems can assume greater autonomy in standardized, informational, administrative, or comparatively low-risk functions where analytical speed, pattern recognition, retrieval, and scalable processing create clear stakeholder value. Conversely, human oversight becomes increasingly important where technological systems approach clinical interpretation, sensitive personalization, emotional vulnerability, uncertain evidence, or consequential decision-making.

The central issue is therefore not the absolute visibility or invisibility of Artificial Intelligence, but the appropriate allocation of technological authority. Computational precision should be employed where it contributes meaningfully to stakeholder utility, while human judgment retains particular importance where contextual interpretation, ethical responsibility, interpersonal sensitivity, and accountability become central to the quality of pharmaceutical communication.

2.2 Digital Skepticism and Ad Avoidance Behaviors: The Role of Generation Z

The pharmaceutical industry's structural transition toward algorithmic precision, predictive data architectures, and hyper personalized omnichannel distribution encounters a distinctive behavioral challenge when intersecting with younger, digitally experienced consumer cohorts such as Generation Z. Members of this demographic have developed within an environment extensively shaped by digital platforms, mobile computing, automated recommendation systems, targeted advertising, and increasingly sophisticated forms of algorithmic mediation. However, familiarity with these technologies should not be interpreted as equivalent to unconditional acceptance of the commercial architectures that operate through them.

This apparent contradiction is particularly relevant to contemporary pharmaceutical communication. Generation Z combines extensive exposure to digital technologies with growing familiarity with targeted advertising, platform-based recommendation systems, privacy debates, and online commercial persuasion. Research on digital advertising has documented significant forms of advertising avoidance and resistance to intrusive promotional communication, while research concerning younger populations indicates that digital and social platforms increasingly influence how health-related information is discovered, evaluated, and shared (Redondo & Aznar, 2018; Lim et al., 2022).

This behavioral resistance should not be interpreted as a universal psychological characteristic or as evidence that an entire generation possesses an identical distrust of corporate communication. Rather, it indicates that younger digital users operate within communication environments in which institutional information competes with peer-generated content, social media creators, independent online resources, recommendation systems, and increasingly sophisticated commercial formats. These conditions can influence the criteria through which credibility, relevance, and persuasive intent are evaluated.

Generation Z has also matured during a period marked by increasing public discussion concerning data privacy, digital surveillance, misinformation, and the commercial use of behavioral information. At the same time, young adults are frequently exposed to health-related information of highly variable scientific quality across digital environments (Mansfield et al., 2025; Lim et al., 2022). The resulting communication context can encourage greater scrutiny of how information is presented, who is responsible for it, and what commercial objectives may underlie a particular interaction.

Consequently, conventional pharmaceutical advertising strategies cannot assume that increasing mathematical precision or targeting granularity will automatically increase acceptance. Hyper personalized communication may create value when it corresponds to an identifiable informational need, but it may also generate discomfort when recipients perceive that sensitive behavioral information has been used in ways that are opaque, disproportionate, or unexpectedly intimate.

This tension is particularly important for automated marketing architectures such as the Next Best Action framework. When Artificial Intelligence systems optimize outreach through behavioral signals, the resulting communication may increase contextual relevance while simultaneously making the underlying use of personal data more visible to the recipient. In some circumstances, highly

specific targeting may therefore prompt questions concerning how the organization obtained the relevant information and why a particular message was directed toward that individual.

Rather than inevitably producing cognitive reactance, hyper personalization should therefore be understood as creating a conditional communication environment. Its effects depend on factors including perceived utility, transparency, data sensitivity, timing, commercial intent, and the recipient's understanding of the technological processes involved. A personalized message may be interpreted as useful support in one context and as intrusive behavioral targeting in another.

This generational friction consequently creates an important limitation for purely optimization-driven approaches to pharmaceutical communication. Technological advancement and algorithmic scalability remain valuable, but they are insufficient on their own to resolve the contemporary trust deficit. Digital precision must operate within communication strategies that make scientific provenance, organizational responsibility, data practices, and stakeholder utility sufficiently visible and understandable (Amini Khiabani & Rezaeian, 2026).

2.2.1 The Algorithmic Intuition and the Trust Deficit

Within this research, the concept of algorithmic intuition is used as an interpretative label for the practical familiarity that digitally experienced users may develop with recurring patterns of online persuasion, sponsored content, recommendation systems, and platform-mediated advertising. The term should not be understood as an innate psychological faculty possessed uniformly by Generation Z. Rather, it describes a form of accumulated digital awareness that can emerge through repeated exposure to commercial and algorithmically curated environments.

Individuals who have spent substantial portions of their lives interacting with digital interfaces may become increasingly familiar with native advertising formats, sponsored posts, recommendation mechanisms, tracking practices, and the stylistic conventions commonly associated with commercial communication. This familiarity can influence how promotional messages are interpreted and may increase sensitivity to persuasive intent, particularly when communication appears excessively polished, highly targeted, or disconnected from recognizable informational utility.

The resulting critical filter is operationally visible in different forms of ad avoidance. Digital users can reduce commercial exposure through both technical and behavioral mechanisms, including ad blockers, browser-level tracking protections, privacy settings, selective attention, rapid scrolling, or the deliberate avoidance of recognizable advertising formats. Research by Redondo and Aznar (2018), for example, demonstrates the relevance of attitudes and motivations underlying online ad-blocking behavior, although such evidence should not be interpreted as establishing a universal or specifically pharmaceutical pattern among Generation Z.

These practices nevertheless have important implications for digital strategy. Traditional pharmaceutical communication touchpoints increasingly operate within environments in which audiences possess greater control over whether promotional content is displayed, attended to, or ignored. Even when a communication asset reaches the intended user, visual and rhetorical characteristics associated with conventional advertising may affect whether the message receives further attention.

This technical and behavioral resistance develops alongside a broader decentralization in the sources through which younger audiences access health-related information. Qualitative research indicates that young adults may use social media as one component of their health-information environment, alongside healthcare professionals, institutional resources, search engines, peer networks, and other

digital sources (Lim et al., 2022). This does not imply that scientific or institutional authority has been systematically replaced by influencer credibility. Rather, health-information evaluation increasingly occurs across a heterogeneous set of sources that provide different forms of expertise, accessibility, experiential knowledge, and perceived authenticity.

Within these decentralized digital spaces, personal narratives and peer experiences may provide forms of relevance that institutional communication cannot always reproduce. A user may value a healthcare professional because of clinical expertise while simultaneously valuing a patient community because it provides experiential understanding of daily life with a particular condition. These forms of credibility are not necessarily interchangeable and may coexist within the same information-seeking process.

The resulting trust environment is therefore more complex than a simple transfer of authority from institutions to social media. Formal clinical validation, scientific evidence, professional expertise, perceived authenticity, personal relevance, and relational accessibility may all contribute differently to credibility depending on the informational context (Lim et al., 2022).

From a digital strategy perspective, this behavioral realignment places an important boundary around the value of computational precision. A Next Best Action system may successfully identify a potentially relevant moment for communication and prioritize an appropriate scientific resource. Nevertheless, the effectiveness of the resulting interaction also depends on whether the recipient understands the source of the information, perceives the intervention as proportionate, and recognizes a legitimate informational purpose behind the communication.

A technically accurate recommendation may therefore remain ineffective if its presentation creates perceptions of excessive commercial targeting or fails to communicate why the information is relevant. Conversely, a clearly sourced and contextually appropriate communication may benefit from personalization without requiring the user to accept an opaque or excessively intimate form of behavioral profiling.

The limitation is not that computational systems necessarily treat every stakeholder as a static data target. Rather, the risk emerges when organizational communication privileges predictive optimization while neglecting the interpretative, relational, and contextual conditions through which stakeholders evaluate the interaction.

This distinction contributes to the broader Data + Empathy framework developed throughout the thesis. Precision should not be subordinated to a rigid requirement for exclusively human delivery, but it should be integrated with communication design that preserves transparency, contextual appropriateness, and stakeholder agency. Human-centeredness therefore concerns the conditions under which data-informed communication is deployed rather than simply the presence or absence of an automated interface.

2.2.2 The Platform Paradox: Visual Validation vs. Scientific Evidence

The cognitive evaluation and subsequent verification of digital health information can be influenced by the technological affordances and structural architectures of the platforms within which content is encountered. Within contemporary media ecology, trustworthiness is not evaluated exclusively through the intrinsic scientific quality of information. The context in which evidence appears, the perceived identity of the source, the format of presentation, and the user's informational objective may all influence how health-related communication is processed (Lim et al., 2022).

This contextual variability creates what this thesis defines as the platform paradox: scientific information may be interpreted differently depending on the communicative environment through which it is encountered. However, platform architecture should not be treated as a deterministic mechanism that automatically determines whether users process information systematically or heuristically.

To interpret this phenomenon, this research draws on the Heuristic-Systematic Model of information processing (Chaiken, 1980). This cognitive framework distinguishes between systematic processing, characterized by greater cognitive effort and detailed consideration of message content, and heuristic processing, in which judgments rely more heavily on accessible cues and cognitive shortcuts. These modes are not mutually exclusive, and individual platforms cannot be assigned permanently to one or the other. User motivation, knowledge, context, information needs, and message characteristics all influence how content is evaluated.

Nevertheless, different platform affordances may create conditions that facilitate particular forms of engagement.

- **YouTube and Higher-Intent Information Processing:** When users access YouTube through an explicit search query concerning a health-related topic, the platform can support comparatively extended engagement with longer-form information. Users may seek detailed explanations, expert commentary, step-by-step educational content, or accounts of personal health experiences (Lim et al., 2022). Long video formats, searchable descriptions, chapters, and links to external resources can provide greater opportunity for the examination of evidence than is typically available in extremely brief content formats. Under these conditions, users may devote greater cognitive effort to evaluating the depth of explanation, source expertise, or presence of supporting information. This does not mean that YouTube inherently produces systematic processing or that the information available on the platform is necessarily scientifically reliable. Rather, its architecture can support more extensive information processing when users approach the platform with sufficiently high informational motivation.
- **Instagram, TikTok, and Rapid Heuristic Evaluation:** Short-form social media architectures such as Instagram and TikTok frequently organize content through continuous algorithmic feeds, brief audiovisual formats, and rapid navigation. Within low-intent browsing contexts, users may have limited time to evaluate each individual health claim and may therefore rely more heavily on immediately accessible cues, including presentation quality, creator familiarity, perceived authenticity, social endorsement, and narrative accessibility. Such cues can function as heuristic signals during rapid information evaluation, but they should not be interpreted as substitutes for scientific validity. Importantly, TikTok and Instagram can also host evidence-based and professionally produced health communication, just as YouTube can host misleading or poorly substantiated material. The relevant distinction therefore concerns the interaction between platform affordances, user motivation, and content characteristics rather than a deterministic separation between scientifically rigorous and aesthetically driven platforms.

YouTube Interaction Framework (Higher Intent Search):

[Clinical Query] → [Long Form Information] → Greater Opportunity for Systematic Processing → Evidence Evaluation

TikTok and Instagram Framework (Rapid Feed Encounter):

[Algorithmic Feed] → [Accessible Visual and Social Cues] → Greater Potential for Heuristic Processing → Rapid Credibility Assessment

This distinction becomes increasingly important with the expansion of Artificial Intelligence-generated content. Generative systems can produce professionally structured text, images, audio, and video capable of reproducing many of the visual and rhetorical conventions traditionally associated with credible digital communication. Consequently, stylistic quality can no longer be assumed to provide reliable information concerning provenance or scientific validity.

AI-generated or AI-modified health misinformation may therefore create additional challenges within communication environments characterized by rapid consumption and limited verification time. A visually persuasive health claim, fabricated correlation, or unsupported wellness recommendation may receive substantial exposure when it is compatible with platform engagement dynamics. However, Generation Z should not be characterized as uniquely susceptible to such content without comparative evidence demonstrating greater vulnerability than other populations.

The more defensible concern is that users across demographic groups may face increasing difficulty evaluating provenance when synthetic content can reproduce familiar visual and communicative markers of authenticity. For younger audiences who frequently encounter health information within social media environments, this challenge becomes particularly relevant because scientifically reliable and unreliable information may coexist within highly similar formats (Mansfield et al., 2025).

The resulting misinformation gap represents an important public-health and corporate communication challenge. Pharmaceutical organizations and regulatory actors frequently operate through formal scientific, legal, and compliance processes, while misleading digital content can circulate rapidly across decentralized platforms. This difference in communication velocity can make immediate institutional response difficult.

It would nevertheless be excessive to assume that pharmaceutical corrections necessarily require weeks or months, that misinformation invariably reaches viral saturation before intervention, or that exposure permanently alters consumer beliefs. The speed and magnitude of misinformation diffusion vary substantially according to platform dynamics, content characteristics, audience response, and the availability of competing information.

The strategic implication is instead that exclusively retrospective correction may be insufficient within high-velocity digital environments. Pharmaceutical organizations, healthcare institutions, platforms, and public authorities may benefit from earlier detection of emerging misinformation patterns, timely evidence-based communication, and communication formats capable of making scientific provenance accessible within the media environments where health-related claims are circulating.

Proactive monitoring should therefore be understood as a complementary communication and governance mechanism rather than as a guarantee that misinformation can be identified or neutralized in real time. The objective is to reduce the temporal distance between the emergence of problematic health narratives and the availability of credible, verifiable information while preserving the scientific and regulatory standards required in healthcare communication (Mansfield et al., 2025).

2.2.3 The Human Element Requirement in the AI Era

A particularly important issue for the theoretical architecture of this thesis concerns the limits of Artificial Intelligence acceptance in sensitive healthcare contexts. A common assumption within

digital strategy is that populations familiar with digital technologies will necessarily prefer increasingly automated and frictionless interactions. However, technological familiarity does not automatically imply willingness to delegate every informational, emotional, or decisional function to Artificial Intelligence.

The distinction becomes particularly relevant in healthcare, where the consequences of communication vary substantially according to context. An automated system used to retrieve general information, explain terminology, or support an administrative function occupies a substantially different role from an Artificial Intelligence system involved in diagnostic interpretation, therapeutic recommendations, emotionally sensitive support, or other high-consequence interactions.

To interpret this distinction, this research integrates the established literature on algorithm aversion. Dietvorst et al. (2015) demonstrated that individuals may become reluctant to rely on algorithms after observing them make errors, even when the algorithm may continue to perform comparatively well. Their research was not specifically conducted within Generation Z populations or healthcare environments and therefore should not be treated as direct evidence of generational resistance to AI-mediated medical care.

Nevertheless, algorithm aversion provides a useful conceptual lens for understanding why technical performance alone may be insufficient to generate sustained confidence in automated systems. In healthcare contexts, perceived stakes, uncertainty, accountability, and the consequences of error can influence whether users consider technological delegation appropriate.

When facing an emotionally sensitive or clinically consequential situation, an individual may require forms of contextual interpretation and interpersonal reassurance that differ substantially from the informational utility expected from a conventional automated tool. This does not mean that every patient necessarily requires human interaction at every stage of a healthcare journey, nor that Artificial Intelligence cannot contribute directly to communication. Rather, the appropriate balance between automated and human involvement is likely to vary with the sensitivity and complexity of the function performed.

Algorithmic Support Infrastructure (Data Synthesis):

[Real World Evidence] + [Pattern Recognition] → High Speed Processing

Human-Governed Interaction (Contextual Responsibility):

[Contextual Interpretation] + [Accountability] → Human-Centered Communication

A related question concerns the use of synthetic empathy. Contemporary language models can generate linguistically empathetic responses and reproduce many verbal conventions associated with reassurance, support, and interpersonal warmth. However, the production of empathetic language should be distinguished from the human capacity to assume moral responsibility, interpret complex personal context, and participate in a genuinely reciprocal relationship.

The concept of the uncanny valley of mind described by Stein and Ohler (2017) provides a useful, although limited, theoretical analogy. Their research suggests that discomfort can emerge when artificial entities approach qualities associated with human mental capacities in ways that remain perceptibly nonhuman. The framework was not developed specifically to evaluate Large Language Models, healthcare communication, or Generation Z and should therefore not be interpreted as evidence that AI-generated empathy necessarily produces psychological repulsion.

Applied cautiously to pharmaceutical communication, the theory nevertheless raises a relevant question: whether highly anthropomorphic or emotionally expressive automated systems may create discomfort when stakeholders perceive an inconsistency between simulated emotional language and the actual nature or authority of the system.

Artificial expressions of concern can be experienced differently according to user expectations, disclosure, context, and task. Some individuals may consider automated reassurance useful, others may perceive it as artificial, and others may remain indifferent. The effect cannot therefore be assumed in advance.

This distinction is consistent with the exploratory field research presented later in Chapter 5. The results do not demonstrate universal rejection of synthetic empathy, but they do indicate substantial ambivalence toward explicitly empathetic AI-mediated health communication and stronger preferences for human responsibility as contextual and clinical sensitivity increases.

Consequently, the relationship between human involvement and Artificial Intelligence should not be framed through a rigid rule requiring all AI systems to remain invisible. The earlier backend/front-end distinction developed in this thesis is more appropriately refined through a principle of risk-proportionate functional asymmetry.

Within this framework, computational intelligence can absorb high-volume functions involving data synthesis, pattern recognition, information retrieval, regulatory support, and other standardized processes in which automation provides demonstrable utility (Sabapathy, 2025). AI may also interact directly with stakeholders in lower-risk contexts when the function is sufficiently bounded, transparent, and appropriate.

Conversely, the requirement for meaningful human involvement increases as the interaction approaches higher levels of clinical uncertainty, emotional vulnerability, sensitive personalization, or consequential decision-making. In these circumstances, human professionals provide not merely an emotionally warmer interface but identifiable accountability, contextual interpretation, professional judgment, and the capacity to respond to circumstances that cannot be reduced to standardized computational outputs (Verma et al., 2024).

The Data + Empathy framework therefore does not position human communication as a cosmetic layer placed over automated decision-making. Nor does it assume that human interaction automatically creates trust. Instead, it argues that technological and human responsibilities should be allocated according to the nature and consequences of the task.

For younger digital audiences, this approach is particularly relevant because extensive familiarity with Artificial Intelligence can coexist with clear expectations concerning transparency, scientific provenance, privacy, and human accountability. Rebuilding or maintaining digital trust consequently requires more than abandoning corporate polish or attempting to imitate interpersonal vulnerability through automated language.

The strategic objective is to use analytical precision to improve the relevance and accessibility of healthcare communication while ensuring that technological authority remains proportionate to contextual risk. Human-centered pharmaceutical communication therefore emerges not from the categorical exclusion of Artificial Intelligence from visible interaction, but from the deliberate assignment of AI and human responsibility according to where each can create legitimate stakeholder value.

2.3 International Regulatory Landscapes and the Regulatory Lag

The accelerating development of digital health communication architectures has generated an increasingly complex relationship between technological innovation and formal legal oversight, a dynamic that can be interpreted through the concept of regulatory lag. To examine this institutional friction, this research incorporates the theoretical framework of the pacing problem developed within technology-governance literature (Marchant, Allenby, & Herkert, 2011). The pacing problem describes the difficulty experienced by legal and regulatory institutions when technological systems, commercial practices, and communication infrastructures evolve more rapidly than the legislative and administrative processes through which regulatory expectations are formally revised.

This temporal asymmetry should not be interpreted as evidence that digital pharmaceutical communication operates in the absence of law or that public authorities are structurally incapable of responding to technological development. Existing legal principles continue to govern emerging communication practices, even when they were established before the appearance of specific technologies such as generative Artificial Intelligence, algorithmic recommendation systems, or real-time behavioral personalization. However, the application of established legal principles to new technological configurations can generate interpretative uncertainty, particularly where regulatory categories were originally developed around more centralized and clearly identifiable communication channels.

Formal legislation, administrative guidance, judicial interpretation, and enforcement procedures also operate through institutional processes that generally require more time than commercial software development or platform-level technological experimentation. Digital platforms and AI-mediated communication systems can be modified continuously, whereas legislative reform, regulatory consultation, and judicial clarification typically involve substantially longer procedural cycles (Marchant et al., 2011). The resulting friction does not create a permanent legal vacuum, but it can produce temporary areas of uncertainty in which organizations must interpret existing rules in relation to technological practices that those rules did not explicitly anticipate.

Temporal Misalignment Matrix:

[Technological and Platform Innovation] → Rapid Iteration →
Continuous / Short-Cycle Development
vs.
[Formal Regulatory Adaptation] → Institutional Procedure →
Comparatively Slower Revision and Enforcement

This tension becomes particularly visible when comparing different international approaches to pharmaceutical promotion. Among the most frequently discussed examples is the divergence between the United States and the European Union concerning advertising of prescription medicinal products to the general public. The two systems reflect different legal and regulatory structures, producing substantially different conditions for direct pharmaceutical communication.

Within the United States, direct-to-consumer advertising of prescription medicines is permitted subject to federal statutory and regulatory requirements. Pharmaceutical promotional communication must remain consistent with applicable regulatory standards concerning truthfulness, substantiation, disclosure, and the presentation of risk and benefit information. FDA prescription-drug advertising requirements, including those reflected in 21 CFR § 202.1, establish important constraints concerning misleading representations and the presentation of information regarding product risks and effectiveness. Constitutional protections applicable to commercial speech also influence the broader

legal environment, but they do not provide pharmaceutical organizations with unrestricted freedom to disseminate medical claims or disregard federal drug-advertising requirements.

Consequently, the American regulatory environment can support comparatively direct communication between pharmaceutical manufacturers and consumer audiences while continuing to subject prescription-drug advertising to substantial regulatory oversight. This configuration provides organizations with broader opportunities for consumer-facing communication than are generally available for prescription medicines within the European Union, but those opportunities remain conditional upon compliance with applicable federal standards.

The European framework adopts a different approach. Directive 2001/83/EC requires Member States to prohibit advertising to the general public of medicinal products that are available only on medical prescription (European Parliament & Council, 2001, Art. 88). This restriction concerns advertising of prescription-only medicinal products rather than all pharmaceutical or health-related communication. The Directive also distinguishes promotional communication from certain forms of non-promotional information. Article 86, for example, excludes information relating to human health or diseases from the advertising framework where there is no reference, even indirectly, to a medicinal product (European Parliament & Council, 2001, Art. 86).

This distinction is strategically important because it prevents the European framework from being reduced to a simple prohibition on communication between pharmaceutical organizations and the public. Life sciences companies may participate in disease-awareness, public-health, corporate, and educational communication, provided that the resulting activity remains consistent with the applicable legal distinction between information and medicinal-product promotion. The regulatory challenge therefore concerns not merely whether communication reaches a public audience, but also its content, purpose, product references, presentation, and broader commercial context.

The resulting international divergence creates a fragmented environment for pharmaceutical digital strategy. A communication asset that may be legally deployable within one jurisdiction may require substantial modification, audience restriction, or complete exclusion in another. Contemporary organizations operating across multiple territories must therefore develop governance systems capable of adapting content, targeting parameters, data-processing practices, and promotional classifications according to the applicable legal environment rather than assuming that a single global digital strategy can be deployed uniformly.

2.3.1 The Jurisdictional Gray Areas of Social Media

The intrinsically cross-border architecture of major digital platforms introduces additional complexity into pharmaceutical and health-related communication. Social-media content can be produced in one jurisdiction, hosted through infrastructure distributed across several territories, promoted by intermediaries operating elsewhere, and accessed by audiences subject to different national legal requirements. This distribution structure can complicate the identification of the relevant regulatory authority and the application of geographically bounded advertising rules to communication that circulates through global networks.

Within this environment, the concept of regulatory arbitrage provides a useful analytical lens, although it should not be assumed that all cross-border digital communication is intentionally designed to exploit legal discrepancies. Regulatory arbitrage may arise when organizations or other commercial actors structure activities in ways that take advantage of differences between national regulatory regimes. Digital platforms can facilitate such practices because the geographical origin of

content, the location of the sponsoring entity, and the jurisdiction of the receiving audience may not coincide (Verma et al., 2024).

The resulting challenge becomes more pronounced when communication is produced not directly by pharmaceutical corporations but by influencers, affiliates, independent creators, telehealth providers, wellness entrepreneurs, or other intermediaries. In such circumstances, the legal relationship between the product manufacturer, the communicator, the digital platform, and the target audience may be less immediately visible than in conventional corporate advertising.

Digital platforms therefore create governance complexity not because national sovereignty disappears, but because enforcement authority remains geographically and institutionally distributed while content circulation is comparatively fluid. Public authorities retain regulatory and enforcement powers within their respective jurisdictions, yet determining responsibility, collecting evidence, identifying commercial relationships, and coordinating responses may become more difficult where communication networks span multiple countries and involve numerous independent actors.

Several jurisdictions have responded by developing more explicit requirements concerning digital endorsements and health-related influencer communication. India provides a relevant example. Consumer-protection authorities have issued guidelines concerning advertising disclosures by celebrities, influencers, and virtual influencers and have subsequently introduced additional guidance for health and wellness communication. These provisions emphasize the need for clear disclosure and require individuals presenting themselves as certified medical or health experts to identify their relevant qualifications when making health-related claims or promoting health products and services (Government of India, Department of Consumer Affairs, 2023).

This regulatory response illustrates an important development in digital governance. Rather than treating influencers exclusively as informal individual speakers outside conventional commercial regulation, authorities increasingly recognize that creator-mediated communication can form part of an advertising and endorsement ecosystem when material connections, promotional relationships, or commercial claims are involved.

Decentralized Influence Loop:

[Global Platform Infrastructure] → [Creator / Affiliate Health Content] → [Cross-Jurisdictional Distribution]



Enforcement Scale Challenge:

[National Regulatory Authorities] ← Jurisdictional and Monitoring Constraints ← [High-Volume Decentralized Feeds]

The presence of formal guidance nevertheless does not ensure complete compliance across decentralized digital environments. Independent creators can publish large volumes of health-related information at comparatively high speed, and the scientific quality of that information may vary considerably (Lim et al., 2022). Regulatory authorities and digital platforms must therefore evaluate communication environments in which promotional, educational, testimonial, and personal content frequently coexist.

This scale difference creates practical enforcement challenges. Traditional pharmaceutical advertising generally involves identifiable corporate sponsors and relatively formalized review structures. Social-media communication, by contrast, can be generated, reproduced, modified, and redistributed by large numbers of independent users. Consequently, public authorities cannot realistically pre-review every individual health-related post before publication.

However, this limitation should not be characterized as a total failure of state regulation. Regulatory authorities retain the ability to investigate, issue warnings, impose sanctions where legally applicable, require corrective measures, and establish clearer standards for commercial actors. The principal challenge concerns the difference between the scale and velocity of decentralized communication and the necessarily selective nature of institutional enforcement.

For pharmaceutical organizations, this environment strengthens the importance of internal governance. Companies operating through affiliates, agencies, creators, patient-support partners, or platform-based campaigns require clearly defined processes concerning sponsorship disclosure, approved claims, scientific substantiation, territorial restrictions, and responsibility for third-party communication. Regulatory compliance therefore increasingly extends beyond reviewing the content produced directly by the organization and toward governing the broader ecosystem through which commercially connected communication may circulate.

2.3.2 Data Privacy vs. Marketing Liberty: The Battle over Tracking

A second major source of regulatory friction concerns the collection, processing, commercialization, and strategic use of prescriber and patient-related data. Data-driven pharmaceutical communication requires organizations to navigate a complex intersection between commercial analytics, privacy expectations, freedom of expression, data-protection law, and the heightened sensitivity associated with health-related information.

In the United States, a significant legal controversy concerning pharmaceutical data use emerged around the commercialization of prescriber-identifying information. Historically, commercial data intermediaries aggregated information concerning prescribing patterns and supplied analytical products to pharmaceutical organizations, allowing commercial teams to tailor detailing strategies more precisely to individual physicians (Fugh-Berman, 2008).

This practice became the subject of state-level regulatory intervention, including legislation adopted by Vermont restricting the sale, disclosure, and use of prescriber-identifying information for pharmaceutical marketing. The dispute ultimately reached the Supreme Court of the United States in *Sorrell v. IMS Health Inc.*, 564 U.S. 552 (2011).

The Court invalidated the Vermont restriction, concluding that the law imposed content- and speaker-based burdens on protected expression and therefore could not survive the applicable First Amendment scrutiny. The decision recognized that prescriber-identifying information could be used by pharmaceutical sales representatives to refine commercial communication and held that Vermont could not selectively restrict that use in the manner established by the statute.

Importantly, *Sorrell* should not be interpreted as creating a permanent regulatory vacuum for pharmaceutical data processing or as granting commercial organizations absolute constitutional protection for every form of health-related data use. The decision concerned a specific Vermont statute and a specific form of prescriber-identifying information. Other privacy, consumer-protection, health-information, contractual, and sector-specific rules can continue to regulate data collection and processing according to their respective legal foundations.

United States Paradigm (Commercial Speech and Prescriber Data):

[Prescriber-Identifying Information] → [Commercial Analytics] →
[Targeted Professional Communication]



[Constitutional and Statutory Constraints]

European Union Paradigm (Data Protection and Promotional
Restrictions):

[Personal / Health-Related Data] → [GDPR Processing Requirements]
→ [Purpose-Limited and Context-Dependent Use]

The commercial use of prescription-related information nevertheless demonstrates how data can increase the precision of pharmaceutical strategy. Historically, data intermediaries enabled organizations to evaluate physician prescribing behavior and allocate detailing resources more selectively (Fugh-Berman, 2008). Contemporary machine-learning systems can extend these analytical practices by combining multiple authorized data sources and identifying patterns across professional interactions, content engagement, and other digital touchpoints.

However, the availability of increasingly sophisticated analytics should not be interpreted as permitting organizations to infer with certainty a clinician's "diagnostic vulnerabilities," future prescribing decisions, or susceptibility to specific promotional techniques. Predictive systems operate probabilistically, and the legitimacy of their use depends on the provenance of the underlying information, applicable legal conditions, organizational purpose, and consequences of the resulting profiling.

The European Union provides a substantially different data-governance environment through Regulation (EU) 2016/679, the General Data Protection Regulation. The GDPR establishes principles including lawfulness, fairness, transparency, purpose limitation, and data minimization, requiring personal data to be adequate, relevant, and limited to what is necessary for the specified processing purpose (Regulation (EU) 2016/679, Art. 5). It also affords particular protection to data concerning health under its framework for special categories of personal data.

These provisions do not categorically prohibit predictive analysis or digital personalization. Rather, they impose legal conditions on how personal data may be collected and processed. Pharmaceutical organizations operating within the European Union must therefore evaluate issues including lawful basis, transparency, purpose compatibility, data minimization, data sensitivity, and the rights of data subjects when designing data-driven communication systems.

The interaction between data protection and pharmaceutical advertising restrictions becomes particularly relevant in disease-awareness communication. Under Directive 2001/83/EC, information concerning human health or diseases falls outside the Directive's definition of medicinal-product advertising when it contains no reference, even indirectly, to a medicinal product (European Parliament & Council, 2001, Art. 86). Consequently, unbranded disease-awareness initiatives can perform legitimate educational and public-health functions.

However, the absence of a product name does not automatically establish that every disease-awareness campaign is non-promotional. The legal characterization depends on the communication as a whole, including whether it indirectly refers to an identifiable medicinal product or operates in a

manner designed to promote its prescription, supply, sale, or consumption. The distinction between legitimate disease education and indirect product promotion is therefore contextual rather than purely lexical.

This distinction substantially changes the strategic interpretation of unbranded communication. Disease-awareness initiatives should not be described as mechanisms through which pharmaceutical companies systematically “circumvent” European advertising prohibitions. In many circumstances they constitute lawful health information. At the same time, where the design, targeting, linkage structure, or surrounding commercial context creates an indirect reference to a medicinal product, regulatory scrutiny may become appropriate.

The strategic challenge for pharmaceutical organizations is consequently not to neutralize privacy or promotional restrictions through semantic restructuring, but to construct communication architectures in which data use, disease education, and commercial objectives remain appropriately separated and legally justified. This requires organizations to evaluate not only what communication is technically possible, but also whether the data-processing and promotional logic supporting that communication remains proportionate, transparent, and compatible with the applicable regulatory framework.

2.3.3 Borderline Contexts and the Failure of Self-Regulation

The difficulty of applying traditional regulatory categories to rapidly evolving digital communication becomes particularly visible within borderline contexts. These environments include products and claims that may sit close to the legal boundaries separating medicinal products, dietary supplements, foods, cosmetics, wellness services, and other regulated consumer-health categories (Catalani et al., 2021). The classification of such products is legally significant because different categories can be subject to substantially different requirements concerning authorization, permissible claims, advertising, evidence, and post-market oversight.

Digital platforms amplify the complexity of these distinctions because users frequently encounter content from multiple regulatory categories within the same feed. A prescription-medicine discussion, a supplement endorsement, a cosmetic claim, a wellness routine, and a personal health testimonial may appear consecutively and be presented through similar audiovisual formats. From the user’s perspective, the regulatory status and evidentiary requirements applicable to each item may not be immediately apparent.

Periods of heightened public attention toward health, physical performance, weight management, immunity, or appearance can further intensify this problem. Commercial actors may respond rapidly to emerging consumer interest by promoting products through creator partnerships, testimonial content, algorithmically distributed advertising, and other digital formats. Where health claims are exaggerated, insufficiently substantiated, or presented in a way that blurs the distinction between wellness support and disease treatment, regulators may face the challenge of evaluating large volumes of content after it has already entered circulation.

This does not mean that commercial actors systematically exploit legal loopholes or that borderline classifications inherently exist to facilitate regulatory avoidance. Product classification depends on applicable law, composition, intended use, claims, and presentation. The governance concern arises when ambiguous or misleading communication takes advantage of consumer uncertainty concerning those distinctions.

[AI-Generated / Influencer Health Content] → Rapid Digital
Diffusion → Potentially Broad Exposure
vs.

[Platform Self-Regulation] → Automated and Human Moderation →
Selective Intervention

[Traditional State Enforcement] → Investigation and Administrative
Review → Subsequent Intervention

The communication challenge becomes particularly relevant among younger audiences because social media forms part of the wider environment through which young adults encounter health information, personal experiences, wellness narratives, and commercial recommendations (Lim et al., 2022). However, this should not be interpreted as evidence that Generation Z uniformly prioritizes influencer testimony over scientific information or that peer-generated content automatically replaces professional medical authority.

Rather, platform environments create situations in which different credibility cues coexist. Scientific evidence, professional credentials, personal testimony, social endorsement, visual quality, creator familiarity, and commercial sponsorship may all contribute differently to how individual users interpret content. The resulting information environment makes transparent disclosure and source evaluation increasingly important.

Platform-based self-regulation plays an important role in this ecosystem. Digital platforms can establish advertising policies, remove prohibited content, restrict specific commercial categories, require sponsorship disclosures, and deploy automated detection systems at a scale unavailable to individual national authorities. These mechanisms represent a meaningful component of contemporary digital governance.

At the same time, platform moderation has structural limitations. Automated systems can process enormous quantities of content, but nuanced health claims often require contextual, scientific, and legal interpretation. A claim may be misleading without containing a prohibited keyword, and an apparently personal recommendation may involve an undisclosed commercial relationship. Similarly, the scientific appropriateness of a health statement may depend on details that cannot be reliably determined through surface-level content classification alone.

Platform self-regulation should therefore not be understood either as entirely ineffective or as a substitute for public regulation. It forms one layer within a broader governance environment that also includes statutory authorities, consumer-protection law, advertising standards, professional regulation, organizational compliance systems, and user-level information literacy.

The principal risk emerges when responsibility for health-information integrity is effectively fragmented across these actors without sufficiently clear accountability. A platform may moderate according to its own policies, a creator may consider themselves responsible only for personal expression, a commercial sponsor may rely on contractual disclosure obligations, and public authorities may intervene only once specific legal thresholds are met. In such circumstances, problematic communication can remain accessible even though no single governance mechanism is capable of addressing every stage of its production and circulation.

This structural limitation strengthens the argument for proactive organizational governance within pharmaceutical and life sciences communication. Organizations cannot assume that platform moderation or retrospective enforcement will independently identify every inappropriate, misleading, or contextually problematic communication before stakeholder exposure. Internal scientific review,

sponsorship governance, data-provenance controls, compliance pre-screening, escalation procedures, and proportionate Human-in-the-Loop oversight therefore remain important components of responsible digital strategy.

Within the Data + Empathy framework developed in this thesis, the appropriate response is not the establishment of a universally human front-end designed to compensate for supposedly unreliable digital platforms. Instead, regulatory and ethical safeguards should be allocated according to the function, sensitivity, and potential consequences of the communication.

Computational systems can contribute to monitoring, information retrieval, content classification, anomaly detection, and compliance pre-screening where these functions are sufficiently bounded and auditable. Human professionals become increasingly important where communication requires interpretation of ambiguous scientific claims, assessment of stakeholder vulnerability, contextual legal judgment, or decisions concerning potentially consequential health information.

The limitations of platform self-regulation therefore reinforce the broader principle of risk-proportionate functional asymmetry. No single technological, organizational, or regulatory mechanism can independently guarantee information integrity across decentralized digital health environments. Responsible pharmaceutical communication instead requires layered governance in which computational capabilities, platform controls, statutory regulation, and identifiable human accountability perform complementary roles.

Within this perspective, the objective of digital governance is not to eliminate every form of uncertainty from online health communication, an outcome that would be neither technically nor institutionally realistic. The more defensible strategic objective is to reduce preventable information risks, increase transparency and traceability, establish clear lines of responsibility, and ensure that greater technological scale does not correspond to weaker standards of scientific and ethical accountability.

Chapter 3: AI Technologies in the Creative and Strategic Process

The tensions identified in Chapter 2 become concrete when Artificial Intelligence is introduced into the production and distribution of pharmaceutical communication. A system may retrieve scientific evidence, assemble previously approved content components, flag a potentially problematic claim, recommend an action, or select a visual variation. These functions differ in both purpose and consequence, even when they are grouped under the broader category of AI-mediated communication. Evaluating them therefore requires attention not only to what can be automated, but also to the stage of the communication process in which automation occurs and to the decisions that computational systems are permitted to influence.

This chapter follows that operational process from content production to personalized distribution. Generative AI, Retrieval-Augmented Generation, modular content architectures, and automated prescreening are examined as possible means of increasing production capacity while maintaining connections to scientific evidence and regulatory review. Once content has been created and validated, predictive architectures such as Next Best Action engines and Dynamic Creative Optimization can support its selection, sequencing, and presentation across different contexts. These capabilities, however, depend on data collection and behavioral inference, bringing the analysis back to questions of privacy, autonomy, and trust. The resulting personalization paradox provides the basis for examining psychological reactance, data minimization, and the privacy-utility trade-off. Throughout the chapter, Artificial Intelligence is treated as a set of function-specific tools operating within evidentiary and governance constraints, rather than as an autonomous substitute for professional judgment.

3.1 Generative AI for Content Automation and Medical Compliance

Within the pharmaceutical sector, the strategic utility of Generative AI becomes particularly significant when it is integrated into a broader Human-Centered Precision communication model. By processing authorized streams of digital interactions, contextual information, and controlled scientific repositories, computational systems can support the adaptation of communication to different stakeholder needs while reducing some of the operational burden associated with large-scale content management (Sabapathy, 2025). Such adaptation should nevertheless be understood as probabilistic and supportive rather than as a mechanism capable of identifying stakeholder anxieties with certainty or guaranteeing that every generated communication will correspond perfectly to an individual's informational requirements.

The strategic challenge is therefore to ensure that algorithmic scalability contributes to communication utility without intensifying the risks of stakeholder alienation, excessive personalization, or scientific inconsistency. This requires the integration of computational tools within a rigorous ethical and operational framework. As articulated in the Data + Empathy framework developed throughout this research, Artificial Intelligence can assume a substantial role in high-speed data synthesis, pattern recognition, information retrieval, and semantic pre-screening, but it should not be assigned a fixed position that is universally invisible or exclusively confined to backend infrastructure. The appropriate degree of AI visibility, autonomy, and human supervision should instead depend on the function performed, the sensitivity of the communication context, and the potential consequences of technological error.

Accordingly, the visible interface of digital outreach does not need to remain categorically under the exclusive sovereignty of human actors. Artificial Intelligence may legitimately interact directly with stakeholders in bounded, standardized, informational, or comparatively low-risk contexts where the

function is transparent and the resulting consequences are limited. Human involvement becomes progressively more important where communication approaches clinical interpretation, emotional vulnerability, ambiguous scientific evidence, sensitive personalization, or consequential decision-making. The resulting architecture therefore reflects the principle of risk-proportionate functional asymmetry rather than a rigid separation between an automated backend and a permanently human frontend.


The utilization of Generative AI in pharmaceutical communication consequently does not imply the automated displacement of human creativity or professional judgment. Rather, it requires a deliberate combination of linguistic agility, evidence-grounding mechanisms, structured content architectures, automated risk-detection tools, and professional oversight. This balanced implementation can transform the creative workspace into a more adaptive environment in which automated iteration coexists with clearly defined institutional safeguards. The subsequent sections of this chapter deconstruct the specific engineering configurations, retrieval mechanisms, and data architectures through which digital marketing strategies can support greater content scalability under a philosophy of Compliance by Design, beginning with Enterprise Large Language Models and Retrieval-Augmented Generation architectures designed to reduce unsupported generation and improve the traceability of AI-assisted pharmaceutical communication.

3.1.1 Enterprise LLMs and Retrieval-Augmented Generation (RAG) in Pharma Copywriting

The deployment of Generative Artificial Intelligence within the pharmaceutical digital outreach ecosystem requires a structural departure from open domain commercial applications. While standard Large Language Models demonstrate high linguistic proficiency, their stochastic nature introduces the risk of semantic hallucinations, which is a computational phenomenon where the neural network generates plausible but clinically unverified medical claims (Ji et al., 2023). Within the healthcare sector, where communication is bound by strict medical accuracy and legal liabilities, such hallucinations represent an unacceptable compliance risk (Ji et al., 2023; Liang & Mackey, 2011). Consequently, modern digital strategies rely on Enterprise Large Language Models coupled with Retrieval Augmented Generation architectures to anchor algorithmic outputs to a deterministic source of truth (Sabapathy, 2025.). This advanced framework directly resolves the fundamental conflict between the static parametric knowledge hidden within deep neural weights and the highly dynamic, constantly updated nature of empirical clinical evidence required for healthcare marketing (Nogueira & Cho, 2019).

The operational mechanics of an advanced Retrieval Augmented Generation framework function through a multi stage data orchestration pipeline that occurs long before the final text generation phase is initiated. The first critical phase of this pipeline involves pre retrieval optimization and database ingestion (Nogueira & Cho, 2019). The internal knowledge base of the pharmaceutical enterprise, which contains extensive clinical guidelines, corporate summaries of product characteristics, historical marketing materials, and regulatory filings, must be subjected to a strict data preprocessing pipeline (Noy et al., 2019). Large text files are structurally fragmented into mathematically manageable units known as chunks. To prevent the loss of medical meaning, developers utilize semantic chunking algorithms that employ sliding window token strategies, ensuring that overlapping text segments preserve the continuity of scientific arguments, dosage instructions, and statistical epidemiological data (Gao et al., 2023, 2024). These chunks are then enriched with unique structural metadata tags, including target demographic parameters, product indications, and jurisdictional safety constraints, creating an organized and legally traceable enterprise data repository (Amugongo et al., 2025).

Once preprocessing is complete, the retrieval phase transitions into the mathematical domain of vector space modeling. When a creative team or an automated marketing system issues a query requesting the generation of a digital asset, such as a specialized medical copy variant or an omnichannel newsletter segment, the system does not interrogate the parametric memory of the Large Language Model directly. Instead, the user prompt is processed by a specialized biomedical embedding encoder, which maps the text into a dense, high dimensional vector space where semantic concepts are represented as mathematical coordinates (Nogueira & Cho, 2019). This user query vector is then projected against the corporate database, where a search engine measures the proximity between the prompt vector and the vectorized segments of the source medical literature (Karpukhin et al., 2020).



This proximity matching relies on advanced similarity metrics, most notably cosine similarity and inner dot product calculations, to determine which text chunks possess the highest semantic relevance to the creative request. However, fetching the raw chunks is insufficient for high risk pharmaceutical communication, necessitating a third phase known as post retrieval processing (Nogueira & Cho, 2019). Because the initial vector search may retrieve redundant information or irrelevant context that could dilute the generation quality, the system routes the top retrieved chunks through an independent cross encoder reranking model (Gao et al., 2023; Nogueira & Cho, 2019). This reranking algorithm assigns an absolute confidence score to each chunk, ordering them based on literal clinical priority and filtering out noisy data before the context window is populated, a method that significantly lowers the inference error rates of the system (Nogueira & Cho, 2019).

Once the most relevant, high confidence chunks are isolated, they are structurally injected into the context window of the Enterprise Large Language Model alongside the original user creative prompt. This engineering constraint redefines the operational boundaries of the generative engine. The Large Language Model is transformed from an autonomous content creator into a constrained context conditional generator and linguistic synthesizer (Lewis et al., 2020). The algorithm is strictly restricted to generating text derived exclusively from the retrieved text segments, converting raw scientific data into engaging marketing prose without inventing unverified product claims or omitting critical safety boundaries (Gao et al., 2023, 2024).

From an engineering perspective, this multi stage architecture successfully mitigates the black box problem of deep neural networks in regulatory healthcare communication (Sabapathy, 2025). Because the generated text is a direct semantic synthesis of the provided context blocks, the system can implement citation mechanisms that link specific output sentences back to the exact source paragraph identifier within the corporate vector database (Amugongo et al., 2025). Furthermore, this architecture allows pharmaceutical companies to achieve unprecedented creative scalability (Noy et al., 2019, 2025). Because the underlying data repository remains isolated from the generative model weights, marketing teams can safely produce hundreds of contextual variations of a single medical concept, tailoring the language to specific physician specialties or digital channels without risking the introduction of unapproved medical claims or violating institutional regulatory integrity.

3.1.2 Modular Content Creation and Hyper-Scalability

The operationalization of Retrieval-Augmented Generation architectures provides an important foundation for grounding AI-assisted communication in controlled scientific information, yet the institutional requirement for digital scalability across diverse stakeholder channels introduces a secondary execution challenge that cannot be resolved through retrieval alone. Traditional creative workflows within pharmaceutical advertising have frequently treated communication assets as relatively self-contained entities. Under this conventional paradigm, a physician detail aid, patient educational brochure, promotional email, or digital advertisement is conceptualized, designed, reviewed, and approved as a complete communication object. While this approach provides a clear unit for Medical, Legal, and Regulatory review, it can create substantial operational complexity when the same scientific message must be adapted across numerous channels, audiences, jurisdictions, formats, and digital contexts.

Contemporary omnichannel strategies can require a continuous stream of textual and visual variations responding to platform specifications, stakeholder characteristics, language requirements, and changing informational contexts. Managing this volume exclusively through repeated creation of independent assets may increase administrative burden and create bottlenecks within review workflows. The challenge is particularly relevant in pharmaceutical communication because adaptation cannot concern visual appearance or tone alone: modifications to wording, target audience, channel, surrounding context, or risk presentation can alter the regulatory interpretation of an otherwise approved scientific message.

To address this operational pressure, digital communication architectures can move from exclusively monolithic asset production toward modular content structures. Within this approach, larger communication assets are decomposed into smaller reusable components that can be stored, classified, retrieved, and assembled within controlled enterprise environments. These components may include therapeutic claims, clinical statistics, approved safety information, references, visual elements, calls to action, or other content units appropriate to the organizational use case.

The concept of modularity should not be interpreted as implying that every pharmaceutical message can be safely divided into completely independent fragments. Scientific and regulatory meaning frequently depends on the relationship between multiple elements, and individual statements can become misleading when removed from their original context. Modularization therefore requires careful definition of which components are genuinely reusable, which relationships must remain preserved, and which combinations require renewed review.

Within an enterprise content-management environment, approved modules can be associated with structured metadata that describe their operational conditions. Metadata may include information concerning intended audience, therapeutic area, product indication, jurisdiction, channel, approval status, expiration date, required references, mandatory safety information, or other internal governance parameters. Such classifications provide the computational infrastructure through which automated systems can determine whether a particular content element is potentially appropriate for reuse within a specific communication context (Sabapathy, 2025).

This metadata architecture effectively functions as a digital governance layer around each reusable component. Rather than treating an approved statement as universally deployable, the system can preserve the conditions under which that statement was reviewed. A clinical claim authorized for communication to a defined professional audience within one jurisdiction, for example, should not automatically become available for general consumer communication or for territories governed by different promotional requirements.

Traditional Monolithic Asset Workflow:

[Draft Asset] → [Human MLR Review] → [Approved Fixed Output]

Modular Content Workflow:

[Controlled Content Modules] → [Module-Level Metadata and Review]
→ [Rule-Governed Assembly] → [Context-Specific Asset] →
[Proportionate Final Review]

The strategic advantage of this architecture lies in the possibility of reusing previously reviewed components without assuming that every subsequent combination becomes automatically compliant. Module-level approval can reduce unnecessary duplication within content-production workflows, but the final communication may still require review because regulatory meaning can emerge from the interaction among individual elements, their visual hierarchy, target audience, or surrounding context.

Consequently, modular content systems should not be interpreted as mechanisms capable of generating thousands of independently deployable assets without additional human assessment. The level of review required depends on how substantially the final configuration differs from previously approved combinations and on the risk profile of the communication. Minor formatting or channel adaptations may require different governance from a newly assembled promotional claim, a new audience configuration, or a communication involving sensitive therapeutic information.

The automated recombination of approved components therefore requires more than simple probabilistic generation or unrestricted asset assembly. Even individually accurate modules can create misleading communication when combined incorrectly. An efficacy statement, for example, may require corresponding qualification, supporting evidence, or safety information, while certain content elements may be inappropriate when presented to particular stakeholder groups or within specific jurisdictions.

Within the architecture proposed in this thesis, structured Knowledge Graphs can provide one possible mechanism for representing these relationships. Knowledge Graphs encode entities and their semantic relationships in a structured form, allowing systems to model connections between claims, evidence sources, therapeutic areas, audiences, regulatory conditions, and required disclosures (Noy et al., 2019). Their relevance to pharmaceutical modular content lies in their potential to complement probabilistic language generation with explicit relational constraints.

A Knowledge Graph should not, however, be interpreted as an infallible or universally adopted regulatory engine. Its effectiveness depends on the quality of the ontology, the completeness and currency of the relationships represented, and the governance procedures through which rules are defined and maintained. Within the conceptual architecture examined here, Knowledge Graphs therefore function as a proposed structured control mechanism rather than as evidence of a standardized pharmaceutical industry implementation.

For example, a graph-based rule system could represent a relationship indicating that the use of a particular therapeutic claim requires the concurrent presence of a specified evidence reference and associated safety information. It could similarly restrict the availability of a content module according to audience, jurisdiction, indication, or approval status. These relationships can help constrain automated assembly by making certain dependencies explicit before a generative model is permitted to adapt the final linguistic or visual presentation.

The resulting architecture creates what can be described as structured flexibility. Generative models retain the capacity to support linguistic adaptation, summarization, formatting, or channel-specific presentation, while structured data systems establish boundaries around the scientific components and relationships that should remain preserved. The Large Language Model therefore operates less as an unrestricted creator of pharmaceutical claims and more as a contextual adaptation layer within a broader governed content environment.

For example, an approved scientific concept could be expressed at different levels of linguistic complexity for a professional portal, an educational resource, or another authorized channel, provided that the adaptation remains consistent with the original evidence, audience requirements, and applicable regulatory constraints. Similarly, visual composition can be adjusted for different devices or media formats while preserving required safety information and source relationships.

This combination of modular content, structured metadata, rule-based relationships, and generative adaptation can substantially increase the efficiency of pharmaceutical content operations. Its value lies not in eliminating Medical, Legal, and Regulatory review, but in allowing review resources to be concentrated more effectively by preserving previously established content relationships, identifying reusable components, and reducing unnecessary repetition across closely related assets.

The strategic outcome is therefore better described as governed scalability rather than unrestricted hyper scalability. Pharmaceutical organizations may increase the speed and flexibility with which digital assets are produced and adapted, but scalability remains conditional upon the integrity of the underlying modules, the accuracy of metadata, the maintenance of relational rules, and the continued availability of proportionate human oversight.

Within the Data + Empathy framework, modular content architectures illustrate how computational capability can expand organizational capacity without automatically expanding technological authority. Artificial Intelligence can support assembly, retrieval, linguistic adaptation, and consistency checking, while organizational governance determines which components may be combined, which adaptations require renewed professional assessment, and which communication contexts justify more intensive human review.

This systematic integration of generative flexibility and structured content governance can therefore reduce some of the temporal friction associated with conventional pharmaceutical content production while preserving the principle that scalability must remain subordinate to scientific accuracy, regulatory appropriateness, and stakeholder context. Rather than eliminating review, modular architectures reorganize it, creating the technical foundation for the automated pre-screening and Compliance-by-Design mechanisms examined in the following section.

3.1.3 Automated Pre-Screening and Compliance by Design

The structural integration of Retrieval-Augmented Generation frameworks and semantic Knowledge Graphs can establish a multilayered ecosystem for generating increasingly grounded and contextually controlled digital assets. However, within the high-risk operational environment of pharmaceutical advertising, relying exclusively on constraints applied during the generation phase is insufficient to ensure regulatory alignment. Even when computational inputs are derived from controlled repositories, the probabilistic nature of generative models can produce ambiguous formulations, unsupported implications, inappropriate combinations of otherwise accurate information, or outputs whose meaning changes when transferred across audiences and contexts (Ji et al., 2023; Gao et al., 2023).

To establish an additional layer of structural safety, the architecture proposed in this thesis incorporates an independent automated prescreening layer operating according to the broader philosophy of Compliance by Design. Rather than functioning as an autonomous regulatory authority, this layer serves as an initial computational gatekeeper that evaluates compiled text and media variants before they enter the final human Medical, Legal, and Regulatory validation phase. Its purpose is to identify potentially problematic content early in the workflow, prioritize assets requiring professional attention, and reduce part of the repetitive screening burden imposed on human review teams.

To understand the potential strategic value of this independent evaluation layer, it is first necessary to consider the limitations of conventional rule-based compliance software. Pharmaceutical organizations have historically relied on systems using controlled dictionaries, keyword detection, predefined rule sets, and regular-expression matching to identify potentially problematic terminology or missing mandatory elements. These approaches remain useful for clearly identifiable conditions, such as detecting prohibited expressions, checking the presence of required statements, or identifying unauthorized terminology.

However, rule-based tools have more limited capacity to evaluate semantic meaning when regulatory risk depends on context rather than on the presence of an isolated word. A promotional asset may avoid prohibited terminology while nevertheless implying an unsupported therapeutic benefit, overstating the certainty of scientific evidence, or creating an unbalanced overall impression through the relationship among multiple statements. Conversely, a keyword may appear in a legitimate scientific context without constituting a compliance problem. The challenge therefore lies not in the complete inadequacy of lexicographical systems, but in their limited ability to interpret complex semantic relationships and contextual meaning at scale.

Within the Compliance by Design architecture proposed in this research, these traditional mechanisms can therefore be complemented by Natural Language Processing systems and evaluator models inspired by the broader LLM-as-a-Judge paradigm (Zheng et al., 2023). Importantly, existing LLM-as-a-Judge research primarily concerns the use of language models to evaluate other model outputs and does not independently establish their reliability as pharmaceutical regulatory authorities. Their application to pharmaceutical compliance should consequently be understood as a proposed adaptation requiring domain-specific validation rather than as an established industry standard.

In this technical configuration, an independent secondary Large Language Model can be architecturally separated from the primary content-generation engine and assigned an evaluative rather than generative role. The objective of this separation is to introduce an additional review perspective into the pipeline. However, architectural separation alone does not eliminate correlated errors, shared model biases, or common limitations arising from similar training data or reasoning architectures. The evaluator therefore constitutes an additional risk-detection mechanism rather than an independent guarantee of correctness.

The secondary model can be configured using selected historical compliance examples, internal organizational policies, applicable regulatory guidance, approved product information, and formalized risk rubrics. Rather than transforming the neural network into an “automated legal expert,” these resources can provide structured criteria against which the evaluator compares generated communication. The resulting system can support preliminary identification of potentially problematic language while leaving the interpretation of ambiguous legal and medical requirements to qualified professionals.

When the primary generation engine compiles advertising variations or omnichannel messaging layouts, the resulting outputs can be routed to this evaluator model for a multidimensional semantic assessment (Zheng et al., 2023). The evaluator may examine dimensions including consistency between claims and supporting evidence, presence of required information, potentially absolute or misleading language, audience appropriateness, internal contradictions, and alignment with predefined organizational rules.

For example, a system may identify whether an efficacy statement appears without an associated safety element expected under the relevant internal communication template, or whether generated language converts a probabilistic clinical finding into an unjustified categorical conclusion. These assessments can complement deterministic rules and structured Knowledge Graph constraints by identifying semantic relationships that would be difficult to capture through keyword matching alone.

The relationship between benefit claims and risk information illustrates the importance of maintaining this distinction. An automated system can assess whether relevant safety information is present, whether it is semantically connected to a benefit statement, and whether its placement appears consistent with predefined organizational requirements. However, regulatory adequacy should not be reduced to a fixed mathematical token distance between therapeutic claims and risk disclosures. The appropriateness and prominence of risk communication depend on the complete communication context, channel, presentation, audience, applicable jurisdiction, and other factors that require professional interpretation.

Accordingly, semantic proximity can function as one computational feature within a broader prescreening architecture, but it should not be treated as a legal test of fair balance. The evaluator can flag assets in which benefit information and safety information appear insufficiently connected or structurally separated, allowing a reviewer to determine whether the overall communication remains appropriate.

The evaluator model can similarly analyze the linguistic formulation of therapeutic claims to identify language that may overstate the underlying evidence. For example, deterministic expressions such as cure, eradicate, guarantee, or eliminate may warrant additional scrutiny when the cited clinical evidence supports only a probabilistic or conditional conclusion. The purpose of automated detection is not to determine autonomously that a regulatory violation has occurred, but to recognize linguistic patterns associated with elevated review risk.

This distinction also applies to the use of confidence or risk scores. Evaluator models can assign internal scores representing the model's assessment of whether an output corresponds to predefined evaluation criteria. Such scores should not be interpreted as mathematical probabilities that a communication asset is legally compliant. They are better understood as operational triage signals that help prioritize review and determine whether additional inspection is warranted.

Thresholds can therefore be calibrated according to organizational risk tolerance and the consequences of false positives and false negatives. A conservative threshold may identify a larger number of potentially problematic assets, increasing review workload but reducing the likelihood that questionable content proceeds unnoticed. A more permissive threshold may improve workflow efficiency while increasing the need for downstream safeguards. The appropriate balance depends on the use case and should itself remain subject to organizational governance.


Compiled Asset Variants

|



AI-as-a-Judge Semantic Prescreening

|



Potential Issues Flagged No Major Issues Detected

| |



Quarantine / Revision Queue Prioritized Routing

and Diagnostic Feedback to Human MLR Review

If the prescreening architecture identifies a potentially unsupported implication, inconsistent claim, missing information element, or other predefined risk indicator, the asset can be temporarily quarantined within the digital asset-management workflow. The system may identify the relevant sentence, attach a diagnostic metadata tag indicating the type of issue detected, and route the material either to a controlled revision stage or directly to professional review.

Automated feedback can also be used to support subsequent revision. For example, the primary generation engine may receive structured instructions indicating that a particular claim lacks sufficient evidentiary support, uses language exceeding the approved source material, or requires an associated disclosure. The system can then generate a revised version subject to the same evidence and content constraints.

However, this process should not operate as an unconstrained recursive loop in which one Artificial Intelligence system repeatedly rewrites content until another AI system assigns a sufficiently high score. Such an architecture could create false confidence by optimizing text toward the evaluator's preferences without necessarily improving genuine regulatory appropriateness. Automated revision should therefore remain bounded by controlled source material, predefined rules, version tracking, and escalation mechanisms that preserve human visibility over substantive changes.

Conversely, when an asset does not trigger major prescreening concerns, it can be prioritized for human Medical, Legal, and Regulatory review rather than considered automatically approved. The absence of computational flags does not demonstrate compliance, because evaluator models may fail to identify subtle scientific, contextual, or jurisdiction-specific problems. The automated layer consequently performs a filtering and prioritization function rather than replacing the final professional assessment.

By screening large volumes of content before manual review, this dual-model architecture can nevertheless improve workflow efficiency. Human compliance officers can devote a greater proportion of their attention to ambiguous, high-risk, or strategically significant communication rather than repeatedly performing elementary checks that can be partially automated. The resulting benefit is therefore not the elimination of regulatory review but the more efficient allocation of specialist expertise across an increasingly complex content environment.

Ultimately, this computational layer reflects the broader logic of the Data + Empathy framework and its principle of risk-proportionate functional asymmetry. Automated systems can absorb portions of the mechanical and high-volume burden associated with rule checking, evidence comparison, semantic anomaly detection, and preliminary content triage. Human professionals retain responsibility for interpreting scientific nuance, evaluating context, resolving legal ambiguity, assessing stakeholder sensitivity, and assuming final accountability for consequential communication.

Compliance by Design should therefore be understood not as the replacement of human regulatory judgment with Artificial Intelligence, but as the integration of computational safeguards earlier in the content lifecycle. Its strategic value lies in creating a layered architecture in which deterministic rules, structured knowledge systems, probabilistic evaluators, and human Medical, Legal, and Regulatory expertise operate as complementary mechanisms. This configuration allows pharmaceutical communication to pursue greater operational scalability while preserving the principle that automated prescreening remains supportive rather than authoritative.

3.2 Predictive Architectures and Next Best Action Engines

The computational synthesis and automated verification of modular healthcare content assets, as detailed in the preceding sections of this research, represent only the foundational production layer of contemporary digital advertising strategy. To achieve effective coordination within increasingly fragmented digital networks, these controlled assets must be connected to distribution architectures capable of supporting more context-aware and data-informed outreach. Having a library of approved atomic components is strategically insufficient if the enterprise lacks the analytical infrastructure required to evaluate which asset may be most appropriate for a particular stakeholder, channel, or moment of interaction. This section analyzes the predictive architectures that support this orchestration process, examining the integration of Next Best Action decision engines and Dynamic Creative Optimization models through which behavioral and contextual data can inform the selection, sequencing, and presentation of digital communication assets. Rather than functioning as an autonomous “operational brain” that deterministically identifies the correct communication for each individual, these systems should be understood as decision-support mechanisms that estimate the potential relevance of different actions within predefined strategic, evidentiary, privacy, and regulatory constraints.

3.2.1 Predictive Behavioral Modeling and Data Ingestion Pipelines

The operational evolution of digital pharmaceutical communication is increasingly associated with the transition from retrospective data analysis toward more dynamic forms of behavioral modeling. Historically, the architectural backbone of commercial health-data utilization relied heavily on the practices of commercial Health Information Organizations, which aggregated prescription tracking records, longitudinal claims data, and pharmacy fulfillment metrics for analytical and commercial purposes (Fugh-Berman, 2008). While these legacy data practices enabled commercial departments


to evaluate past behavioral patterns and adjust localized engagement strategies retrospectively, they were inherently limited by the temporal distance between data collection, interpretation, and subsequent strategic execution. Contemporary digital communication architectures can reduce this delay through machine-learning pipelines designed to ingest and process multiple data streams with greater frequency, thereby supporting more responsive forms of stakeholder analysis (Sabapathy, 2025).

These computational infrastructures may integrate heterogeneous data streams captured across the broader enterprise ecosystem (Verma et al., 2024). Depending on the technical configuration and legal basis for processing, data-ingestion pipelines can incorporate signals derived from customer relationship management systems, website navigation patterns, virtual event participation, newsletter interactions, scientific-content downloads, and other digital touchpoints. Distributed data-processing architectures can support the handling of higher-volume interaction streams, although their implementation depends on organizational infrastructure, privacy requirements, system design, and the specific communication use case.

Within the behavioral modeling architecture proposed in this thesis, selected interaction signals can be transformed into structured features that contribute to the updating of stakeholder profiles. For example, repeated engagement with specific scientific topics, prolonged consultation of certain clinical resources, or recurring searches concerning a therapeutic area may indicate increased informational interest. These signals should not be interpreted as direct evidence of a stakeholder's clinical intent, diagnostic activity, or future prescribing behavior. Rather, they function as probabilistic indicators that may support the identification of emerging informational patterns.

The resulting profile therefore represents an analytical approximation rather than a complete or real-time representation of stakeholder needs. Dynamic behavioral profiling can complement traditional static segmentation by allowing organizations to incorporate more recent interaction data into communication planning, while the interpretation of those data remains subject to privacy constraints, data-quality limitations, contextual uncertainty, and organizational governance (Sabapathy, 2025).

4. DATA INGESTION PIPELINE ARCHITECTURE



Furthermore, predictive behavioral modeling may support the anticipation of potential informational needs before they are explicitly articulated. This does not imply that the system can reliably infer a clinician's underlying clinical challenges or future therapeutic decisions. Instead, machine-learning models can identify patterns within historical and current interaction data that are statistically associated with particular categories of informational interest.

For instance, if a healthcare professional demonstrates increasing engagement with diagnostic guidelines, scientific publications, or clinical evidence concerning a specific condition or therapeutic area, the system may classify this activity as a signal of heightened informational relevance. The digital ecosystem can then prioritize potentially useful scientific evidence, safety information, educational resources, or support materials for subsequent review and delivery. Such prioritization should remain constrained by applicable data-protection requirements, consent conditions, audience eligibility, and internal compliance rules.


In this sense, predictive modeling can shift pharmaceutical digital communication from purely retrospective segmentation toward a more responsive form of decision support. Its strategic value lies not in predicting individual clinical behavior with certainty, but in helping organizations interpret evolving patterns of engagement and allocate appropriate informational resources more efficiently. Within the broader Data + Empathy framework, behavioral data therefore serve as one input into communication planning rather than as a definitive representation of stakeholder intention, preserving the distinction between computational inference and meaningful human interpretation.

3.2.2 The Next Best Action Orchestration Framework

The computational realization of dynamic behavioral profiling finds one of its principal operational applications within the Next Best Action framework, which supports the orchestration of omnichannel pharmaceutical communication. Rather than relying exclusively on legacy automated scheduling parameters or static decision trees, Next Best Action systems can integrate multiple contextual variables to estimate which communication intervention may be most appropriate at a particular point in the stakeholder journey. These variables may include the timing of outreach, the available communication channel, and the type of scientific or educational content to be delivered (Sabapathy, 2025). The objective is therefore not to identify a universally “correct” action, but to prioritize among several possible actions according to expected relevance, stakeholder context, organizational objectives, and applicable governance constraints.

Within the architecture proposed in this thesis, Next Best Action decision engines can be supported by probabilistic predictive models, optimization techniques, and, in more advanced configurations, approaches such as reinforcement learning or multi-armed bandit algorithms. These techniques should not be interpreted as constituting a universal technological standard across pharmaceutical organizations, but as possible computational methods through which competing actions may be compared. The engine can assign relative scores to potential interventions according to variables such as recent interaction patterns, historical channel responsiveness, contextual signals, content eligibility, and the estimated utility of each communication option. These analytical variables are then cross-referenced with compliance constraints, consent conditions, channel permissions, and broader enterprise objectives (Sabapathy, 2025).

For instance, if a healthcare professional performs a specialized query within an authorized clinical portal, the system may evaluate the relative suitability of several subsequent interactions. Depending on the available permissions and communication context, a secure portal notification, an educational resource, or a conventional email may receive different relevance scores. Within the framework proposed in this thesis, the orchestration engine can simultaneously identify content modules whose metadata correspond to the therapeutic area, jurisdiction, audience, and informational context involved, such as an epidemiological dataset, a peer-reviewed clinical publication, or approved safety information.



This automated decision-support process can help pharmaceutical organizations reduce indiscriminate communication and replace purely linear broadcast strategies with more context-aware engagement models (Verma et al., 2024). However, the recommendations produced by a Next Best Action system remain probabilistic assessments derived from incomplete behavioral signals. A high predicted relevance score does not demonstrate that a stakeholder actually requires a particular message, while a low score does not prove that communication would be inappropriate. The system therefore functions as an orchestration and prioritization mechanism rather than as an autonomous interpreter of stakeholder intention.

Within the framework proposed in this thesis, a potential outreach action can also be suppressed or deferred when predefined criteria indicate insufficient expected relevance, inadequate consent, inappropriate timing, excessive communication frequency, or unresolved compliance concerns. Such thresholds should be understood as organizational decision rules rather than mathematical guarantees that an intervention will or will not generate psychological reactance. Their purpose is to introduce restraint into automated outreach and reduce unnecessary or poorly contextualized communication.

By subordinating communication frequency and channel selection to a combination of informational relevance, stakeholder permissions, and governance conditions, Next Best Action systems can support a more selective approach to pharmaceutical engagement. Their strategic contribution therefore lies not in maximizing the volume of personalized touchpoints, but in improving the allocation of communication opportunities while preserving the possibility that, in some situations, the most appropriate next action is no action at all. This principle is consistent with the broader Data + Empathy framework, in which computational systems support prioritization while human and institutional governance determine the acceptable boundaries of automated engagement.


3.2.3 Dynamic Creative Optimization and Algorithmic Visual Strategy

The strategic utility of the Next Best Action engine can be extended through Dynamic Creative Optimization frameworks, which translate distribution decisions into dynamically assembled visual communication strategies. Within contemporary digital environments, visual design does not necessarily need to remain a single static asset developed and reviewed exclusively as a monolithic entity. Within the conceptual model proposed in this research, Dynamic Creative Optimization can

support the structured assembly of modular graphic components, including typographic hierarchies, imagery, interface layouts, content blocks, and other predefined visual elements (Sabapathy, 2025; Jones et al., 2024). Once the underlying orchestration system identifies an eligible content asset and communication channel, the optimization framework can assemble approved visual components into a layout adapted to the characteristics of the selected channel, device, audience category, and informational context.

This dynamic visual adaptation represents one possible application of the Data + Empathy framework within the creative process. Behavioral and contextual data can inform decisions concerning information density, visual hierarchy, presentation format, and interface organization, with the objective of improving accessibility and reducing unnecessary communication friction. However, these data should not be interpreted as providing direct access to an individual's current cognitive or emotional state. Dynamic Creative Optimization can estimate which presentation characteristics may be appropriate according to observed interaction patterns, but the resulting design remains a probabilistic adaptation rather than a precise representation of the user's psychological condition.

The framework can operate by separating content logic from selected presentation rules, allowing the visual interface to adapt predefined graphical properties according to contextual and engagement signals. Such adaptation may include the prioritization of particular content formats, the organization of information within a page, or the selection among previously approved visual templates. The purpose is therefore not to generate unrestricted visual variation, but to support controlled flexibility within predefined design and governance boundaries.



For example, where interaction signals suggest sustained engagement with detailed scientific information, the interface may prioritize structured clinical data, comparative information, references, and clearly presented safety material. Conversely, communication occurring within shorter-form or mobile-first environments may require a more concise visual hierarchy and immediate informational clarity. These adaptations should not be interpreted as proof that the system has identified a specific “systematic processing state” or that a particular layout will necessarily reduce cognitive load. Rather, they represent design choices informed by channel characteristics, prior engagement patterns, and established principles of information presentation.

Because dynamic visual assembly operates within the broader governance architecture described in the preceding sections, modular elements can be associated with metadata, structured rules, and

automated prescreening mechanisms intended to preserve important relationships among claims, evidence, safety information, and audience conditions. Knowledge Graphs and other rule-based systems can therefore support the maintenance of these relationships across dynamically generated layouts, while automated screening can identify configurations that require additional review. These mechanisms contribute to regulatory consistency but do not independently guarantee medical accuracy or compliance.

The role of the visual strategist consequently evolves from the exclusive production of fixed assets toward the design and governance of adaptable visual systems. Creative professionals define typographic hierarchies, component relationships, design boundaries, visual rules, and accessibility principles within which computational systems can generate permitted variations (Jones et al., 2024; Sabapathy, 2025). Human expertise therefore remains important not only for aesthetic quality but also for determining whether dynamically assembled communication is appropriate to its audience and context.

Within the Data + Empathy framework, Dynamic Creative Optimization should therefore be understood as a mechanism for controlled visual adaptation rather than as an autonomous system capable of reading or responding directly to a stakeholder's internal cognitive state. Data can inform the selection and organization of visual alternatives, while human-centered design and governance establish the boundaries within which those alternatives remain scientifically responsible, contextually appropriate, and consistent with the requirements of pharmaceutical communication.

3.3 The Personalization Paradox: Balancing Relevance and User Privacy

The architectural capability to orchestrate digital outreach through Next Best Action engines and to assemble customized visual layouts through Dynamic Creative Optimization models, as detailed in the preceding section, represents a significant development within contemporary advertising strategy (Sabapathy, 2025). However, the practical implementation of this predictive ecosystem introduces an important tension when algorithmic personalization intersects with user perceptions of privacy, autonomy, and informational control. This tension can be described through the broader concept of the personalization paradox: while stakeholders may value communication that is relevant, accessible, and responsive to their informational context, the same users may become uncomfortable when the degree of tracking, inference, or behavioral profiling required to produce that relevance becomes excessively visible or difficult to understand (Verma et al., 2024).

Within healthcare communication, this tension is particularly significant because personalization may involve information related to health interests, therapeutic concerns, professional activity, or other sensitive behavioral signals. The perceived usefulness of a personalized message therefore cannot be evaluated independently from the methods through which the underlying data were collected, processed, and translated into communication. Greater precision does not automatically generate greater trust; depending on context, the same personalization mechanism may be interpreted as helpful, intrusive, opaque, or inappropriate.

The personalization paradox should consequently be understood not as a deterministic rejection of algorithmic targeting, but as a context-dependent relationship between perceived utility and perceived intrusion. Factors such as the sensitivity of the information involved, transparency regarding data use, user control, communication purpose, institutional credibility, and the degree of inference performed by the system may all influence whether personalization is experienced as valuable or invasive.

This section examines the multidimensional dimensions of this tension by analyzing psychological reactance in response to perceived algorithmic intrusion, the privacy-utility trade-off created by data-


intensive personalization architectures, and the role of the Data + Empathy framework in establishing more proportionate boundaries between computational precision and human-centered communication. Rather than attempting to eliminate personalization or maximize it indiscriminately, the objective is to identify the conditions under which data-driven relevance can be pursued without undermining autonomy, privacy, and stakeholder trust.

3.3.1 Algorithmic Reactance and the Privacy Paradox in Healthcare

To understand the resistance that may emerge when personalization makes underlying tracking practices especially visible, the phenomenon can be interpreted through Psychological Reactance Theory. The theory proposes that individuals may experience a motivational response when they perceive that their freedom of choice, autonomy, or control is being restricted or threatened (Brehm, 1966). Within digital environments, this response can become relevant when users perceive that personal data are being collected, inferred, or used in ways they did not expect, particularly when those practices influence the information or persuasive messages they subsequently receive.

In pharmaceutical communication, this form of reactance may occur when personalization is perceived as crossing the boundary between useful relevance and excessive behavioral profiling (Amini Khiabani & Rezaeian, 2026). A message tailored to a user's apparent health interests may be interpreted positively when it provides timely and useful information, but the same message may become uncomfortable if the user becomes aware that sensitive searches, browsing patterns, or other behavioral signals have been used to infer personal concerns without sufficient transparency or control.

The resulting response should not be described as a uniform or inevitable rejection of personalized communication. Perceived intrusiveness varies according to the sensitivity of the data involved, the transparency of the process, the relationship between the user and the communicating organization, the perceived usefulness of the message, and the degree of control available to the individual. Personalization can therefore generate both utility and discomfort depending on the context in which it is experienced.



This tension is closely related to the broader privacy paradox, which describes the discrepancy that can exist between users' stated concerns regarding privacy and their actual willingness to disclose personal information in exchange for convenience, personalization, or access to digital services. Within health-related contexts, the relevance of this paradox becomes particularly important because personal and behavioral data may concern sensitive aspects of an individual's wellbeing, symptoms, therapeutic interests, or healthcare decisions.

For digitally experienced users, repeated exposure to platform recommendation systems, targeted advertising, data-breach reporting, and discussions surrounding algorithmic influence may contribute

to greater awareness of the mechanisms underlying personalization (Mansfield et al., 2025). This awareness should not be interpreted as evidence that Generation Z possesses a universal or involuntary “algorithmic intuition.” Rather, some users may become more attentive to recognizable indicators of targeting, sponsorship, recommendation logic, or behavioral tracking.

When an Enterprise Large Language Model or personalization system dynamically adapts educational material or disease-awareness content according to observed health-related interests, the precision of the resulting communication may therefore influence how the user interprets the interaction. If the basis of the personalization is unclear, the user may question how much information has been collected and how it has been used. Conversely, transparent explanations, clear boundaries around data use, and meaningful privacy controls may reduce uncertainty and support a more acceptable form of personalized communication.

The relationship between personalization and trust is therefore not determined solely by the sophistication of the underlying technology. Highly precise targeting does not automatically increase credibility, and visible personalization does not necessarily undermine it. Trust depends on the broader interaction among perceived relevance, transparency, institutional legitimacy, data sensitivity, and user autonomy.

Within the Data + Empathy framework, this distinction is particularly important. Behavioral and contextual data can support more relevant communication, but the interpretation and use of those data should remain proportionate to the sensitivity and purpose of the interaction. Human-centered governance is therefore required not only when an algorithm produces an incorrect output, but also when technically accurate personalization risks becoming psychologically intrusive or ethically disproportionate.

Algorithmic reactance can consequently be understood as a potential signal that personalization has exceeded the user's perceived boundary of acceptable influence. The appropriate strategic response is not to eliminate personalization entirely, but to reduce unnecessary inference, strengthen transparency, preserve meaningful user control, and align the intensity of data-driven adaptation with the actual informational value provided to the stakeholder.

3.3.2 The Privacy Utility Tradeoff and Data Minimization Architectures

The computational management of the personalization paradox requires digital strategists to address the broader privacy-utility trade-off. This tension reflects the fact that predictive performance, responsiveness, and contextual relevance may improve when machine-learning systems are provided with richer and more granular data, while the collection and processing of personal information simultaneously increase legal, ethical, and reputational risks. Within pharmaceutical communication, this balance is particularly sensitive because data may concern health-related interests, professional behavior, therapeutic preferences, or other categories of information whose use requires careful governance (Regulation (EU) 2016/679).

The General Data Protection Regulation establishes principles including purpose limitation, data minimization, transparency, and heightened protection for special categories of personal data such as health information. Consequently, the strategic objective cannot be to maximize data collection simply because additional data may improve predictive performance. Organizations must instead determine which information is actually necessary for a specific communication function and whether the expected analytical value is proportionate to the privacy implications of processing that information.

Pharmaceutical organizations therefore face a significant operational tension. More granular stakeholder profiles may increase the precision of certain personalization systems, yet the continuous aggregation of behavioral and health-related data is neither technically necessary for every form of useful personalization nor automatically compatible with responsible data governance. The relationship between analytical utility and data volume is consequently not linear: beyond a certain point, additional information may introduce greater legal exposure, complexity, bias, and perceived intrusiveness without producing equivalent communication value.

The Optimization Tradeoff Matrix:

Higher Analytical Utility ← May Benefit From → Greater Data Granularity

vs.

Privacy and Regulatory Protection ← Requires → Purpose Limitation
+ Data Minimization + Appropriate Safeguards

This reframing is important because privacy-preserving architectures should not be interpreted as a secondary constraint imposed after a predictive system has already been designed. Instead, privacy considerations can be incorporated directly into the architecture through Privacy by Design and data-minimization principles. The objective is to achieve sufficient analytical utility with the smallest quantity and sensitivity of data reasonably required for the intended communication purpose.

Several technical approaches may contribute to this objective depending on the use case, infrastructure, and legal conditions involved. Federated learning, for example, can allow certain machine-learning models to be trained or updated across distributed data environments without necessarily transferring all raw local data to a central repository (Yang et al., 2019). Differential privacy can introduce mathematically controlled noise into aggregated outputs in order to reduce the risk that information about specific individuals can be inferred from analytical results (Dwork & Roth, 2014). Edge-processing architectures may similarly allow selected computations to occur closer to the user's device or local environment rather than requiring every interaction signal to be transmitted directly to centralized enterprise systems.

These technologies should not be treated as universally appropriate or as automatically compliant solutions. Their effectiveness depends on implementation quality, threat models, organizational capability, the type of data involved, and the specific analytical objective. Federated learning, for instance, does not inherently eliminate privacy risk, while differential privacy introduces trade-offs between privacy protection and analytical accuracy. Similarly, local or edge processing can reduce certain forms of centralized data exposure without removing the need for transparent governance, consent where required, or lawful processing conditions.

Within the conceptual architecture examined in this thesis, privacy-preserving approaches can therefore be used selectively to reduce unnecessary centralization and limit exposure of raw personal data. A system may process certain interaction patterns locally, aggregate only the information required for a particular analytical function, or use anonymization and pseudonymization techniques where appropriate. The resulting architecture seeks to preserve useful signals while reducing the quantity of identifiable information retained within central enterprise repositories.

The strategic implication is that personalization does not need to depend on unrestricted behavioral surveillance. In many communication contexts, sufficiently useful adaptation may be achieved

through lower-resolution signals, contextual information, explicit user preferences, or aggregated behavioral patterns. This is particularly relevant within healthcare communication, where the marginal benefit of additional personalization must be weighed against the sensitivity of the data required to achieve it.

The privacy-utility trade-off should therefore be interpreted as an optimization problem governed by proportionality rather than as a binary conflict between analytical performance and legal compliance. The objective is not to maximize predictive accuracy at any cost, nor to eliminate all data-driven personalization, but to identify the minimum level of information required to generate meaningful utility within acceptable legal and ethical boundaries.

Within the Data + Empathy framework, this principle reinforces the distinction between computational possibility and organizational appropriateness. Data determine what a system may be capable of inferring, but governance determines whether those inferences should be generated, retained, or used in communication. Human-centered evaluation becomes particularly important where increasingly granular personalization offers limited additional utility while substantially increasing privacy risk or the perception of surveillance.

Privacy-preserving architectures therefore contribute to the broader principle of risk-proportionate functional asymmetry. Low-risk personalization may rely on limited contextual or aggregate data, while more sensitive applications require progressively stronger safeguards, narrower purposes, and greater human oversight. The aim is not to construct the most data-intensive personalization system technically possible, but to create one whose analytical sophistication remains proportionate to the communication objective and to the rights of the individuals whose data make that personalization possible.


3.3.3 Resolving the Paradox: Risk-Proportionate Functional Asymmetry and the Data + Empathy Framework

Ultimately, the management of algorithmic reactance and the optimization of the privacy-utility trade-off can be addressed through the core conceptual model developed in this research, defined as the Data + Empathy framework (Verma et al., 2024). This thesis argues that the personalization paradox is unlikely to be effectively addressed either by reducing the sophistication of contemporary machine-learning systems or by returning to the less adaptive broadcast communication models characteristic of the legacy mass-marketing era (Jones et al., 2024). Instead, the central challenge lies in determining how computational capability should be allocated across communication functions in a manner that preserves relevance while protecting autonomy, transparency, and trust.

Rather than establishing a rigid operational boundary between an invisible computational backend and a permanently human front-end interface, the framework proposes a model of risk-proportionate functional asymmetry. Within this architecture, the appropriate role of Artificial Intelligence depends on the function being performed, the sensitivity of the data involved, the degree of uncertainty, the potential consequences of error, and the relational intensity of the communication context. The objective is therefore not to define a universally fixed location for Artificial Intelligence within the communication infrastructure, but to determine where automation provides meaningful utility and where human involvement becomes increasingly necessary.

Data + Empathy Framework

Risk-Proportionate Functional Asymmetry



Within this model, Artificial Intelligence can support high-speed data synthesis, pattern recognition, information retrieval, cross-encoder reranking, modular content organization, predictive modeling, and semantic compliance prescreening (Gao et al., 2023; Zheng et al., 2023; Sabapathy, 2025). In comparatively low-risk and standardized contexts, computational systems may also interact more directly with stakeholders through functions such as information retrieval, navigation assistance, administrative support, or clearly bounded educational communication. These applications remain subject to transparency, evidentiary constraints, privacy requirements, and appropriate escalation mechanisms. As the interaction becomes more clinically sensitive, emotionally complex, ambiguous, or consequential, however, the need for meaningful human involvement increases.

Human contribution should therefore not be understood as a purely emotional front-end layer added after automated systems have already determined the communication strategy. Human-centered evaluation plays a broader role by interpreting technological possibilities in relation to stakeholder context, assessing whether personalization is proportionate to the informational need, evaluating uncertainty, and determining whether a technically feasible intervention is also ethically and relationally appropriate. Within the Data + Empathy framework, empathy does not represent a computational simulation of emotional behavior, but rather the organizational capacity to evaluate data-driven possibilities through the perspective of the people affected by them.

The interaction among these dimensions can be expressed through three complementary functions. Data create computational possibility by identifying patterns, organizing evidence, and supporting analytical precision. Empathy evaluates whether those possibilities are appropriate in relation to human context, vulnerability, autonomy, and informational need. Governance determines whether, where, and under what conditions those possibilities should become organizational action. In this sense, the objective is not to make Artificial Intelligence increasingly human in appearance, but to make the organizational use of Artificial Intelligence more human-centered.

The integration of human-centered evaluation can also support more transparent forms of personalization by encouraging clearer explanations of data use, appropriate disclosure of AI involvement, and meaningful mechanisms through which stakeholders can exercise control over their privacy preferences. These measures do not guarantee trust, but they can improve the conditions under which personalized communication is interpreted as legitimate and useful rather than opaque or excessively intrusive.

The Data + Empathy framework therefore reframes the personalization paradox as a problem of technological role allocation rather than as a simple opposition between automation and human interaction. The appropriate balance does not consist in maximizing automation or categorically restricting it, but in assigning computational and human functions according to risk, sensitivity, uncertainty, and consequence. Through this risk-proportionate configuration, pharmaceutical organizations can preserve the analytical advantages of Artificial Intelligence while reducing the conditions under which personalization becomes intrusive, disproportionate, or insufficiently accountable.

Chapter 4: The Regulatory Wall and Borderline Contexts

Technical capability alone does not determine whether a communication practice is legally or ethically acceptable. The systems examined in Chapter 3 can generate, adapt, select, and distribute health-related content at considerable speed, but they operate across legal categories that vary according to jurisdiction, audience, product classification, and communicative purpose. These boundaries become particularly difficult to identify in digital environments, where educational material, disease-awareness initiatives, wellness claims, influencer content, and commercial promotion may circulate through the same platforms and adopt similar formats. Regulatory assessment must therefore consider not only the wording of an individual claim, but also its sponsorship, targeting strategy, product associations, and broader distribution context.

Chapter 4 examines the governance problems that arise when data-driven communication develops more rapidly than the mechanisms responsible for overseeing it. Indirect advertising and borderline health communication reveal the difficulty of distinguishing legitimate information from commercial influence, while platform moderation and national regulation distribute responsibility among actors with different powers and incentives. The semaglutide case illustrates how these limitations can become visible in practice. The analysis then turns to Compliance by Design and AI-assisted prescreening as possible preventive safeguards, before addressing algorithmic transparency, misinformation, and the continuing need for identifiable human accountability. Regulation is therefore treated not simply as an external restriction on innovation, but as part of the governance architecture that determines how technological capabilities may be used and who remains responsible for their consequences.

4.1 Surrogate and Hidden Advertising: Regulatory Gaps in High-Exposure Contexts

The regulation of prescription-drug communication creates substantially different opportunities for direct public engagement across jurisdictions. In the United States, direct-to-consumer advertising of prescription medicines is permitted subject to specific regulatory requirements concerning scientific substantiation, accuracy, and the presentation of benefit and risk information (FDA, 2023; 21 CFR § 202.1). Within the European Union, by contrast, advertising to the general public of prescription-only medicinal products is prohibited, while promotion directed toward healthcare professionals remains permitted under defined regulatory conditions (European Parliament & Council, 2001, Art. 88).

These restrictions do not exclude pharmaceutical organizations from broader forms of public health communication. Life sciences organizations may develop unbranded disease-awareness, educational, corporate, and public-health initiatives that are not centered on the promotion of an identifiable prescription medicinal product (Jones et al., 2024; Verma et al., 2024). Unbranded disease-awareness communication can include educational campaigns, disease-awareness initiatives, or lifestyle-related information that omit explicit proprietary brand names or product references while addressing a particular therapeutic or health domain (Liang & Mackey, 2011; Verma et al., 2024). Under certain circumstances, however, the distinction between informational communication and indirect promotion may become less clear, particularly where an initiative is closely connected to the sponsoring organization's commercial portfolio.

Within digital environments, disease-awareness initiatives may incorporate search-engine optimization, targeted content distribution, and social-listening analytics to identify emerging informational needs, public concerns, and health-related discussions (Sabapathy, 2025).

Organizations can use these signals to develop educational resources addressing symptoms, disease progression, diagnostic pathways, or available forms of care. Such practices can improve the contextual relevance of communication, but the use of behavioral or search-related information does not in itself determine whether the resulting initiative should be classified as promotional.

The distinction becomes more complex where disease-awareness communication is strategically aligned with therapeutic areas in which the sponsoring organization has substantial commercial interests (Jones et al., 2024). Such alignment may generate indirect commercial benefits by increasing awareness of a condition, encouraging diagnostic consultation, or expanding public attention toward a therapeutic category (Verma et al., 2024). However, these possible commercial effects should not be treated as evidence that disease-awareness initiatives necessarily function as surrogate advertising. Their legal and ethical characterization depends on the communication as a whole, including its content, purpose, product associations, targeting practices, and surrounding commercial context.

This creates an important challenge for digital communication ethics. On the one hand, unbranded disease-awareness campaigns can perform legitimate public-health functions by improving health literacy, reducing stigma surrounding underdiagnosed conditions, and encouraging individuals to seek appropriate professional medical advice (Jones et al., 2024). On the other hand, concerns may arise where data-driven targeting, selective presentation of therapeutic pathways, or other persuasive techniques are used in ways that amplify health-related anxieties or blur the distinction between education and commercial influence (Liang & Mackey, 2011). The relevant governance problem therefore lies not in disease-awareness communication itself, but in the conditions under which informational communication may acquire an indirect promotional function.

This ambiguity can complicate institutional oversight, particularly where communication occupies the boundary between public-health education, commercial persuasion, and medicinal-product promotion. Supervisory authorities such as the Food and Drug Administration or national European regulatory bodies may consequently need to evaluate not only the explicit wording of a communication asset, but also its broader purpose, sponsorship, targeting strategy, product associations, and digital distribution context.

Regulatory complexity becomes particularly visible within borderline contexts involving products situated close to the legal distinctions among foods, cosmetics, dietary supplements, wellness products, and medicinal products (Catalani et al., 2021). During periods of heightened public attention to physical performance, immune support, weight management, or wellness, including international sporting events or public-health crises, commercial communication may increasingly respond to emerging consumer concerns. In such environments, persuasive health claims can circulate rapidly, creating additional challenges where product classification, evidentiary requirements, and the distinction between general wellness claims and medicinal claims are not immediately apparent to consumers (Rodríguez Hernández et al., 2024).

In the United States, dietary supplements are generally not subject to the same pre-market approval requirements applicable to new drugs and operate within a regulatory framework that places substantial emphasis on post-market oversight (DSHEA, 1994; FDA, 2022). Commercial actors may use structure-and-function claims concerning the intended effect of a product on the body, subject to applicable requirements and provided that such communication does not improperly present the product as preventing or treating disease and remains adequately substantiated (FDA, 2022; FTC, 2023).

Digital platforms can extend the reach and visibility of these forms of health-related marketing. Through influencer networks on platforms such as Instagram and TikTok, commercial content can

become integrated into media environments in which promotional communication, personal experience, lifestyle narratives, and user-generated content coexist (Deohans et al., 2026; Lim et al., 2022). Influencers may incorporate wellness products or dietary practices into personal narratives in ways that can make the commercial character of the communication less immediately distinguishable from ordinary lifestyle content, particularly where sponsorship relationships are insufficiently visible (Gell et al., 2026).

The interpretation of such content may also be influenced by platform-specific credibility cues. Younger digital users can encounter health communication in environments where social endorsement, creator familiarity, perceived authenticity, visual presentation, and interpersonal identification coexist with more systematic forms of evidence evaluation (Chantre et al., 2025). These cues may affect how promotional messages are interpreted, but they should not be understood as evidence that Generation Z uniformly relies on heuristic processing or systematically privileges peer recommendations over scientific information. The governance concern is instead that decentralized creator networks can increase the volume and speed of health-related claims circulating outside conventional corporate review structures, including claims whose commercial relationships or evidentiary basis may not be immediately transparent (Deohans et al., 2026).

The increasing volume of unbranded, indirect, and creator-mediated health communication also exposes practical limitations in exclusively retrospective enforcement models (Gomes & Sultan, 2024). Regulatory authorities frequently rely on monitoring, administrative action, warnings, and legal proceedings that necessarily operate through formal investigative procedures. Digital communication, by contrast, can circulate rapidly and achieve substantial exposure before regulatory evaluation is completed (Vosoughi et al., 2018). In some circumstances, therefore, potentially misleading or insufficiently substantiated claims may already have reached significant audiences before corrective action becomes possible. This temporal asymmetry does not demonstrate that retrospective enforcement is ineffective, but it highlights the value of complementary preventive and organizational safeguards.

For life sciences organizations, these conditions strengthen the importance of proactive internal governance alongside formal legal compliance. The objective should not be to replace public regulation with organizational self-regulation, but to integrate scientific review, transparency, data governance, sponsorship controls, and proportionate technological safeguards earlier in the communication lifecycle. Within the Data + Empathy framework developed in this thesis, data-driven capabilities can support the identification of informational needs and the adaptation of communication, while human-centered evaluation examines whether a proposed intervention is appropriate in relation to stakeholder vulnerability, autonomy, context, and informational utility. Governance subsequently determines whether, where, and under what conditions those technological capabilities should be translated into organizational action.

By integrating transparency, scientific accuracy, proportionality of data use, and respect for stakeholder autonomy into digital communication strategies, life sciences organizations can create more favorable conditions for credible and responsible engagement. Such measures do not guarantee stakeholder trust or eliminate regulatory risk, but they can support a form of technological innovation in which commercial objectives remain compatible with public-health responsibilities and identifiable organizational accountability.

4.2 The Role of Digital Platforms: Algorithmic Self-Regulation vs. National Regulatory Authorities

The governance of digital health communication is characterized by an operational tension between platform-based content governance and statutory enforcement by public regulatory authorities (Gomes & Sultan, 2024). As digital platforms have become increasingly important environments for the discovery and circulation of health-related information, technology companies have developed automated and human moderation systems intended to identify content that violates platform policies (Lim et al., 2022). These governance mechanisms operate within commercial environments in which user engagement, network activity, and advertising revenue constitute relevant organizational incentives. Such incentives may not always coincide with public-health priorities concerning scientific accuracy, appropriate risk communication, and patient protection, creating potential tensions between commercial platform objectives and health-information governance (Gomes & Sultan, 2024).

Commercial social-media platforms can employ automated moderation mechanisms including keyword filters, computer-vision classifiers, and risk-scoring systems to process large volumes of user-generated content (Gomes & Sultan, 2024). These technologies can support the identification of clearly defined policy violations at a scale that would be difficult to achieve through exclusively manual review. Their limitations become more apparent, however, where the regulatory or scientific significance of communication depends on context rather than on readily identifiable textual or visual features. Nuanced health claims, indirect commercial relationships, unbranded promotional communication, and context-dependent omissions may require forms of scientific, legal, or interpretative assessment that automated classification systems cannot consistently resolve on their own (Lim et al., 2022; Gomes & Sultan, 2024).

Public regulatory authorities, including the U.S. Food and Drug Administration and national European supervisory bodies, operate through a different institutional logic. Their activities are grounded in statutory mandates and formal procedures intended to evaluate regulated health communication according to applicable legal and evidentiary requirements (Directive 2001/83/EC, 2001; U.S. Food and Drug Administration, 2023). At the same time, the enforcement of such standards across decentralized and cross-border digital networks can present practical difficulties related to jurisdiction, identification of responsible actors, evidence gathering, and the time required for administrative or legal intervention (Marchant et al., 2011). Digital content may circulate rapidly across platforms and jurisdictions, while formal investigations, warning procedures, and legal proceedings necessarily operate through more structured institutional processes. As a result, potentially problematic communication may achieve substantial exposure before regulatory intervention occurs (Vosoughi et al., 2018).

This temporal and organizational asymmetry suggests that neither platform moderation nor retrospective public enforcement can independently address every risk associated with digital health communication (Gomes & Sultan, 2024; Verma et al., 2024). Platform governance offers scale and speed but remains limited in its ability to interpret complex scientific, commercial, and regulatory contexts, while public authorities retain formal legal powers but necessarily operate within jurisdictional and procedural constraints (Marchant et al., 2011). The resulting governance environment therefore requires complementary forms of responsibility rather than reliance on a single regulatory mechanism.

For life sciences organizations, this creates a rationale for incorporating preventive governance mechanisms earlier in the communication lifecycle. Internal scientific review, structured content controls, automated compliance pre-screening, escalation procedures, and proportionate human

oversight can complement both platform-level moderation and statutory regulation. Within the Data + Empathy framework developed in this thesis, computational systems may assist with high-volume monitoring, classification, evidence retrieval, and preliminary risk identification, while human-centered evaluation considers contextual appropriateness, stakeholder vulnerability, uncertainty, and the possible consequences of communication. Governance then determines whether, where, and under what conditions technological outputs should be translated into organizational action. Compliance-by-Design should therefore be understood as a supportive layer within a broader governance architecture rather than as a substitute for either professional judgment or public regulatory authority (Sabapathy, 2025; Verma et al., 2024).

4.2.1 Case Study: The Semaglutide Communication Crisis

The widespread circulation of GLP-1 (glucagon-like peptide-1) receptor agonist narratives across digital media provides a relevant illustrative case for examining contemporary challenges in healthcare communication governance. Semaglutide, marketed under brand names including Ozempic for Type 2 diabetes and Wegovy for weight management, became the subject of extensive online discussion involving pharmaceutical information, weight-loss narratives, personal experiences, telehealth services, influencer content, and commercial promotion. The resulting communication environment differs substantially from conventional pharmaceutical advertising, in which identifiable sponsors operate through formal Medical, Legal, and Regulatory review processes and are subject to defined requirements concerning scientific substantiation, promotional claims, and risk communication.

In digital environments, communication concerning GLP-1 therapies has circulated through a more decentralized network of actors, including healthcare professionals, patients, influencers, telehealth providers, online vendors, and compounding entities. This fragmentation complicates the identification of communicative responsibility because not all participants occupy the same legal, institutional, or commercial position. The significance of the case therefore lies less in a complete displacement of institutional oversight than in the coexistence of formally regulated pharmaceutical communication with high-volume decentralized content that may fall under different regulatory categories and levels of supervision.

From an analytical perspective, the rapid diffusion of semaglutide-related content also illustrates the operational differences between traditional broadcast communication and algorithmically mediated digital environments. Conventional pharmaceutical advertising generally relies on comparatively identifiable sponsors and planned distribution channels, whereas platforms such as TikTok and Instagram organize content through recommendation systems that respond to multiple engagement and contextual signals. Content involving visible transformations, personal testimonies, lifestyle narratives, or emotionally salient experiences may consequently receive substantial exposure when users interact with and redistribute it. However, such visibility should not be interpreted as evidence that recommendation algorithms systematically privilege promotional narratives over clinically reliable information. The distribution of health-related content results from the interaction among platform design, user behavior, content characteristics, network dynamics, and moderation practices.

Applying Chaiken's (1980) Heuristic-Systematic Model provides one interpretative perspective on this communication environment. Short-form and rapidly consumed content may create conditions

in which users rely more heavily on accessible cues such as creator familiarity, visual presentation, personal testimony, and social endorsement when evaluating information. This does not imply that social-media users inevitably process health information heuristically or that such cues replace systematic evaluation of clinical evidence. Rather, platform affordances and low-intent browsing contexts may increase the relevance of heuristic cues, particularly when detailed scientific information and risk disclosures are less immediately accessible.

Algorithmic Amplification vs. Traditional Communication Channels:

Traditional DTC Model:

Centralized Sponsor → Internal MLR Review → Regulated Promotional Communication → Planned Distribution

Algorithmically Mediated Network Model:

Decentralized Actors → Engagement and Contextual Signals → Recommendation Systems → Potentially Expanded Exposure and Heuristic Evaluation

The case also illustrates limitations associated with platform-based content governance. Commercial platforms use combinations of automated detection systems, policy rules, user reporting, and human moderation to identify content that may violate platform standards. These mechanisms can operate at considerable scale, but their effectiveness may be reduced where the problematic character of health communication depends on subtle contextual factors, undisclosed commercial relationships, indirect product references, or scientifically misleading implications rather than easily identifiable prohibited terms.

Similar challenges arise in relation to commercial disclosure. Mechanisms such as #ad or #sponsored labels can increase the visibility of material relationships when they are used clearly and appropriately. Their effectiveness may nevertheless be reduced when disclosures are difficult to notice, ambiguously presented, or integrated into content whose commercial and personal elements are difficult for users to distinguish. The resulting governance concern is not that platform disclosure mechanisms are inherently ineffective, but that transparency depends on both the existence and the practical visibility of the disclosure within the communication context.

The regulatory response to semaglutide-related communication further illustrates the complexity of applying established pharmaceutical and consumer-protection rules to decentralized digital environments. Under the Federal Food, Drug, and Cosmetic Act and related regulatory provisions, false or misleading promotional communication concerning prescription medicines may trigger regulatory scrutiny, while health-related commercial claims may also fall within the competence of other authorities according to the actor, product, and communication involved. Online vendors, telehealth intermediaries, influencers, manufacturers, and compounding entities should therefore not be treated as a single regulatory category, since the applicable obligations depend on their role and relationship to the communication.

Regulatory authorities have responded to aspects of the broader GLP-1 communication environment through public safety communications, enforcement actions, warning letters, and guidance concerning unapproved or misleadingly marketed products and health claims. The Food and Drug Administration has addressed risks associated with certain unapproved semaglutide

products and online distribution practices, while the Federal Trade Commission maintains substantiation and disclosure requirements relevant to commercial health communication. These interventions illustrate the continued relevance of statutory oversight even within decentralized digital environments.

Illustrative Regulatory Response Pathways:

Potential FD&C Act / FDA Issues → Misleading or Non-Compliant Drug-Related Communication
→ Regulatory Review / Warning or Safety Communication

FTC Act / Health-Claim Requirements → Insufficiently Substantiated Commercial Claims or Inadequate Disclosure → Consumer-Protection Enforcement

The semaglutide case can also be interpreted through the pacing-problem framework developed by Marchant et al. (2011). Digital communication can be produced, reproduced, and redistributed at a speed that differs substantially from the procedural timelines associated with formal investigation, evidentiary assessment, regulatory review, and enforcement. This temporal asymmetry does not demonstrate that regulatory institutions are incapable of responding effectively, nor does it establish that every widely circulated GLP-1 narrative preceded regulatory intervention. It nevertheless illustrates how rapidly evolving communication practices can create periods in which potentially problematic content achieves substantial visibility before formal institutional responses are completed.

The broader market context surrounding GLP-1 therapies also demonstrates why causal interpretation requires caution. Increased public attention to semaglutide coincided with strong demand, supply pressures, extensive media coverage, off-label interest, telehealth expansion, and the emergence of unapproved or compounded alternatives. Digital communication formed part of this environment, but these developments cannot be attributed exclusively to recommendation algorithms or to delays in regulatory intervention. The case is therefore better understood as illustrating the interaction of technological, commercial, medical, and institutional factors rather than as demonstrating a single causal pathway from algorithmic amplification to market disruption.

From a governance perspective, the principal implication is that platform moderation and retrospective public enforcement perform important but necessarily different functions. Neither mechanism can independently govern every stage of decentralized health communication. Platforms can support high-volume detection and distribution controls, while public authorities provide formal legal interpretation and enforcement. Pharmaceutical and life sciences organizations consequently have an additional responsibility to govern the communication activities, partners, and digital systems that fall within their organizational control rather than assuming that external moderation or subsequent enforcement will identify every potential risk.

This case therefore provides a useful transition toward the Compliance-by-Design framework examined in the following section. The objective of proactive internal governance is not to create an automated mechanism capable of guaranteeing regulatory compliance before publication. Instead, computational tools can support evidence retrieval, semantic screening, anomaly detection, and preliminary risk identification, while organizational governance determines when additional scientific, legal, or human review is required. Within the broader Data + Empathy framework, the

relevant principle remains risk-proportionate functional asymmetry: technological capabilities can assist with scalable monitoring and prescreening, whereas the level of human involvement should increase with scientific ambiguity, regulatory uncertainty, stakeholder vulnerability, and the potential consequences of communication.

4.3 Compliance by Design: Leveraging AI as a Pre-Screening Tool for Regulatory Alignment

The operational limitations of platform-based content moderation and retrospective legal enforcement, as examined in the preceding sections, support the development of additional preventive safeguards within pharmaceutical digital communication (Gomes & Sultan, 2024; Marchant et al., 2011). Rather than relying exclusively on external platform filters or post-publication administrative intervention, life sciences organizations can incorporate internal review mechanisms earlier in the content lifecycle (Sabapathy, 2025). Within this thesis, Compliance by Design is proposed as a conceptual governance approach in which computational tools assist with preliminary screening, evidence retrieval, semantic analysis, and risk identification before communication assets proceed to formal human review.

One possible implementation involves the use of a secondary Large Language Model as an evaluator of outputs produced by a primary generation system (Zheng et al., 2023). The evaluator can be architecturally separated from the generating model and configured with access to relevant guidance, internal policies, structured compliance criteria, and approved evidence repositories. Such separation may introduce an additional layer of scrutiny, but it should not be assumed to eliminate correlated errors or shared model biases. Models developed from similar architectures, training distributions, or reasoning patterns may still reproduce related weaknesses, making independent human and organizational oversight necessary.

Within the proposed architecture, generated advertising variations or omnichannel communication assets can be routed to the evaluator for preliminary semantic screening. The system may identify the presence of benefit statements, associated risk information, unsupported claims, missing contextual elements, or inconsistencies with predefined organizational rules. In this role, semantic similarity, textual proximity, and structured rule matching can operate as screening indicators that help identify assets requiring further review. They should not, however, be interpreted as direct computational measurements of legal fair balance, since regulatory adequacy depends on contextual, scientific, and qualitative considerations that cannot be reduced to semantic distance alone (U.S. Food and Drug Administration, 2023; Directive 2001/83/EC, 2001; Zheng et al., 2023).

Compliance-by-Design Pre-Screening Architecture

[Compiled Digital Assets] → [Secondary Evaluator Model]

↓

Preliminary Risk Assessment

→ Potential concern identified

- Diagnostic metadata attached
- Asset flagged or temporarily withheld

- Routed for revision and/or enhanced human review

→ No major predefined concern identified

- Preliminary screening completed
- Prioritized routing to human MLR review

Within this architecture, confidence or risk scores may be used as internal prioritization mechanisms rather than as estimates of legal compliance. A lower score can indicate that an asset contains features requiring additional review, while a higher score can indicate that fewer predefined concerns were detected during automated screening. These scores therefore support triage rather than determining whether a communication is legally compliant. Where appropriate, diagnostic feedback generated by the evaluator may also support iterative content refinement, including feedback-based optimization approaches such as Reinforcement Learning from AI Feedback (Bai et al., 2022). Such mechanisms remain supportive and experimental rather than substitutes for professional regulatory assessment.

By filtering and prioritizing content before formal review, AI-assisted pre-screening can potentially reduce part of the repetitive workload faced by compliance professionals and direct attention toward assets presenting greater scientific, regulatory, or contextual uncertainty (OECD, 2023; Sabapathy, 2025). The purpose of the dual-model architecture is therefore not to automate approval, but to support scalable risk detection while preserving human responsibility for final interpretation and decision-making.

Table 4.1 Comparative Governance Models for Pharmaceutical Digital Communication

Governance Model	Primary Objective	Strengths	Operational Limitations	Pharmaceutical Implications
Platform Self-Regulation	Enforce platform policies, manage user safety, and moderate large volumes of digital content.	High-scale moderation, rapid processing, and automated detection of defined textual, visual, or behavioral signals (Gomes & Sultan, 2024).	Limited ability to interpret complex scientific context, indirect promotional intent, nuanced risk communication, or ambiguous commercial relationships (Gomes & Sultan, 2024; Lim et al., 2022).	Can complement other safeguards but should not be treated as the primary mechanism for pharmaceutical regulatory review.
National Regulatory Authorities	Protect public health and enforce applicable statutory requirements for medicinal-product and health-related communication.	Formal legal authority, investigative powers, regulatory guidance, and enforceable sanctions (Directive 2001/83/EC; U.S. Food and Drug Administration, 2023).	Formal procedures, jurisdictional constraints, and potential temporal asymmetry between digital circulation and regulatory intervention (Marchant et al., 2011; Vosoughi et al., 2018).	Provides authoritative legal oversight but may be complemented by preventive organizational controls.
Compliance by Design	Integrate computational pre-screening, structured rules, evidence retrieval, and semantic analysis into internal communication workflows.	Scalable preliminary screening, risk identification, structured triage, and earlier detection of potentially problematic assets (Bai et al., 2022; Zheng et al., 2023).	Residual hallucination risk, model bias, interpretability limitations, false positives and false negatives, and inability to independently determine legal compliance (Ji et al., 2023; Sabapathy, 2025).	Can support prioritization and early risk detection before formal MLR review, without replacing professional judgment.

Governance Model	Primary Objective	Strengths	Operational Limitations	Pharmaceutical Implications
Human-in-the-Loop Governance	Provide contextual, scientific, ethical, and regulatory evaluation where automated assessment is insufficient.	Supports qualitative interpretation, contextual judgment, scientific validation, accountability, and evaluation of stakeholder implications (OECD, 2023; Chantre et al., 2025).	Higher resource requirements, slower review processes, reviewer fatigue, and operational complexity (Jones et al., 2024; Sabapathy, 2025).	Maintains identifiable organizational responsibility and provides review for issues that cannot be reliably resolved through computational screening alone.

Analytical Discussion: Toward a Layered Governance Architecture

The comparative characteristics summarized in Table 4.1 indicate that no single governance mechanism is capable of addressing all the scientific, regulatory, operational, and ethical challenges associated with pharmaceutical digital communication. Platform moderation offers scalability and speed but remains limited when interpretation depends on scientific or commercial context. Public regulatory authorities possess formal legal authority but necessarily operate through institutional procedures and jurisdictional frameworks. Computational pre-screening can identify predefined patterns and potential risks at scale, while remaining vulnerable to model error, incomplete context, and interpretability limitations (Gomes & Sultan, 2024; Marchant et al., 2011; Ji et al., 2023; Sabapathy, 2025).

These limitations support a layered governance architecture combining technical, legal, organizational, and human safeguards throughout the communication lifecycle. Within such an architecture, Compliance by Design does not replace statutory regulation or professional review. Instead, computational mechanisms can support evidence retrieval, semantic screening, rule-based checks, and preliminary identification of potentially problematic assets before publication (Zheng et al., 2023; Verma et al., 2024). Their contribution lies in assisting prioritization and early risk detection rather than determining whether communication is legally compliant.

Human oversight remains particularly important where communication involves scientific ambiguity, complex benefit-risk interpretation, vulnerable audiences, uncertain evidence, or significant clinical consequences. AI-assisted auditing systems can support consistency checks and identify explicit or predefined concerns, but they do not possess legal authority, moral agency, or independent contextual understanding (OECD, 2023; Zheng et al., 2023). Medical, Legal, and Regulatory professionals therefore remain responsible for interpreting scientific evidence, assessing applicable regulatory requirements, and deciding whether communication is appropriate for publication.

Importantly, the role assigned to Artificial Intelligence within this governance model should not be interpreted as a permanent separation between an invisible computational backend and an exclusively human-facing communication layer. AI visibility, autonomy, and human supervision should instead vary according to the function being performed and the associated level of risk. In standardized and comparatively low-risk activities, computational systems may operate with greater autonomy or

visibility, provided that appropriate transparency and escalation mechanisms are maintained. As scientific uncertainty, stakeholder vulnerability, emotional sensitivity, regulatory ambiguity, or potential clinical consequences increase, the need for meaningful human involvement correspondingly increases.

This risk-proportionate functional asymmetry connects Compliance by Design with the broader Data + Empathy framework developed in this thesis. Data-driven systems provide computational capabilities for retrieval, analysis, classification, and preliminary risk detection. Human-centered evaluation considers whether the resulting technological possibilities are appropriate in relation to stakeholder context, vulnerability, autonomy, and informational need. Governance then determines whether, where, and under what conditions those capabilities should be translated into organizational action.

Accordingly, sustainable governance in pharmaceutical digital communication should not rely exclusively on platform moderation, retrospective regulatory enforcement, automated compliance screening, or human review considered in isolation. A more robust approach combines these mechanisms according to their respective functions and limitations, preserving identifiable human and organizational accountability while using computational systems where they can provide meaningful analytical support.

4.4 Ethical Dilemmas: Algorithmic Transparency and the Risk of AI-Generated Misinformation

The strategic integration of generative language architectures, predictive recommendation engines, and automated decision systems into pharmaceutical communication, while offering significant operational capabilities, introduces ethical dilemmas that center on algorithmic transparency, information integrity, and public trust (Amini Khiabani & Rezaeian, 2026; Sabapathy, 2025). Within the healthcare domain, where communication may influence health-related perceptions, therapeutic adherence, and interactions with healthcare professionals, the deployment of Artificial Intelligence cannot be evaluated solely through metrics of computational efficiency or engagement optimization (Jones et al., 2024; Verma et al., 2024). Instead, the technological transformation of life sciences outreach creates a complex ethical landscape in which the pursuit of message relevance intersects with the necessity of preserving human agency, scientific accuracy, transparency, and informed choice (Amini Khiabani & Rezaeian, 2026).

At the core of the ethical tension surrounding generative Artificial Intelligence in health media is the phenomenon of Natural Language Generation hallucination, whereby probabilistic language models may generate fluent and apparently plausible statements that are nevertheless clinically unsupported, inaccurate, or fabricated (Ji et al., 2023). Because linguistic fluency does not constitute evidence of factual validity, Large Language Models can articulate inaccurate scientific claims in a form that appears authoritative to the recipient (Ji et al., 2023). Within general consumer advertising, an incorrect descriptive output may primarily affect communication quality; however, within the life sciences sector, unsupported statements regarding drug efficacy, dosage instructions, contraindications, or therapeutic side effects may generate substantially more serious consequences for health-related understanding and decision-making (Ji et al., 2023; Liang & Mackey, 2011).

This ethical vulnerability is heightened by the persuasive characteristics of modern synthetic media. Natural language models can reproduce sophisticated scientific terminology, professional linguistic

structures, and empathetic forms of expression that resemble conventional medical communication (Ji et al., 2023). When digital consumers or healthcare professionals encounter algorithmically generated content that lacks adequate source grounding, the apparent credibility of the language may make it more difficult to distinguish scientifically supported information from inaccurate or unsupported generation (Ji et al., 2023; Liang & Mackey, 2011). While technical frameworks such as Retrieval-Augmented Generation and semantic Knowledge Graphs can condition language models on controlled enterprise repositories and structured evidence, they do not eliminate the residual risks of retrieval error, semantic drift, incorrect synthesis, or inappropriate contextual adaptation (Gao et al., 2023; Lewis et al., 2020; Sabapathy, 2025). Automated generation therefore cannot be treated as an autonomous substitute for proportionate medical, regulatory, and professional review.

A secondary ethical dilemma involves the limited interpretability of deep-learning architectures, frequently characterized in technology-governance literature through the concept of the black-box problem (Amini Khiabani & Rezaeian, 2026). As multi-layer neural networks process high-velocity streams of behavioral and contextual data to support Next Best Action recommendations or dynamically configure communication assets, the relationship between specific input variables and the resulting strategic recommendation may remain difficult to interpret or communicate clearly to end users and regulatory reviewers (Sabapathy, 2025). This limited transparency creates a structural tension with ethical principles of health communication that emphasize informed consent, accountability, transparency, and respect for individual autonomy (Amini Khiabani & Rezaeian, 2026).

When users recognize that their visual interface or content feed has been dynamically adapted without understanding how their personal or health-related data were collected and processed, the resulting personalization may generate uncertainty or psychological reactance (Chantre et al., 2025; Mansfield et al., 2025). This possibility is particularly relevant among digitally experienced audiences, including Generation Z, although familiarity with algorithmically curated environments should not be interpreted as evidence of a universal ability to identify or understand personalization mechanisms. Perceived opacity in the underlying tracking infrastructure may transform otherwise useful contextual relevance into a perception of excessive surveillance or intrusive profiling (Amini Khiabani & Rezaeian, 2026; Chantre et al., 2025). When personalization operates without adequate explanations of data provenance and purpose, users may question whether sensitive health interests are being used primarily to improve informational support or to optimize commercial persuasion, potentially weakening digital trust (Amini Khiabani & Rezaeian, 2026; Verma et al., 2024).

The ethical challenges surrounding algorithmic opacity and content validity are further amplified by the speed at which synthetic and misleading information can circulate across decentralized digital networks (Chou et al., 2018; Vosoughi et al., 2018). Social-media environments can facilitate the rapid diffusion of emotionally salient or controversial information, creating conditions in which unsupported health narratives, manipulated media, and misleading medical advice may reach substantial audiences before their reliability has been adequately assessed (Chou et al., 2018; Vosoughi et al., 2018). The spread of such material should not be attributed exclusively to recommendation algorithms, since diffusion reflects the interaction of platform architecture, network dynamics, user behavior, content characteristics, and moderation practices. Nevertheless, generative technologies can reduce the cost and increase the speed of producing plausible synthetic content,

potentially increasing the volume of information requiring verification. By the time institutional authorities, platforms, or pharmaceutical organizations identify misleading claims, those narratives may already have circulated widely, reinforcing the limitations of exclusively retrospective intervention (Gomes & Sultan, 2024; Marchant et al., 2011; Vosoughi et al., 2018).

To address these ethical dilemmas without suppressing technological innovation, this thesis argues that life sciences organizations should move beyond passive legal compliance toward an active framework of algorithmic responsibility anchored in the Data + Empathy paradigm (Sabapathy, 2025; Verma et al., 2024). Within this conceptual framework, Artificial Intelligence is not assigned a universally fixed position within an invisible backend, nor is visible communication categorically reserved for human actors. Instead, the allocation of computational and human functions should follow the principle of risk-proportionate functional asymmetry. Artificial Intelligence can perform comparatively autonomous tasks where communication is standardized, bounded, transparent, and associated with limited consequences, while meaningful human involvement becomes progressively more important as interactions become clinically sensitive, emotionally complex, uncertain, or consequential. Data create computational possibility; empathy evaluates that possibility in relation to human context; and governance determines whether, where, and under what conditions it should become organizational action. By integrating transparent data practices, proportionate technological role allocation, and meaningful professional accountability, life sciences organizations can reduce the risks associated with algorithmic opacity and AI-generated misinformation while supporting innovation consistent with public-health protection and ethical integrity.

4.5 Toward Responsible AI Governance in Pharmaceutical Digital Communication

The analysis conducted throughout Chapter 4 demonstrates that the governance of AI-driven pharmaceutical digital communication cannot be effectively addressed through isolated technological solutions or exclusively reactive policy measures. Central to this challenge is the recognition that regulatory lag is likely to remain a recurring structural feature of the digital health ecosystem (Marchant et al., 2011). Advances in machine learning, algorithmic recommendation systems, and generative language architectures can develop and diffuse across commercial environments considerably faster than formal legislative and regulatory processes are able to adapt (Marchant et al., 2011; Sabapathy, 2025). This pacing problem does not imply that existing legal frameworks become irrelevant or ineffective, but rather that periods of temporal misalignment may emerge in which novel marketing practices, influencer collaborations, and algorithmic personalization techniques develop before specific regulatory guidance or enforcement practices have fully adapted to them (Marchant et al., 2011; Verma et al., 2024).

Within this environment, commercial platform self-regulation represents an important but incomplete layer of governance (Gomes & Sultan, 2024). Social-media platforms operate through business models in which engagement, interaction, advertising revenue, and user retention constitute significant organizational incentives. These commercial structures can create tensions between content moderation objectives and broader platform-growth priorities, although they should not be interpreted as evidence that moderation systems are categorically designed to preserve harmful or misleading content. Automated platform-governance mechanisms may nevertheless face substantial difficulties when evaluating nuanced health communication, including implicit promotional claims, contextual omissions, borderline disease-awareness content, and communications whose regulatory

meaning depends on the relationship among multiple elements rather than on the presence of individual prohibited terms (Gomes & Sultan, 2024).

These operational limitations reinforce rather than diminish the role of national and supranational regulatory authorities. Public supervisory bodies, including the U.S. Food and Drug Administration and competent European authorities, possess statutory enforcement powers and public-health mandates that private technology platforms do not replicate (Directive 2001/83/EC; U.S. Food and Drug Administration, 2023). Although regulatory agencies face practical difficulties in continuously monitoring high-volume and decentralized digital environments, their legal authority establishes mandatory boundaries for pharmaceutical promotion and provides institutional mechanisms for investigation, enforcement, and corrective action. Responsible governance therefore requires interaction between public regulation and organizational safeguards rather than the substitution of one by the other (Marchant et al., 2011; Sabapathy, 2025).

To reduce the temporal distance between high-speed digital communication and retrospective external oversight, this thesis argues that Compliance by Design should be interpreted as an enterprise-wide governance approach rather than merely as a technical software mechanism (Sabapathy, 2025; Verma et al., 2024). Under this interpretation, legal requirements, scientific evidence, risk-management protocols, structured content rules, and automated pre-screening mechanisms are incorporated earlier within the content-development lifecycle. Knowledge Graph constraints, deterministic rule systems, Retrieval-Augmented Generation, and evaluator models may contribute to identifying potentially problematic content before public deployment, while structured metadata can support the traceability of claims, evidence, audience conditions, and approval status (Zheng et al., 2023). These mechanisms should be understood as tools for risk reduction and prioritization rather than as systems capable of autonomously determining legal compliance.

Accordingly, automated pre-screening does not eliminate the requirement for professional oversight. Probabilistic language models remain susceptible to hallucination, retrieval error, semantic drift, contextual misinterpretation, and limitations in model interpretability (Ji et al., 2023; Sabapathy, 2025). Even when evaluator systems are trained or configured with domain-specific criteria, their outputs remain dependent on the quality of their source material, evaluation rubrics, technical implementation, and underlying model architecture. They should therefore not be relied upon to resolve independently complex legal ambiguity, scientific uncertainty, or ethically sensitive communication contexts. Human compliance professionals, medical specialists, legal experts, and communication practitioners remain necessary for evaluating scientific validity, contextual appropriateness, stakeholder vulnerability, and consequential decisions for which institutional accountability must remain clearly identifiable (OECD, 2023; Sabapathy, 2025).

Reconnecting with the core contribution of this research, the Data + Empathy framework can consequently be understood as an overarching governance philosophy rather than solely as a tactical communication model. Its central proposition is not that Artificial Intelligence should remain permanently confined to an invisible computational backend, but that technological roles should be allocated according to their function, sensitivity, uncertainty, and potential consequences. In comparatively standardized and low-risk contexts, AI systems may perform more direct informational, analytical, or administrative functions. As interactions become more clinically

sensitive, emotionally complex, legally ambiguous, or consequential, the requirement for meaningful human involvement increases.

This risk-proportionate functional asymmetry replaces the rigid distinction between a computational backend and an exclusively human front-end. Human contribution extends beyond visible empathetic dialogue and includes the evaluation of whether a proposed technological use is proportionate, whether available data should be used at all, whether an automated recommendation requires escalation, and whether the resulting communication remains appropriate to the stakeholder and context. Data create computational possibility; empathy evaluates that possibility in relation to human needs and vulnerability; governance determines whether, where, and under what conditions the resulting action should be permitted. The objective is therefore not to make Artificial Intelligence more human in appearance, but to make its organizational deployment more human-centered.

Critically evaluated, the future development of pharmaceutical AI governance depends on recognizing that statutory regulation, platform moderation, Compliance by Design, and Human-in-the-Loop oversight perform complementary rather than interchangeable functions (Gomes & Sultan, 2024; Marchant et al., 2011; Sabapathy, 2025). Statutory authorities establish enforceable legal boundaries; platforms govern access and distribution within their infrastructures; enterprise systems can embed preventive safeguards into content workflows; and human professionals provide contextual interpretation, accountability, and escalation where automated systems are insufficient. Sustainable digital transformation therefore requires organizations to move beyond compliance understood as a final checkpoint and toward governance structures in which transparency, scientific accuracy, privacy, and accountability are considered throughout the entire technological lifecycle (Amini Khiabani & Rezaeian, 2026; Verma et al., 2024).

Within this perspective, public trust and organizational credibility are not guaranteed by either technological sophistication or the presence of human oversight alone. They are more plausibly supported when stakeholders can understand the role of Artificial Intelligence, when personalization remains proportionate to its informational purpose, when scientific claims remain traceable to appropriate evidence, and when clear mechanisms of accountability are maintained across automated and human decision-making processes (OECD, 2023; Sabapathy, 2025). Transparent governance therefore functions not simply as a regulatory constraint on innovation, but as one of the conditions under which AI-mediated pharmaceutical communication can remain socially and institutionally credible.

Ultimately, the conceptual model proposed in this chapter demonstrates that responsible AI governance in life sciences communication is simultaneously a technological, organizational, regulatory, and strategic challenge. Architectures such as Retrieval-Augmented Generation, structured Knowledge Graphs, automated pre-screening, and Next Best Action systems can increase the capacity to manage evidence, scale communication, and identify potential risks earlier in the content lifecycle, but their effectiveness depends on how pharmaceutical organizations integrate these tools within broader governance processes. The transition toward responsible AI therefore shifts attention away from the maximization of automation and toward the appropriate allocation of technological roles under transparent and accountable institutional conditions. This perspective provides the conceptual bridge to the empirical analysis developed in Chapter 5, where selected case

studies and field research are used not to validate the framework conclusively, but to examine how its central principles are reflected, challenged, and refined within real-world communication contexts.

Chapter 5: Empirical Analysis: Case Studies and Field Research

The theoretical and normative frameworks developed in the preceding chapters indicate that the integration of Artificial Intelligence into healthcare communication involves an ongoing negotiation among technological scalability, regulatory compliance, and stakeholder trust (Jones et al., 2024; Sabapathy, 2025). While computational architectures such as Retrieval-Augmented Generation, Dynamic Creative Optimization, and predictive Next Best Action engines provide significant operational capabilities, their real-world deployment is mediated by the cognitive expectations of digital-native cohorts and the institutional boundaries established by regulatory frameworks (Marchant et al., 2011; Verma et al., 2024).

To transition from conceptual governance models to real-world applications, this section conducts a comparative analysis of three international communication initiatives:

1. **Eurofarma's *Scrolling Therapy*** (developed with Dentsu Creative, 2023);
2. **The American Society of Clinical Oncology's (ASCO) *The Most Beautiful Sound*** (developed with Grey Health & Wellness, 2023); and
3. **Dove's *The Code*** (developed with Ogilvy, 2024).

These initiatives were identified through purposive sampling to construct a balanced comparative matrix across healthcare and adjacent wellness domains. Rather than analyzing conventional promotional advertising, the selection examines three distinct configurations of computational and communicative technology:

- **Technology-enabled rehabilitation and patient adherence (*Scrolling Therapy*)**, in which machine learning operates as a supportive rehabilitation tool;
- **Sensory data translation for supportive care (*The Most Beautiful Sound*)**, in which complex biophysical data are translated into an acoustic resource intended to provide a more tangible and emotionally meaningful representation of therapeutic processes; and
- **Voluntary boundary-setting against algorithmic visual representation (*The Code*)**, which provides an external benchmark for the ethical governance of generative AI and authentic human representation.

Importantly, Dove is not a pharmaceutical manufacturer, and *The Code* is not a medicinal product campaign. It is incorporated here as an adjacent-sector benchmark from the personal care and wellness industry because its publicly stated commitment regarding the use of generative AI in representations of real women provides a relevant external reference point for examining algorithmic bias, aesthetic homogenization, and authentic human representation. These issues are also relevant to younger digital audiences and to health-related communication environments in which credibility, identity, and psychological well-being are increasingly influenced by digitally mediated representations (Mansfield et al., 2025).

Methodologically, these three cases are not presented as empirical proof or formal validation of the Data + Empathy framework. Because case studies of this nature reflect specific creative, organizational, and technological contexts, they cannot establish generalizable causal laws. Instead, they serve as purposively selected real-world applications used to illustrate, examine, and critically

interpret how different configurations of computational capability, human-centered communication, and ethical governance can be understood through the Data + Empathy framework.

Where analytically appropriate, the cases are examined through the theoretical lenses of the Heuristic-Systematic Model (Chaiken, 1980), Algorithm Aversion (Dietvorst et al., 2015), and the Uncanny Valley of Mind (Stein & Ohler, 2017). This comparative approach enables an assessment of the communicative mechanisms that may support or undermine stakeholder trust in technology-mediated health environments.

5.1 Comparative Analysis of Selected AI-Driven Initiatives

Purposive Sampling and Empirical Scope Matrix

Initiative & Organization	Sector & Category Classification	Analytical Configuration Examined
Eurofarma: <i>Scrolling Therapy</i> (Dentsu Creative, 2023)	Pharmaceutical/ Neurological Care	Functional therapeutic utility: technology-enabled rehabilitation supported by computer vision
ASCO: <i>The Most Beautiful Sound</i> (Grey Health & Wellness, 2023)	Healthcare / Clinical Oncology and Supportive Care	Sensory data translation: biophysical information translated into an acoustic experience
Dove: <i>The Code</i> (developed with Ogilvy; Dove, 2024)	Adjacent Benchmark / Personal Care and Wellness	Ethical boundary-setting: critical reflection on synthetic human representation and generative AI

5.1.1 Case 1: Eurofarma - Scrolling Therapy (Functional AI as a Clinical Rehabilitation Utility)

In 2023, the pharmaceutical company Eurofarma, in partnership with Dentsu Creative and with institutional support from patient advocacy organizations including the Associação Brasil Parkinson, introduced *Scrolling Therapy*. The initiative was developed as an experimental mobile application designed to address hypomimia—commonly referred to as “facial masking”—in individuals living with Parkinson’s disease (Dentsu Creative, 2023). Rather than approaching Artificial Intelligence primarily as a mechanism for commercial targeting, automated content generation, or promotional optimization, the project applied computational technology to a specific behavioral and rehabilitative challenge within the patient journey.

Hypomimia is a frequent neurological manifestation of Parkinson’s disease and is characterized by diminished spontaneous facial expression and reduced facial mobility. Beyond its motor manifestations, reduced facial expressivity can influence interpersonal communication because facial movements constitute an important component of non-verbal social interaction (Dentsu Creative, 2023; Tickle-Degnen et al., 2014). Consequently, the relevance of hypomimia extends beyond the mechanical execution of facial movements: difficulties in expressing recognizable emotional cues may affect how patients communicate with caregivers, family members, healthcare professionals, and

their broader social environment. This communicative dimension makes the condition particularly relevant to the conceptual perspective adopted in this thesis, in which technological effectiveness is evaluated not exclusively in terms of computational performance but also according to its capacity to respond to identifiable human needs.

From a clinical rehabilitation perspective, managing facial stiffness may involve consistent facial exercises, including repeated smiling, eyebrow elevation, and exaggerated facial movements. According to the campaign documentation, adherence to these exercises represents a substantial challenge, with approximately 3% of patients reported to maintain regular practice (Dentsu Creative, 2023). A behavioral obstacle identified in the project's background is the emotional burden associated with conventional mirror-based exercises: repeatedly performing facial movements in front of a mirror may require patients to confront visible manifestations of their condition, potentially contributing to emotional distress and discontinuation of the routine (Dentsu Creative, 2023).

This difficulty reveals an important distinction between the theoretical availability of a therapeutic activity and its practical integration into everyday patient behavior. A rehabilitation exercise may be clinically relevant while simultaneously generating psychological or behavioral friction that limits consistent execution. From this perspective, the communication and design challenge is not simply to inform patients that an exercise should be performed, but to restructure the interaction through which the patient experiences that activity. *Scrolling Therapy* addresses this tension by attempting to integrate the rehabilitative movement into an existing digital habit rather than asking the individual to create an entirely separate behavioral routine.

Behavioral Adherence Barrier in Conventional Rehabilitation:

[Mirror-Based Facial Exercises] → [Direct Confrontation with Impairment] → [Emotional Distress] → [Potential Discontinuation]

Scrolling Therapy Behavioral Reframing:

[Facial Movement Tracking] → [On-Device Computer Vision] → [Command Translation] → [Habitual Social Browsing]

The contrast between these two pathways is analytically significant. In the conventional pathway, the therapeutic exercise itself constitutes the primary object of the patient's attention. In the alternative configuration proposed by *Scrolling Therapy*, the exercise becomes functionally embedded within another activity. The technological intervention therefore does not eliminate the rehabilitative task; rather, it modifies the context in which that task is performed. This distinction provides an example of how digital innovation can intervene at the level of behavioral experience rather than merely increasing the quantity of information delivered to patients.

To address this adherence challenge, *Scrolling Therapy* converts facial rehabilitation exercises into digital navigation inputs for social media platforms, specifically Facebook and Instagram (Dentsu Creative, 2023). Built on machine-learning and computer-vision technologies, the mobile application uses the smartphone's front-facing camera to track facial movements, mapping approximately 10,000 points across 43 facial muscles (Dentsu Creative, 2023). In this configuration, computer vision operates as an intermediary between physical movement and digital interaction: facial expressions

that would otherwise constitute isolated rehabilitation exercises acquire an additional functional purpose within the interface.

The application translates five core facial gestures into specific social media actions: a sustained smile functions as a “like” command; a surprised expression involving eyebrow elevation triggers a downward scroll; a disgusted facial gesture initiates an upward scroll; a sad facial expression activates video playback; and an angry expression executes a swipe (Dentsu Creative, 2023). The system also provides visual feedback intended to help users perform the required facial movements while interacting with their digital feeds.

This gesture-to-command architecture represents more than a technical interface modification. It transforms the relationship between therapeutic repetition and digital behavior by assigning immediate functional consequences to movements that might otherwise be perceived exclusively as rehabilitation exercises. The patient's interaction with the technology is therefore structured around an activity already present in contemporary digital routines, while the computational layer translates that interaction into a rehabilitative opportunity. The strategic relevance of this configuration lies in its attempt to reduce the experiential distance between healthcare-related behavior and everyday life.

From a data-governance and privacy perspective, the technical architecture relies on localized, on-device processing (Dentsu Creative, 2023). Rather than requiring facial movement data to be processed through an external cloud-based architecture, the computer-vision system operates locally on the user's mobile device. This architectural decision is particularly relevant because facial movement analysis may involve information that users could reasonably perceive as sensitive within a health-related context.

While this design choice should not be conflated with formal regulatory certification under a specific statutory framework, it represents a privacy-oriented approach consistent with the data-minimization principle established by the General Data Protection Regulation (Regulation (EU) 2016/679). By processing facial movement data locally on the user's device, the architecture reduces the need to transmit potentially sensitive information to external servers. In strategic terms, this configuration also illustrates how privacy considerations can be incorporated at the level of technological architecture rather than addressed exclusively through downstream disclosure mechanisms or consent interfaces.

On-Device Computational Pathway:

[Device Camera Input] → [Local Machine Learning / Computer Vision] → [On-Device Command Translation] → [Social Media Navigation]

The relevance of this architecture to pharmaceutical digital communication becomes clearer when considered against the broader personalization models examined in Chapters 2 and 3. Many contemporary AI-driven marketing architectures derive strategic value from collecting, aggregating, and interpreting increasingly granular behavioral data. *Scrolling Therapy* represents a substantially different configuration: computational sophistication is directed toward enabling a functional interaction rather than expanding the informational profile available for commercial persuasion. The

case therefore provides a useful counterpoint to data-intensive personalization strategies in which greater technological capability is frequently associated with greater extraction and analysis of user information.

When interpreted through the conceptual lens of this thesis, *Scrolling Therapy* provides a relevant real-world illustration of the Data + Empathy framework. The initiative positions computational technology not as a front-facing commercial instrument designed to maximize advertising exposure or execute predictive product targeting, but as an enabling functional utility. The computational system performs the technical tasks of real-time facial movement processing and gesture translation, while the user experience remains centered on integrating rehabilitation into a familiar everyday activity.

The “Data” component of the framework is visible in the computer-vision architecture, which transforms facial movement into structured computational inputs and subsequently into interface commands. The “Empathy” dimension does not depend on an artificial system attempting to reproduce human emotion through synthetic dialogue. Instead, it emerges from the design decision to address a concrete source of friction in the patient experience: the difficulty of sustaining repetitive rehabilitation while confronting the visible manifestations of a progressive condition. In this sense, empathy is operationalized through the structure of the interaction rather than simulated through anthropomorphic AI behavior.

This distinction is particularly important to the argument developed throughout this thesis. Human-centered AI does not necessarily require the technological system itself to appear human, emotionally expressive, or conversational. A computational architecture may support an empathetic objective precisely by remaining functionally limited and unobtrusive. *Scrolling Therapy* therefore provides an example of how empathy can be embedded into technological design through an understanding of behavioral context, rather than generated synthetically at the visible communication interface.

Viewed through the theoretical lens of Algorithm Aversion (Dietvorst et al., 2015), the design of *Scrolling Therapy* also offers an interesting contrast to contexts in which users are required to rely directly on algorithmic judgment. Rather than positioning the computational system as an autonomous decision-maker, the application assigns it a narrowly defined assistive function: translating facial movements into navigation commands. The algorithm does not diagnose the patient, determine a therapeutic pathway, or attempt to replace the authority of healthcare professionals.

Consequently, the case should not be interpreted as empirical evidence that *Scrolling Therapy* eliminates algorithm aversion. No such causal conclusion can be established from the campaign itself. Instead, its architecture illustrates a configuration in which some of the conditions commonly associated with resistance to algorithmic judgment are less central to the interaction. The user is not required to surrender a consequential clinical decision to an opaque automated system; computational capability is instead subordinated to a transparent and identifiable functional task. This distinction supports the broader proposition developed in this thesis that the role assigned to Artificial Intelligence may be as important as its technical sophistication in shaping stakeholder perceptions.

Furthermore, viewed through the Heuristic-Systematic Model (Chaiken, 1980), the initiative can be interpreted as a reframing of the cognitive demands associated with rehabilitation. Rather than presenting facial exercises as a separate therapeutic task requiring sustained deliberate attention,

Scrolling Therapy embeds them within an already familiar digital activity. This integration may reduce the perceived effort associated with the exercise while preserving its functional purpose.

The relevance of this reframing should not be interpreted as evidence that social media browsing is inherently or exclusively heuristic. Rather, the theoretical value of the Heuristic-Systematic Model in this context lies in distinguishing between different levels of cognitive effort and engagement. By incorporating rehabilitation into an existing behavioral routine, the application attempts to reduce the additional cognitive and motivational burden associated with initiating a separate exercise session. The technology therefore operates not by persuading the patient through increasingly sophisticated messaging, but by altering the environment in which the desired behavior occurs.

From a strategic communication perspective, this represents a meaningful departure from conventional models of pharmaceutical digital outreach. Rather than beginning with the question of how Artificial Intelligence can increase message relevance, conversion probability, or promotional efficiency, *Scrolling Therapy* begins with an identifiable patient problem and subsequently assigns computational technology a specific functional role in addressing it. The direction of the relationship is therefore reversed: patient need determines technological application rather than technological capability determining the communication strategy.

By anchoring rehabilitation to an existing digital routine, the initiative illustrates how pharmaceutical organizations can create patient value without relying exclusively on overt promotional messaging (Verma et al., 2024). It also suggests that brand-related value in healthcare may emerge indirectly through the provision of useful services and experiences rather than solely through persuasive communication. Importantly, the case itself does not establish a measurable long-term increase in brand equity; instead, it provides a strategic example of how utility-oriented innovation can contribute to a broader patient-centered positioning.

In strategic terms, *Scrolling Therapy* therefore suggests that life sciences organizations can deploy Artificial Intelligence to address concrete functional and emotional challenges within the patient journey. Rather than using computational capabilities primarily for persuasive targeting, the initiative represents a utility-oriented approach in which technological innovation supports patient experience and therapeutic engagement. Its relevance to the Data + Empathy framework lies not merely in the fact that AI operates in the background, but in the broader allocation of technological and human value: computational precision is assigned to a clearly delimited functional task, while the design of the overall experience remains oriented toward reducing patient friction, preserving agency, and responding to an identifiable human need.

For this reason, *Scrolling Therapy* constitutes a particularly relevant first case within the comparative analysis. It establishes one possible configuration of responsible AI-driven healthcare communication in which technological sophistication is not measured by the visibility or autonomy of the artificial system, but by its capacity to integrate unobtrusively into an existing patient behavior. The subsequent cases allow this configuration to be compared with alternative approaches in which data are translated into emotionally meaningful experiences or in which organizations deliberately establish boundaries around the visible use of generative Artificial Intelligence.

5.1.2 Case 2: ASCO - The Most Beautiful Sound (Translational Acoustic Data as Emotional Oncology Support)

In 2023, the American Society of Clinical Oncology (ASCO), in collaboration with Grey Health & Wellness, introduced *The Most Beautiful Sound* during the ASCO Annual Meeting (Grey Health & Wellness, 2023). The initiative was conceptualized as a multisensory health communication project designed to translate microscopic biophysical phenomena into an accessible auditory experience for oncology patients, clinicians, caregivers, and broader healthcare audiences (Grey Health & Wellness, 2023). Unlike conventional pharmaceutical communication strategies centered on product information, treatment claims, or audience targeting, the initiative explored whether complex scientific information could be transformed into a sensory experience capable of connecting biomedical knowledge with the emotional dimension of serious illness.

In clinical oncology, disease progression and therapeutic response are conventionally evaluated and communicated through multiple forms of medical evidence, including imaging technologies, laboratory biomarkers, histological examination, and statistical indicators. These forms of evidence are indispensable to clinical decision-making, but their technical complexity also creates a communication challenge: scientific information generated at the molecular or cellular level is not necessarily intuitive or immediately meaningful to patients without professional interpretation. The distinction between the scientific accuracy of information and its experiential accessibility is particularly relevant in oncology, where patients are simultaneously required to process complex medical information and navigate the emotional uncertainty associated with diagnosis and treatment.

From a communication perspective, this creates a broader question that extends beyond the transmission of medical facts. The challenge is not simply whether scientific information is available, but how complex biological processes can be represented without compromising their evidentiary basis while making them more tangible to non-specialist audiences. *The Most Beautiful Sound* addressed this problem through sensory translation, using computational processing not to simplify the scientific process into a promotional claim, but to create an alternative perceptual interface through which microscopic biological activity could be experienced.

Conventional Scientific Communication Interface:

[Complex Biological and Clinical Data] → [Technical Representation] → [Professional Interpretation] → [Patient Understanding]

Acoustic Translation Framework:

[Apoptotic Cellular Movement] → [Biophysical Measurement] → [Algorithmic Sonification] → [Auditory Representation]

To create this sensory representation, the project collaborated with researchers to examine cellular death, or apoptosis, in breast and lung cancer cells exposed to therapeutic intervention (Grey Health & Wellness, 2023). According to the project documentation, the scientific process involved advanced microscopy and the measurement of cellular and molecular phenomena associated with apoptosis. These measurements provided the underlying data from which the auditory representation was subsequently constructed.

The project documentation identifies Stimulated Raman Scattering (SRS) microscopy as part of this scientific process (Grey Health & Wellness, 2023). SRS microscopy enables researchers to examine molecular characteristics of biological samples through their vibrational properties, providing detailed information without necessarily relying on conventional fluorescent labeling. Within the context of *The Most Beautiful Sound*, the relevance of this technology lies not simply in its technical sophistication, but in its role as the scientific foundation from which an otherwise imperceptible biological process could be translated into another sensory modality.

During apoptosis, malignant cells undergo a series of structural and biochemical transformations. Rather than presenting these microscopic phenomena exclusively through laboratory imagery or numerical measurements, the initiative processed the resulting data through a sonification procedure, mapping selected characteristics of the observed biological process into frequencies accessible to human hearing (Grey Health & Wellness, 2023).

It is methodologically important to distinguish this process from the literal recording of cancer cells producing audible sounds. The resulting auditory material should instead be understood as **sonification**: a computational translation in which scientific data are systematically mapped onto acoustic parameters. This distinction is essential because it preserves the epistemic boundary between the underlying biological phenomenon and its communicative representation. The sound does not constitute an unmediated acoustic recording of cellular death; rather, it represents a technologically constructed sensory interpretation derived from measured biological information.

Translational Sonification Architecture:

[In-Vitro Cellular Process] → [Biophysical Measurement / SRS Microscopy] → [Cellular Data] → [Algorithmic Sonification] → [Auditory Representation]

This distinction between biological evidence and communicative representation is particularly significant within pharmaceutical and healthcare communication. Translating scientific information into an emotionally accessible format creates opportunities for patient engagement, but it also introduces an ethical responsibility to ensure that the representation does not become detached from the scientific process from which it originates. In *The Most Beautiful Sound*, the communicative value of the auditory experience therefore derives from the relationship between empirical laboratory data and its subsequent sensory translation, rather than from the sound functioning independently as evidence of a clinical outcome.

When evaluated through the conceptual model developed in this thesis, *The Most Beautiful Sound* illustrates a distinct configuration of the Data + Empathy framework. In this case, computational data processing is not primarily deployed to optimize commercial targeting, monitor purchasing behavior, or generate individualized persuasive messages (Sabapathy, 2025). Instead, computation functions as a translational mechanism between scientific observation and human experience.

The **Data** dimension is represented by the biophysical information generated through the observation and measurement of cellular processes. Computational processing organizes and translates this information into an alternative sensory format. The **Empathy** dimension, by contrast, emerges from the communicative intention to make an otherwise microscopic and abstract biological process more accessible and emotionally meaningful to people confronting cancer. As in the previous Eurofarma

case, empathy is therefore not generated by an artificial agent attempting to imitate human emotional behavior. It is embedded in the design objective that determines how technological capability is applied.

This distinction reinforces an important element of the Data + Empathy framework. The human-centered value of Artificial Intelligence and computational technologies does not necessarily depend on anthropomorphizing the system or creating synthetic conversational intimacy. Computational systems can contribute to empathetic communication by processing complexity in ways that enable healthcare professionals, patients, and institutions to encounter scientific information through more accessible forms. The technology remains instrumental; the human significance derives from the purpose assigned to it and from the context in which its output is interpreted.

The computational infrastructure of *The Most Beautiful Sound* consequently remains largely in the background, performing the mathematical transformation required to convert measured biological information into an auditory representation. The visible encounter, by contrast, is the sensory experience itself. This configuration differs from applications in which Artificial Intelligence occupies the front-facing communicative role by generating advice, simulating emotional dialogue, or producing personalized persuasive content.

Viewed through the Heuristic-Systematic Model (Chaiken, 1980), the initiative also provides an interesting example of how the same underlying subject matter may be encountered through different modes of information processing. Clinical oncology necessarily relies on systematic forms of evidence, professional interpretation, and careful evaluation. The auditory representation does not replace these processes and should not be interpreted as an alternative to clinical explanation. Instead, it adds a sensory layer through which an otherwise abstract biological phenomenon can become more immediately perceptible.

From this perspective, the sonification can be interpreted as providing a comparatively accessible heuristic representation alongside, rather than instead of, systematic medical information. The value of this distinction is important. A heuristic cue in healthcare communication should not be confused with evidence of individual therapeutic efficacy. Hearing the sonified representation of experimentally observed cellular processes does not establish that a particular patient's treatment is working, nor does it replace imaging, laboratory analysis, clinical assessment, or physician interpretation. Its communicative function is instead to create an intuitive sensory connection with a scientific phenomenon that would ordinarily remain inaccessible outside specialized laboratory environments.

This complementary relationship between systematic evidence and sensory accessibility makes the case particularly relevant to the communication challenges examined throughout this thesis. Healthcare organizations frequently face the problem of communicating technically complex information without either overwhelming audiences or reducing scientific nuance to emotionally persuasive simplifications. *The Most Beautiful Sound* illustrates one possible response: maintaining the underlying scientific reference while changing the modality through which that information is encountered.

The institutional context of the initiative is also relevant. *The Most Beautiful Sound* was associated with ASCO, a professional scientific organization, rather than being structured as a conventional

branded pharmaceutical product campaign (Grey Health & Wellness, 2023). This distinction does not automatically exempt a health communication initiative from every applicable regulatory or ethical requirement, nor should it be interpreted as evidence that professional-society communication is inherently neutral. It does, however, position the project differently from direct promotional communication centered on the commercial benefits of a specific medicinal product.

The absence of a conventional product-centered promotional structure allows the case to be analyzed primarily as an example of scientific and patient-oriented communication. Its strategic objective is therefore better understood in terms of translating scientific complexity into an accessible experience than in terms of stimulating demand for an identifiable treatment. This distinction is particularly valuable within the comparative structure of this chapter because it demonstrates that computational innovation in healthcare communication can serve functions beyond personalization, targeting, and commercial conversion.

At the same time, the emotional dimension of the initiative requires analytical caution. *The Most Beautiful Sound* was designed to provide a meaningful and potentially reassuring sensory representation of a difficult biological process, but the existence of such an intention should not be confused with evidence of a measurable psychological outcome. Without controlled evaluation of patient responses, it would be methodologically inappropriate to conclude that the project reduced anxiety, increased hope, or improved trust among oncology patients as a generalizable effect. The case instead demonstrates how emotional considerations can inform the design of scientific communication while preserving a clear distinction between communicative intention and empirically measured psychological impact.

This distinction also prevents empathy from becoming synonymous with emotional persuasion. Within the Data + Empathy framework, empathetic communication should not rely on manipulating patient vulnerability or replacing complex evidence with reassuring symbolism. Its purpose is instead to recognize the emotional and cognitive circumstances in which scientific information is received and to design communication accordingly. In this respect, *The Most Beautiful Sound* provides a particularly useful example because the emotional component is derived from the presentation of scientific information rather than from the fabrication of a synthetic narrative detached from biological evidence.

Strategically, the initiative therefore suggests that healthcare organizations can use computational processing as an intermediary between scientific complexity and human understanding. Such an approach may be particularly valuable in therapeutic areas characterized by high informational complexity and emotional vulnerability, provided that the distinction between scientific evidence and its communicative representation remains transparent. The case also illustrates that innovation in healthcare communication does not necessarily require greater personalization or increasingly autonomous Artificial Intelligence. Technological sophistication can instead be directed toward making existing scientific knowledge more interpretable, tangible, and meaningful.

Compared with *Scrolling Therapy*, *The Most Beautiful Sound* consequently represents a different operationalization of the Data + Empathy framework. In the Eurofarma case, computational technology is integrated into patient behavior to support a functional rehabilitative activity. In the ASCO case, computational processing mediates between scientific data and sensory experience. The

underlying principle remains comparable, but the mechanism changes: the first case emphasizes functional utility, whereas the second emphasizes translational accessibility and emotional meaning.

This distinction contributes to the broader comparative argument of the chapter. Data + Empathy should not be understood as prescribing a single technological architecture in which every AI system must perform the same backend function. Rather, the framework identifies a broader strategic principle: computational capability acquires healthcare communication value when its role is clearly connected to an identifiable human need and when technological processing remains subordinate to scientific integrity, ethical responsibility, and meaningful stakeholder experience.

Within these boundaries, *The Most Beautiful Sound* provides a relevant real-world illustration of how data-intensive scientific processes can be transformed into patient-oriented communication without making computational technology itself the central protagonist of the interaction. Its significance lies not in demonstrating that sonification produces measurable therapeutic outcomes, but in showing how scientific organizations can experiment with alternative forms of communication that connect objective biological information with the emotional realities of serious illness.

5.1.3 Case 3: Dove - The Code (Critical Benchmark: Ethical Resistance to Algorithmic Aesthetic Homogenization)

In 2024, marking the twentieth anniversary of its *Campaign for Real Beauty*, Dove, a personal care brand owned by Unilever, launched *The Code* in collaboration with Ogilvy (Dove, 2024; Unilever, 2024). The initiative emerged within a broader public debate concerning the rapid diffusion of generative Artificial Intelligence across advertising, creative production, and digitally mediated representations of human appearance. Unlike the previous cases examined in this chapter, Dove operates primarily within the fast-moving consumer goods (FMCG) and personal care sector. It is not a pharmaceutical manufacturer, and *The Code* does not constitute a medicinal product campaign.

Its inclusion within this comparative analysis therefore requires a different methodological justification. *The Code* is purposively incorporated as an **adjacent-sector critical benchmark** rather than as direct evidence of pharmaceutical AI practice. Personal care, beauty, wellness, dermatology, and health-related communication increasingly coexist within the same digital environments, where users encounter commercial, informational, influencer-generated, and health-oriented content through interconnected social media feeds. Consequently, representational practices developed outside the pharmaceutical industry can provide strategically relevant insights into questions of authenticity, algorithmic bias, human representation, and responsible technological adoption that also affect life sciences communication.

This relevance becomes particularly evident when considering younger digital audiences. Digital-native users encounter representations of health, appearance, skin, body image, wellness, and self-optimization within media environments in which the distinction between organically produced human imagery and synthetically generated content is becoming progressively less visible. Research concerning younger populations has highlighted the importance of examining the potential psychological consequences associated with digitally mediated content, misinformation, social comparison, and emerging AI-generated environments (Chantre et al., 2025; Mansfield et al., 2025). Within this context, *The Code* provides an opportunity to examine how a commercial organization

can respond not by maximizing the visible use of generative AI, but by publicly defining boundaries around its deployment.

Generative Image Pipeline and Representation Risk:

[Text-to-Image Prompt] → [Generative Model / Training-Data Patterns] → [Synthetic Human Representation] → [Potential Reproduction of Aesthetic Bias]

The Code Boundary-Setting Response:

[Recognition of Representational Risk] → [Corporate AI Commitment] → [Real Beauty Prompt Guidelines] → [Human-Centered Representation]

The campaign specifically addressed the capacity of text-to-image generative models to reproduce and potentially amplify established aesthetic conventions embedded within the visual material on which such systems are trained (Dove, 2024). Generative image models do not construct representations of beauty independently from existing cultural material. Their outputs emerge from statistical patterns learned from large datasets, meaning that recurrent visual conventions, demographic imbalances, and historically dominant representations may influence the types of bodies, faces, skin characteristics, and aesthetic attributes generated in response to broad prompts.

Within its campaign materials, Dove illustrated this problem by experimenting with generic beauty-related prompts and examining the visual outputs produced by generative AI systems (Dove, 2024). The resulting representations were used to highlight the tendency of generative systems to reproduce narrow or idealized conceptions of beauty when prompts do not explicitly introduce greater representational diversity. Rather than treating these outputs as neutral technological artifacts, *The Code* framed them as evidence of a broader communication challenge: the growing ability of computational systems to participate in the construction and circulation of culturally meaningful representations of human appearance.

This distinction is particularly relevant because generative AI changes the scale at which idealized imagery can be produced. Conventional image manipulation already enabled advertisers and content creators to modify physical appearance extensively. Generative systems, however, introduce the possibility of producing entirely synthetic individuals and environments at comparatively low marginal cost and with substantial creative flexibility. The ethical question therefore extends beyond whether a photograph has been digitally retouched. It concerns whether commercial communication should replace human representation itself with probabilistically generated visual substitutes, particularly in contexts where authenticity, vulnerability, identity, and physical experience form part of the communicative relationship.

From a health and well-being perspective, these questions become particularly significant when synthetic imagery intersects with body image, skin appearance, aging, and digitally mediated standards of attractiveness. Existing research on digital media and younger audiences provides reason to examine carefully how highly curated or idealized representations may interact with social comparison and psychological well-being (Mansfield et al., 2025). However, it would be methodologically inappropriate to infer from *The Code* itself that exposure to generative imagery directly causes body dysmorphia or other specific psychological outcomes. The campaign is therefore

treated here as a communication and governance case rather than as clinical evidence of the psychological effects of generative AI.

To respond to these representational concerns, Dove adopted two complementary strategic interventions. First, the brand publicly established a corporate commitment concerning the use of generative Artificial Intelligence in representations of women within its advertising (Dove, 2024; Unilever, 2024). This commitment functions as a form of voluntary organizational boundary-setting: instead of waiting for comprehensive external regulation to determine every permissible application of generative imagery, the company articulated an internal standard intended to preserve the role of authentic human representation within its communication strategy.

Second, Dove introduced its *Real Beauty Prompt Guidelines*, designed to encourage users and creators to formulate prompts that produce more diverse and representative visual outputs (Dove, 2024). This component is analytically important because the campaign does not adopt a purely anti-technological position. Dove's strategy distinguishes between different uses of generative AI: it establishes limits around synthetic representation in its own brand communication while simultaneously engaging with generative systems as objects that can be interrogated, influenced, and critically understood.

The combination of these two interventions produces a more nuanced governance configuration than a simple rejection of Artificial Intelligence. On one side, the organization identifies a domain in which it considers authentic human representation strategically and ethically significant. On the other, it acknowledges that generative technologies are already embedded within the wider digital environment and therefore provides users with an educational mechanism intended to expose and potentially counter some of the representational tendencies produced by these systems.

Boundary-Setting Configuration:

[Generative AI Capability] → [Assessment of Representational Risk] → [Organizational Boundary] → [Human Representation in Brand Communication]

[Generative AI Capability] → [Critical Engagement] → [Prompt Guidelines] → [More Deliberate Use of Generative Systems]

When analyzed through the conceptual model developed in this thesis, *The Code* illustrates a **boundary-setting configuration** of the Data + Empathy framework. This differs substantially from both *Scrolling Therapy* and *The Most Beautiful Sound*. Eurofarma uses computational technology as a functional mechanism embedded within rehabilitation, while the ASCO initiative employs data processing as a bridge between microscopic scientific information and sensory experience. Dove, by contrast, is relevant precisely because the organization identifies circumstances in which expanding the visible role of generative technology may conflict with its human-centered communication objectives.

The Data + Empathy framework should therefore not be interpreted as requiring organizations to deploy Artificial Intelligence wherever technological implementation is possible. Responsible AI strategy also involves decisions concerning where automation should be limited, where human representation should remain visible, and where computational capabilities may create tensions with stakeholder expectations. In this case, the “Empathy” dimension is expressed partly through the

preservation of authentic human representation, while the “Data” dimension involves recognizing and critically examining the representational patterns generated by computational systems.

This interpretation expands the framework beyond the idea that responsible AI consists exclusively of placing technology in an invisible backend. The strategic question is instead one of **role allocation**. Computational systems may perform valuable analytical, organizational, or supportive functions while organizations simultaneously establish areas in which synthetic generation is considered inappropriate or strategically undesirable. From this perspective, technological restraint can itself constitute an element of AI governance.

Viewed through the Heuristic-Systematic Model (Chaiken, 1980), *The Code* can also be interpreted as an intervention that encourages audiences to move from immediate visual reception toward more deliberate consideration of how synthetic imagery is produced. Highly polished visual content can function as an immediately accessible communication cue without requiring audiences to examine the technical processes responsible for its construction. By exposing the relationship between prompts, generative systems, and resulting representations, the campaign makes the production mechanism itself part of the message.

The analytical relevance of the Heuristic-Systematic Model here does not depend on assuming that Instagram, TikTok, or other visual platforms are inherently governed by heuristic processing. Such a claim would oversimplify both user behavior and Chaiken's theoretical model. Rather, the campaign illustrates how a representation that may initially be encountered primarily at the level of visual appearance can subsequently become the object of more deliberate scrutiny once audiences are invited to consider its synthetic origin, underlying training patterns, and representational implications.

In this sense, *The Code* attempts to transform algorithmic production from an invisible technical process into an explicit object of communication. This is particularly relevant to AI literacy. As generative outputs become increasingly realistic, the ability to evaluate not only what an image represents but also **how and why it was produced** becomes an important component of critical digital engagement. The campaign therefore operates simultaneously at the level of brand positioning and public technological awareness.

The relationship with the **Uncanny Valley of Mind** requires greater theoretical caution. Stein and Ohler (2017) examined psychological responses to artificial entities perceived as possessing human-like mental characteristics. Their framework was not developed specifically to explain reactions to AI-generated beauty advertising or synthetic photographs. Consequently, *The Code* should not be presented as empirical evidence that generative imagery triggers the Uncanny Valley of Mind.

Nevertheless, the theory provides a limited interpretative analogy for one aspect of the broader authenticity problem examined in this thesis. Both phenomena concern tensions that can emerge when artificial systems approximate dimensions traditionally associated with human beings while their artificial origin remains relevant to perception. In the context of generative advertising, however, the more appropriate analytical issue is not necessarily “uncanniness” in the strict theoretical sense, but the potential discrepancy between apparent human authenticity and computational fabrication. For this reason, Stein and Ohler's framework can inform the discussion conceptually, but should not be treated as a direct causal explanation of audience responses to *The Code*.

A related distinction applies to psychological reactance. It would be inappropriate to infer that exposure to synthetic imagery necessarily generates reactance or brand rejection. Audience responses depend on multiple contextual variables, including disclosure, perceived authenticity, prior attitudes toward AI, the purpose of the communication, and individual technological familiarity. What *The Code* demonstrates more directly is that Dove itself identified synthetic human representation as sufficiently relevant to its brand positioning to establish a public organizational boundary around its use.

This distinction between **observed organizational action** and **inferred audience response** is methodologically important. The former is directly observable in the campaign: Dove established a policy and communicated a position concerning generative AI. The latter would require independent empirical research measuring consumer perceptions, trust, authenticity judgments, or behavioral responses. The present case analysis therefore uses the campaign to examine corporate governance and communication strategy rather than to claim psychological outcomes that the available evidence does not establish.

From a governance perspective, *The Code* is particularly relevant to the regulatory-lag problem developed in Chapter 4. Generative AI technologies have diffused through advertising and digital communication more rapidly than comprehensive sector-specific regulatory frameworks have been able to define every permissible representational practice. Within such an environment, organizations face decisions that cannot always be deferred until external authorities establish detailed rules for every technological application.

Dove's response illustrates one possible form of **proactive self-governance**. The company established an internal boundary based on its existing brand philosophy and representational objectives rather than treating the absence of comprehensive prohibition as equivalent to unrestricted strategic permission. This does not mean that voluntary corporate standards can replace statutory regulation. As argued in Chapter 4, private self-regulation lacks the democratic legitimacy, enforcement authority, and universal applicability of formal regulatory frameworks. It does, however, illustrate how organizational governance can operate within the temporal space created by regulatory lag.

This distinction is directly relevant to pharmaceutical communication. Life sciences organizations increasingly encounter generative technologies capable of producing synthetic patient imagery, automated narratives, personalized educational materials, conversational interfaces, and emotionally calibrated communication at scale. The technical availability of these capabilities does not by itself determine whether their use is ethically appropriate in every context. Therapeutic areas involving visible symptoms, chronic illness, oncology, mental health, dermatological conditions, disability, or other highly personal experiences may require particular sensitivity regarding the substitution of authentic patient representation with synthetic equivalents.

The strategic lesson from *The Code* is therefore not that pharmaceutical organizations should reproduce Dove's policy or categorically reject generative AI. Such a conclusion would exceed what an adjacent-sector case can support. Rather, the case suggests that life sciences organizations may benefit from defining explicit organizational boundaries concerning **where, when, and for what purposes** generative technology is considered appropriate.

This approach is consistent with the broader governance architecture developed in Chapter 4. Compliance by Design addresses whether content satisfies applicable regulatory and scientific requirements; Human-in-the-Loop governance preserves professional oversight and contextual judgment; and the Data + Empathy framework introduces an additional strategic question concerning whether a technically permissible application is also compatible with human-centered communication objectives. *The Code* illustrates that responsible technological governance can therefore include deliberate non-use as well as deployment.

From the perspective of stakeholder trust, this distinction is particularly significant. Authenticity cannot be assumed to result automatically from the presence of a human image, just as synthetic content cannot be assumed automatically to generate distrust. Nevertheless, organizations operating in sensitive health environments must consider whether the mode of representation is congruent with the vulnerability and lived experience of the people depicted. Where communication depends heavily on testimony, emotional identification, or representations of illness, undisclosed synthetic substitution may create ethical questions that extend beyond conventional measures of advertising efficiency.

The campaign therefore broadens the meaning of the Data + Empathy framework. In *Scrolling Therapy*, computational capability serves **functional rehabilitation utility**. In *The Most Beautiful Sound*, computational processing serves **scientific-to-sensory translation**. In *The Code*, the relevant strategic contribution lies in **ethical boundary-setting around synthetic representation**. These three configurations demonstrate that the framework is not dependent on a single technology or operational architecture. Instead, it concerns the governance of the relationship between computational capability and identifiable human needs.

Importantly, this comparative progression also prevents the framework from becoming an unconditional endorsement of AI adoption. If Data + Empathy were interpreted merely as a method for making increasingly sophisticated Artificial Intelligence appear more human, it would risk reproducing precisely the trust problems identified in the preceding chapters. The framework instead requires organizations to determine whether technological deployment genuinely contributes to stakeholder value and whether particular human functions or representations should remain outside the domain of automation.

For pharmaceutical and life sciences organizations, *The Code* consequently functions as a useful external benchmark rather than a directly transferable pharmaceutical model. It illustrates how a major commercial organization can publicly articulate limits around a rapidly diffusing technology and integrate those limits into its broader communication identity. The relevance of the case lies less in the specific prohibition adopted by Dove than in the governance principle underlying it: organizations do not need to treat every technically feasible AI application as strategically desirable.

In strategic terms, this case therefore suggests that responsible AI adoption may require both **technological capability and technological restraint**. Computational intelligence can generate significant value when used for evidence synthesis, operational support, personalization, and other appropriate functions, but human-centered governance also requires explicit consideration of contexts in which automation could compromise authenticity, representation, or stakeholder expectations. Within the Data + Empathy framework, the central strategic objective is consequently not maximum AI deployment, but appropriate AI deployment.

As the third case in the comparative analysis, *The Code* completes a progression from **AI as functional utility**, through **computation as sensory translation**, to **AI as an object of ethical boundary-setting**. This progression provides the foundation for the cross-case synthesis developed in Section 5.1.4, where the three initiatives can be compared systematically according to technological visibility, stakeholder value, communicative mechanism, governance safeguards, and their respective interpretations through the Data + Empathy framework.

Table 5.1 Comparative Framework of Analyzed Global AI-Driven Campaigns

Analytical Dimension	Eurofarma: Scrolling Therapy	ASCO: The Most Beautiful Sound	Dove: The Code
Primary Campaign Objective	Support facial rehabilitation and adherence by integrating facial exercises into everyday digital behavior for people living with Parkinson’s disease and hypomimia.	Translate microscopic biological processes associated with cancer-cell apoptosis into an auditory representation intended to make complex scientific information more tangible and perceptually accessible.	Address narrow algorithmic representations of beauty and establish an explicit organizational boundary concerning the synthetic representation of women.
Sector & Health Category	Pharmaceutical / Neurological Care and Technology-Enabled Rehabilitation	Healthcare / Clinical Oncology, Scientific Communication and Supportive Care	Adjacent-Sector Benchmark / Personal Care, Beauty and Wellness Communication
Technological / Computational Configuration	Functional Computer-Vision Utility: on-device machine learning processes facial movements and translates predefined expressions into social-media navigation commands.	Computational Sonification / Translational Accessibility: experimentally observed cellular and biophysical information is transformed into an auditory representation.	Generative-AI Governance / Ethical Boundary-Setting: generative AI is examined as a representational technology, while the organization establishes explicit limits concerning AI-generated representations of real women.
Primary Target Stakeholder	People living with Parkinson’s disease, with secondary relevance to caregivers and neurological-care stakeholders.	Oncology patients, caregivers, healthcare professionals, and the broader oncology community.	Women and girls, beauty consumers, creative and advertising professionals, and broader audiences exposed to AI-mediated representations of beauty.
HSM Interpretative Lens	Primarily Functional Rather Than Persuasive: the initiative reduces friction by embedding a rehabilitative activity within a familiar digital behavior; it should not be interpreted as direct evidence of a specific heuristic or systematic processing route.	Sensory Accessibility: sonification may provide an accessible perceptual cue alongside, rather than in place of, systematic engagement with the underlying scientific evidence.	Systematic Ethical Reflection: the initiative encourages deliberate consideration of algorithmic beauty standards, synthetic representation, and the implications of generative AI for authenticity.
Regulatory & Ethical Safeguards	On-device processing reduces the need to transmit facial-movement data to external infrastructure, representing a privacy-oriented design choice conceptually consistent with data-minimization principles without establishing formal GDPR certification.	Epistemic boundaries distinguish the underlying biological evidence from its sonified representation, avoiding the implication that the auditory output independently demonstrates therapeutic effectiveness or patient outcomes.	Voluntary organizational boundary-setting concerning AI-generated representations of real women, complemented by guidance intended to encourage more representative generative prompting practices.

Analytical Dimension	Eurofarma: Scrolling Therapy	ASCO: The Most Beautiful Sound	Dove: The Code
Data + Empathy Alignment	Empathy as Friction Reduction: computational capability addresses a concrete difficulty within the rehabilitation experience without requiring the system to simulate human emotion.	Empathy as Translational Accessibility: computational processing creates an alternative sensory interface through which complex scientific information can become more perceptually accessible.	Empathy as Representational Responsibility: organizational governance considers whether technological representation remains consistent with authenticity, diversity, and the human subjects represented in communication.

5.1.4 Cross-Case Synthesis: Key Strategic and Operational Findings

A comparative examination of the three initiatives detailed in Table 5.1 highlights several strategic dynamics regarding how computational technologies and human-centered communication interact across healthcare and adjacent domains. Rather than evaluating these initiatives as empirical validations of a single operational model, the cross-case synthesis identifies four comparative themes that illustrate the practical relevance of the Data + Empathy framework. The comparison is therefore interpretative rather than causal: its purpose is not to establish that one technological configuration systematically produces superior stakeholder outcomes, but to examine how different organizations allocate computational capabilities, human involvement, and ethical boundaries in response to distinct communication challenges.

Although the three initiatives differ substantially in sector, technological architecture, target stakeholder, and communicative objective, they share a broader strategic characteristic. In each case, technological capability is subordinated to a defined human-centered purpose rather than treated as an end in itself. Eurofarma employs computer vision to restructure a rehabilitation experience; ASCO uses computational translation to make microscopic scientific information perceptible through another sensory modality; and Dove establishes boundaries around a form of generative technology when its application may conflict with the organization's representational objectives. Their comparative relevance therefore lies not in technological similarity, but in the different ways through which organizations determine the appropriate role of computation within stakeholder communication.

1. The Strategic Value of AI as an Enabling Backend Utility

The first comparative theme concerns the visibility, authority, and functional role assigned to computational technology. Across the cases examined, the most relevant strategic distinction is not simply whether Artificial Intelligence is present or absent, but where it is positioned within the stakeholder experience and what degree of communicative or decisional authority it receives.

In *Scrolling Therapy*, computer-vision algorithms operate on the user's device to translate facial movements into digital navigation commands (Dentsu Creative, 2023). The computational layer is essential to the functioning of the application, yet it does not occupy the role of an autonomous clinician, diagnostic authority, or synthetic communication agent. Its purpose is narrowly defined: recognizing facial movement and converting it into an interface action. The technology therefore enables the experience without becoming the primary object of the patient's interaction.

This configuration is consistent with the broader principle developed through the Data + Empathy framework: technological sophistication does not necessarily require technological visibility. The value generated by the system derives from what the computational capability enables for the patient rather than from the patient's continuous awareness of the underlying machine-learning architecture. In this respect, *Scrolling Therapy* illustrates how Artificial Intelligence can operate as an infrastructural mechanism supporting a human-centered objective while maintaining limited decisional authority.

A comparable, although technically distinct, configuration emerges in *The Most Beautiful Sound*. Here, computational processing is used to translate biological information associated with experimentally observed cellular processes into an auditory representation (Grey Health & Wellness, 2023). Once again, the computational mechanism performs an enabling role. The mathematical transformation required for sonification is indispensable to the project, but the algorithm itself is not presented as the communicative protagonist. The stakeholder encounters the resulting sensory representation rather than an artificial agent claiming clinical authority.

The comparison between these first two cases is particularly informative because it demonstrates that “backend utility” does not refer to a single technological architecture. Eurofarma uses computer vision to connect physical behavior with a digital interface, whereas the ASCO initiative uses computational processing to connect scientific measurement with sensory perception. Their technical functions are different, but their strategic positioning is comparable: computation mediates between two domains while remaining subordinate to the purpose of the interaction.

This observation also clarifies the meaning of “backend” within the Data + Empathy framework. The term should not be understood exclusively in its narrow software-engineering sense. Rather, it describes a strategic configuration in which computational intelligence performs analytical, translational, organizational, or functional work without automatically assuming the visible human-facing role traditionally occupied by clinicians, patients, caregivers, or other trusted communicators. The central question is therefore not whether the algorithm is technically located on a server, device, or cloud infrastructure, but whether computational capability becomes the dominant communicative actor.

The Code introduces a different but complementary perspective. Unlike Eurofarma and ASCO, Dove's initiative is not primarily an example of computational technology being deployed as an enabling healthcare utility. Instead, it provides an adjacent-sector benchmark concerning the circumstances in which an organization may deliberately restrict the visible application of generative technology. Dove's public position regarding AI-generated representations of women reflects a strategic judgment that technological feasibility does not necessarily imply communicative desirability (Dove, 2024; Unilever, 2024).

The relevance of this case lies in the distinction between AI capability and AI role allocation. Generative models are technically capable of producing increasingly realistic synthetic human imagery, but *The Code* illustrates an organizational decision to establish limits around that capability within a specific representational context. This does not demonstrate that front-facing generative AI inherently produces stakeholder skepticism, nor does it establish that authentic human imagery

automatically generates trust. Instead, it demonstrates that the placement and visibility of AI can itself become an object of strategic governance.

Taken together, the three cases therefore suggest a broader principle: responsible AI adoption requires organizations to determine not only what computational systems can perform, but also which functions should be delegated to them and which should remain visibly human. This distinction is particularly relevant to life sciences communication, where technological efficiency frequently intersects with clinical authority, emotional vulnerability, privacy expectations, and the credibility of health-related information.

The implications extend directly to the personalization architectures examined in Chapter 3. Retrieval-Augmented Generation, predictive Next Best Action systems, automated prescreening, Dynamic Creative Optimization, and other computational infrastructures may create significant operational value without necessarily becoming autonomous front-facing communicators. Their analytical capabilities can support evidence retrieval, content selection, compliance processes, audience understanding, and decision support while human professionals retain responsibility for contextual interpretation and consequential communication (Sabapathy, 2025; Verma et al., 2024).

This does not imply that every visible application of Artificial Intelligence is inappropriate. Conversational systems, generative interfaces, and other human-facing technologies may provide legitimate value under appropriate conditions. Rather, the comparative evidence presented here suggests the importance of questioning the assumption that increasing technological capability should automatically correspond to increasing technological visibility. The appropriate degree of visibility depends on the communicative objective, stakeholder vulnerability, regulatory environment, and level of authority assigned to the system.

The distinction becomes especially important in pharmaceutical environments because communication is rarely evaluated exclusively according to engagement or conversion metrics. Scientific substantiation, regulatory compliance, patient safety, informed interpretation, and professional accountability remain central constraints. A system capable of generating persuasive or personalized content at scale may therefore require a different governance architecture from a system performing data processing or operational support in the background.

From the perspective of the Data + Empathy framework, the comparative value of these cases consequently lies in their illustration of functional asymmetry between machine and human roles. Computational systems are particularly valuable when assigned tasks suited to their analytical and processing capabilities, whereas empathy, contextual interpretation, accountability, and the representation of lived experience remain areas in which human involvement retains strategic importance. This does not establish an absolute technological boundary; instead, it provides a governance principle for determining how responsibilities should be allocated.

The first cross-case finding can therefore be summarized not as the superiority of invisible AI, but as the strategic value of purpose-limited computational enablement.

Scrolling Therapy illustrates AI supporting behavioral rehabilitation; *The Most Beautiful Sound* illustrates computation translating scientific complexity; and *The Code* illustrates the deliberate restriction of synthetic representation where an organization considers authentic human presence

central to its communication philosophy. Together, these cases support a model in which Artificial Intelligence is evaluated according to its contribution to a clearly defined stakeholder objective rather than according to the extent of its visibility or autonomy.

This interpretation is consistent with the broader governance perspective developed in Chapter 4. Responsible AI deployment does not require organizations to choose between technological innovation and human-centered communication. It requires them to define the relationship between the two. Within this relationship, computational capability can expand analytical and operational possibilities, while human judgment determines the objectives, boundaries, and contexts through which those capabilities are deployed (OECD, 2023; Verma et al., 2024).

Accordingly, the comparative analysis suggests that one of the most relevant strategic questions for pharmaceutical organizations is not simply “*Where can AI be implemented?*”, but rather “*Where does AI create meaningful stakeholder value, and what degree of visibility and authority should it receive in that context?*” This shift from technological adoption to technological role allocation constitutes the first major operational implication emerging from the three cases and provides the foundation for the subsequent comparative themes.

2. Functional Utility versus Synthetic Persuasion

A second comparative finding concerns the distinction between functional technological utility and the use of Artificial Intelligence as a mechanism for synthetic persuasion. The three initiatives suggest that the strategic relevance of computational technology cannot be evaluated exclusively according to the sophistication of the underlying model or the degree of personalization it enables. Equally important is the type of value that the technology is designed to create for the stakeholder. Within the cases examined, this distinction emerges particularly clearly when the functional configurations represented by *Scrolling Therapy* and *The Most Beautiful Sound* are compared with the representational concerns addressed by Dove’s *The Code*.

In *Scrolling Therapy*, technological value is directly connected to an identifiable patient-related function. Computer vision is not introduced primarily to increase exposure to a branded message, optimize promotional conversion, or generate individualized persuasive content. Instead, it translates facial movements into digital navigation commands and integrates rehabilitation exercises into an activity already embedded within the user’s everyday digital routine (Dentsu Creative, 2023). The computational capability therefore has an instrumental purpose that can be understood independently from the commercial identity of the organization responsible for the initiative.

This distinction is strategically relevant because it shifts the starting point of AI implementation. Rather than beginning with an available technology and subsequently identifying opportunities for its commercial deployment, the initiative begins with a specific source of patient friction: the difficulty of maintaining repetitive facial rehabilitation exercises and the emotional burden associated with conventional mirror-based practice. Technology is subsequently introduced as one possible mechanism for restructuring that experience. In this configuration, stakeholder need precedes technological application.

The Most Beautiful Sound reflects a comparable principle through a substantially different mechanism. Its computational architecture does not perform a rehabilitative function, but instead

addresses a communication problem generated by the complexity and perceptual inaccessibility of microscopic biological information. The sonification process translates measured cellular phenomena into an auditory representation, allowing an otherwise imperceptible scientific process to be encountered through another sensory modality (Grey Health & Wellness, 2023). Computational processing therefore creates value through translation rather than through persuasive message generation.

The comparison between these two initiatives reveals an important characteristic of utility-oriented AI. Utility does not necessarily require that the technology perform a clinical intervention or directly modify a health outcome. It may also consist of reducing behavioral friction, improving the accessibility of complex information, translating scientific data into more intelligible forms, or enabling stakeholders to interact with healthcare-related processes in ways that would otherwise be difficult. What connects the Eurofarma and ASCO cases is therefore not their technical architecture, but the fact that computational capability is assigned to a clearly identifiable problem within the stakeholder experience.

This approach contrasts with a different potential use of Artificial Intelligence in digital communication: the generation of synthetic persuasive material at scale. Generative and AI-driven systems can support increasingly scalable forms of content production, personalization, and audience adaptation, expanding the operational capabilities available to pharmaceutical commercial functions (Jones et al., 2024; Sabapathy, 2025; Verma et al., 2024). As discussed in the preceding chapters, such capabilities may contribute to the operational efficiency of contemporary pharmaceutical communication architectures (Sabapathy, 2025; Verma et al., 2024).

However, technical scalability and stakeholder utility are not equivalent concepts. A system may significantly increase the volume, speed, or personalization of communication without necessarily increasing the substantive value delivered to the recipient. This distinction becomes especially important in healthcare environments, where the stakeholder is not merely a consumer exposed to a commercial message but may simultaneously be a patient managing illness, a caregiver navigating uncertainty, or a healthcare professional evaluating scientifically consequential information.

The strategic issue is therefore not whether generative technology should be categorically excluded from pharmaceutical communication. Such a conclusion would be inconsistent with the broader argument of this thesis and would overlook legitimate applications of generative systems in information retrieval, content adaptation, educational support, workflow assistance, and other appropriately governed contexts. Rather, the comparative cases indicate the importance of distinguishing between AI that primarily increases an organization's communicative capacity and AI that directly addresses an identifiable stakeholder need.

These two objectives are not mutually exclusive. A well-designed system may achieve both. Nevertheless, the distinction provides a useful governance criterion because it changes the question asked during technological implementation. A capacity-oriented approach asks what additional content, targeting, automation, or personalization the organization can produce through Artificial Intelligence. A utility-oriented approach asks what specific stakeholder problem the technology is expected to address and whether computational intervention represents an appropriate mechanism for doing so.

This distinction can be represented conceptually as follows:

Capacity-Oriented AI Logic:

[Technological Capability] → [Increased Content / Personalization Capacity] → [Stakeholder Exposure] → [Expected Communication Outcome]

Utility-Oriented AI Logic:

[Identified Stakeholder Need] → [Assessment of Appropriate Technological Role] → [AI-Enabled Function] → [Stakeholder Value]

The difference between these pathways is not merely semantic. In the first configuration, technological capability can become the starting point of the strategy, creating a risk that implementation is justified primarily because automation is available. In the second, the stakeholder problem provides the starting point, and Artificial Intelligence is evaluated according to whether it contributes meaningfully to its resolution. Within the Data + Empathy framework, the latter sequence is particularly significant because it prevents technological sophistication from becoming detached from human purpose.

Dove's *The Code* provides an adjacent-sector counterpoint to this distinction. The campaign focuses specifically on generative visual systems and the capacity of such technologies to construct synthetic representations of human beauty (Dove, 2024; Unilever, 2024). The concern articulated by Dove is not that generative Artificial Intelligence lacks creative capability. On the contrary, the relevance of the campaign depends precisely on the increasing sophistication of that capability. The governance question arises because technically convincing synthetic representation may not always be consistent with the organization's objectives concerning authenticity and representational diversity.

In this sense, *The Code* introduces a useful distinction between creative efficiency and representational legitimacy. A synthetic human image may be generated efficiently, adapted rapidly, and optimized for specific aesthetic parameters, while simultaneously raising questions concerning who is being represented, which visual norms are being reproduced, and whether a synthetic substitute is appropriate within the communicative context. The campaign therefore illustrates that greater generative capability does not resolve the normative question of whether that capability should be used.

This observation acquires additional significance when transferred cautiously to pharmaceutical and healthcare communication. Synthetic patient imagery, AI-generated testimonials, artificially constructed narratives of illness, or automated emotional messaging could theoretically enable organizations to produce highly scalable and personalized communication. Yet the strategic value of these applications cannot be inferred solely from their technical performance. In contexts involving disease, vulnerability, treatment expectations, or lived patient experience, the authenticity and provenance of representation may become part of the ethical meaning of the communication itself.

This does not imply that synthetic representation is inherently deceptive. Generative content can be transparently disclosed, ethically designed, and appropriately used in educational or illustrative contexts. Nor does the comparison establish that audiences invariably prefer human-generated

communication. The cases analyzed in this chapter do not provide the empirical evidence necessary to support either categorical conclusion. Instead, they suggest that the appropriateness of synthetic communication depends on its function, disclosure, context, and relationship to stakeholder expectations.

The distinction is particularly relevant to the personalization paradox examined in Chapter 3. Artificial Intelligence enables organizations to make communication increasingly adaptive, predictive, and individually relevant (Sabapathy, 2025; Verma et al., 2024). At the same time, greater personalization may require more extensive data processing and can intensify concerns regarding privacy, perceived surveillance, and the appropriate boundaries of automated communication (OECD, 2023; Verma et al., 2024). Consequently, maximizing personalization does not automatically maximize perceived value.

Functional utility provides a different basis for personalization. Rather than personalizing primarily to increase persuasive effectiveness, an organization can use data and computational capabilities to adapt a service, educational resource, or supportive function to a stakeholder's circumstances. In such configurations, personalization becomes valuable because it modifies the usefulness of the interaction rather than merely increasing the probability that a promotional message will attract attention.

This distinction does not eliminate commercial objectives. Pharmaceutical organizations remain commercial entities, and communication strategies inevitably operate within broader organizational goals. The relevant strategic question is therefore not whether commercial value should disappear, but how commercial objectives are connected to stakeholder value creation. Utility-oriented applications offer one possible model in which organizational value may emerge indirectly from the provision of useful, credible, or supportive experiences rather than exclusively from direct persuasive pressure.

Within the Data + Empathy framework, this relationship can be conceptualized as a shift from persuasion-first personalization to value-first personalization. The distinction does not reject persuasion as a legitimate component of marketing communication. Instead, it establishes a sequencing principle: computational capabilities should first be evaluated according to the human problem they address, the evidence supporting their function, and the ethical implications of their deployment. Promotional or organizational value should subsequently emerge within those boundaries rather than determine them retrospectively.

The Eurofarma case illustrates this sequence particularly clearly. The technological architecture is meaningful because it responds to a pre-existing rehabilitative challenge. The ASCO case follows a comparable logic at the level of communication: sonification is meaningful because it creates an alternative interface between complex biological information and human perception. Dove, by contrast, demonstrates the inverse governance decision: when a technological capability is perceived as potentially conflicting with an established human-centered representational objective, the organization may decide to restrict rather than maximize its deployment.

The comparison therefore produces three distinct but complementary strategic responses to technological capability:

[Stakeholder Friction] → [Deploy AI as Functional Utility] –
Eurofarma

[Scientific Complexity] → [Deploy Computation as Translational
Utility] – ASCO

[Representational Risk] → [Establish Limits on Generative
Deployment] – Dove

Together, these configurations reinforce the proposition that the strategic value of Artificial Intelligence cannot be reduced to automation intensity. In some contexts, greater automation may improve accessibility or functionality; in others, the most responsible decision may involve limiting automation and preserving human representation. The common variable is not the amount of AI deployed but the relationship between technological capability and the stakeholder objective it is intended to serve.

This finding also has implications for the evaluation of AI investments within pharmaceutical organizations. Conventional performance metrics may prioritize efficiency gains, content-production speed, audience reach, click-through rates, conversion indicators, or cost reductions. These measures remain operationally relevant, but they may provide an incomplete assessment of AI-driven communication in health-related contexts. A utility-oriented model would require complementary indicators capable of evaluating whether the technological intervention improves information accessibility, reduces stakeholder friction, supports comprehension, preserves perceived autonomy, or contributes to other clearly defined patient or professional needs.

Such an approach would also create a stronger connection between technological governance and measurement. If the objective of an AI application is explicitly defined in relation to stakeholder utility, organizations can evaluate the system against that objective rather than relying exclusively on generic engagement metrics. This is particularly important in pharmaceutical communication because high engagement does not necessarily indicate high informational quality, ethical appropriateness, or patient value. A highly engaging message may still be misleading, intrusive, emotionally manipulative, or clinically irrelevant.

In regulated health communication, engagement metrics alone cannot establish informational quality, regulatory appropriateness, or stakeholder value, particularly where AI-mediated communication intersects with scientific evidence and patient vulnerability (OECD, 2023; Jones et al., 2024).

Accordingly, the comparative cases suggest that responsible pharmaceutical AI strategy should distinguish between what attracts stakeholder attention and what creates stakeholder value. These outcomes may overlap, but they should not be treated as interchangeable. The Data + Empathy framework places particular emphasis on the latter: data and computational intelligence provide the analytical and operational capacity, while empathy requires organizations to evaluate those capabilities from the perspective of the stakeholder's actual informational, functional, and emotional circumstances.

This finding also contributes to the broader argument concerning trust developed throughout the thesis. Trust should not be treated as an automatic consequence of either technological innovation or human-centered branding. Neither *Scrolling Therapy* nor *The Most Beautiful Sound* empirically

demonstrates that functional AI necessarily increases trust, just as *The Code* does not establish that synthetic representation necessarily destroys it. The more defensible conclusion is that clear functional purpose, appropriate technological boundaries, transparency, and alignment with stakeholder needs provide conditions under which trustworthy communication may be more plausibly supported.

For life sciences organizations, the resulting strategic implication is therefore not to replace generative persuasion with functional AI in every circumstance. It is to establish a more disciplined hierarchy of technological objectives. Before asking how Artificial Intelligence can generate more content, personalize more interactions, or automate more communication, organizations should identify the stakeholder problem to be addressed and determine whether the proposed computational intervention creates substantive value in relation to that problem.

The second cross-case finding can consequently be summarized as a transition from technology-led communication to stakeholder-led technological deployment. Eurofarma and ASCO illustrate how computational systems can be designed around functional or interpretative needs, while Dove demonstrates that organizations may also create value by establishing boundaries when technological capability conflicts with representational priorities. Within the Data + Empathy framework, technological innovation becomes strategically meaningful not when it maximizes synthetic output, but when its role is proportionate to the human problem it is intended to address.

This principle provides a direct transition to the third comparative finding. If the strategic value of Artificial Intelligence depends on the stakeholder problem it addresses rather than on the amount of automation it enables, the next analytical question concerns how human-centered value itself is defined and operationalized across different stakeholder experiences. This issue is examined in the following section through the relationship between computational capability, empathy, and stakeholder experience.

3. Human-Centered Value Creation and Stakeholder Experience

A third comparative finding concerns the role of human-centered value creation in shaping the stakeholder experience of AI-mediated communication. If the previous finding established that technological deployment should be evaluated according to the stakeholder problem it is intended to address, the present comparison examines how that value is translated into the design of the interaction itself. Across the three cases, empathy does not emerge as a standardized technological feature or as the capacity of an artificial system to simulate human emotion. Rather, it can be interpreted as an organizational and design principle through which computational capabilities are configured around the functional, informational, emotional, and representational circumstances of the people involved.

This distinction is particularly important within the Data + Empathy framework. The term *empathy* should not imply that Artificial Intelligence possesses human emotional understanding or that increasingly anthropomorphic systems necessarily produce more empathetic interactions. Research on algorithmic judgment and human responses to automated systems indicates that the relationship between technological performance and user acceptance is considerably more complex, particularly when algorithms are perceived as imperfect or are assigned evaluative roles traditionally associated with human judgment (Dietvorst et al., 2015). Similarly, research on the *Uncanny Valley of Mind*

suggests that attributing human-like mental characteristics to artificial agents can produce ambivalent responses rather than automatically strengthening perceived human connection (Stein & Ohler, 2017). Within the framework proposed in this thesis, empathy is therefore located primarily in the design and governance of the technological interaction, rather than attributed to the computational system itself.

This conceptualization becomes visible in *Scrolling Therapy*. The human-centered dimension of the initiative does not derive from an AI system attempting to reproduce empathetic conversation or emotional behavior. Instead, it emerges from the way in which the technological architecture responds to a specific difficulty within the patient's rehabilitation experience. By translating facial exercises into social-media navigation commands, the application restructures the relationship between the rehabilitation task and an activity already familiar to many users (Dentsu Creative, 2023). The empathetic component can therefore be interpreted as a design response to behavioral and emotional friction rather than as a property of the computer-vision model.

This distinction contributes to a broader understanding of patient-centered technological innovation. In this configuration, the system does not need to appear more human in order to become more responsive to human circumstances. Its value instead derives from recognizing that the experience of a rehabilitation exercise is influenced not only by its clinical objective but also by the context in which it is performed, the effort required to sustain it, and the emotional experience associated with repetition. Consequently, the relevant human-centered question is not simply whether the underlying computer vision operates accurately, but whether the interaction has been structured in a manner that corresponds to the practical circumstances of the person expected to use it.

The Most Beautiful Sound illustrates a different dimension of human-centered value creation. In this case, the central issue is not behavioral friction but the perceptual distance between highly specialized biological information and the individuals attempting to understand or emotionally relate to it. Through sonification, experimentally observed cellular phenomena are computationally translated into an auditory representation (Grey Health & Wellness, 2023). The technological system therefore operates as an intermediary between scientific measurement and human sensory experience.

Importantly, the human-centered value of this translation should not be interpreted as evidence that the auditory representation improves clinical outcomes, reduces anxiety, or increases trust. The available case material does not establish these effects. Its analytical significance instead lies in the attempt to make a complex and otherwise imperceptible process more accessible through a sensory modality that differs from conventional textual or visual scientific communication. In this context, empathy can be understood as an effort to consider how scientific information is experienced by the recipient, while preserving the distinction between the underlying biological evidence and its communicative representation.

This distinction is also consistent with the broader communication problem addressed by the Heuristic-Systematic Model. Chaiken's (1980) framework differentiates between relatively effortful systematic processing and forms of judgment based on more accessible informational cues. The ASCO case should not be interpreted as demonstrating that auditory communication necessarily produces heuristic processing, nor that sensory representation should substitute systematic engagement with scientific evidence. Rather, the case illustrates how an additional representational

layer can make complex information more perceptually accessible while the scientific evidence itself remains available for more systematic interpretation. Human-centered communication, from this perspective, concerns not the simplification of evidence to the point of distortion, but the design of multiple pathways through which stakeholders may encounter complexity.

Dove's *The Code* introduces a third dimension of stakeholder experience: representation and authenticity. Unlike Eurofarma and ASCO, Dove does not primarily use computational technology to facilitate a patient activity or translate scientific information. Instead, the initiative interrogates the consequences of using generative systems to construct representations of human appearance and establishes an explicit organizational boundary concerning the synthetic representation of women in its communication (Dove, 2024; Unilever, 2024).

The relevance of this case to human-centered value creation lies in the recognition that stakeholder experience is shaped not only by the informational content of communication but also by who is represented, how that representation is constructed, and whether the audience can understand its provenance. In digitally mediated environments, representations of appearance can interact with broader questions concerning identity, self-perception, social comparison, and psychological well-being, particularly among younger populations exposed to highly visual digital environments (Mansfield et al., 2025). This does not establish that AI-generated beauty imagery directly causes specific psychological harms. It does, however, reinforce the importance of treating digital representation as a meaningful component of the broader environment within which stakeholder experiences are formed.

The three initiatives therefore reveal different dimensions of what this thesis defines as empathy:

Eurofarma: Empathy as Friction Reduction

Human-centered design responds to the practical and emotional circumstances surrounding a rehabilitation activity.

ASCO: Empathy as Translational Accessibility

Computational processing creates an alternative sensory interface through which complex scientific information can become more perceptually accessible.

Dove: Empathy as Representational Responsibility

Organizational governance considers whether technological representation is consistent with authenticity, diversity, and the human subjects represented by the communication.

These configurations are analytically distinct, but they share an important characteristic. In none of the three cases does empathy require the computational system itself to become an autonomous emotional actor. Instead, empathy is introduced through decisions made around the system: which problem is selected, which technological function is deployed, how the resulting interaction is structured, what information is made accessible, which representations are considered appropriate, and where technological boundaries are established.

This observation is particularly relevant to pharmaceutical communication because the pursuit of increasingly human-like Artificial Intelligence can easily be conflated with the pursuit of more human-centered communication. The two objectives are not equivalent. Conversational interfaces, synthetic voices, virtual humans, personalized language generation, and other anthropomorphic

technologies may create more naturalistic forms of interaction, but their human-like appearance or behavior does not independently establish that the underlying communication is empathetic, appropriate, or trustworthy. The broader literature on algorithm aversion and anthropomorphic artificial agents cautions against assuming a linear relationship between technological human-likeness and positive user evaluation (Dietvorst et al., 2015; Stein & Ohler, 2017).

Within a healthcare context, this distinction acquires additional importance because communication frequently occurs under conditions of informational asymmetry, uncertainty, and personal vulnerability. Stakeholders may be required to interpret complex scientific information, evaluate treatment-related choices, understand potential risks, or navigate interactions involving sensitive personal data. Consequently, the human-centered quality of an AI-mediated interaction cannot be evaluated exclusively through conversational fluency, personalization accuracy, or engagement metrics. Responsible AI frameworks similarly emphasize the importance of transparency, accountability, human oversight, and attention to the individuals affected by automated systems (OECD, 2023).

The Data + Empathy framework therefore proposes a broader interpretation of personalization. Conventional computational personalization can be understood primarily in terms of **relevance**: selecting the content, channel, timing, or message most likely to correspond to an individual's predicted characteristics or behavior. AI-driven pharmaceutical commercial operations increasingly enable these forms of data-informed targeting and engagement (Sabapathy, 2025; Verma et al., 2024). However, relevance alone does not establish human-centered value.

A message may be highly relevant from the perspective of an algorithmic prediction while remaining intrusive, poorly timed, emotionally inappropriate, or insufficiently transparent from the perspective of the recipient. Conversely, a less extensively personalized intervention may create substantial stakeholder value if it responds appropriately to a meaningful functional or informational need. This distinction reinforces the personalization paradox discussed in Chapter 3: increasing the computational precision of communication does not necessarily produce a proportional increase in perceived stakeholder benefit, particularly where personalization requires extensive data processing or creates uncertainty concerning how individual information is being used (OECD, 2023; Verma et al., 2024).

Within the framework proposed here, Data and Empathy therefore perform complementary but non-interchangeable functions. Data provides the informational basis through which organizations can identify patterns, segment stakeholders, predict potential needs, retrieve relevant evidence, and adapt communication. Empathy introduces a second evaluative layer by asking whether the resulting intervention is appropriate to the human circumstances in which that information will be received.

This relationship can be conceptualized as:

Data → What can be known, predicted, or adapted

Empathy → What should be communicated, how it should be experienced, and where technological intervention should stop

The distinction is intentionally normative. Data can identify that a stakeholder is statistically likely to respond to a particular message, but it cannot independently determine whether delivering that

message is ethically appropriate. A predictive model may identify an optimal moment for engagement, but optimization does not automatically establish that intervention at that moment respects autonomy, privacy, emotional circumstances, or the broader context of care. These questions require governance criteria that extend beyond predictive accuracy and commercial performance (OECD, 2023).

For this reason, empathy within the proposed framework should not be reduced to emotional tone. A communication can use warm, reassuring, or personalized language while remaining structurally insensitive to the stakeholder's needs. Conversely, a highly technical intervention may be human-centered even when it contains no explicit emotional language. *Scrolling Therapy* illustrates this distinction particularly clearly: the computer-vision architecture does not need to communicate sympathy in order for the design of the application to respond to a human difficulty. Similarly, the computational sonification used in *The Most Beautiful Sound* does not itself possess emotional understanding; the human-centered dimension derives from the decision to create a more accessible sensory relationship with complex biological information.

This distinction has important implications for the role of generative AI in pharmaceutical communication. Large language models and other generative systems can increasingly produce linguistically fluent, context-sensitive, and apparently personalized responses. Such capabilities may be useful in appropriately governed applications involving information retrieval, content adaptation, workflow support, or stakeholder education. Nevertheless, linguistic fluency should not be treated as equivalent to empathy. A system can generate language that appears compassionate without possessing contextual understanding of the patient's lived circumstances or responsibility for the consequences of the interaction.

The practical implication is not that AI-generated communication must always remain invisible or emotionally neutral. Rather, the level of human involvement should correspond to the sensitivity, uncertainty, and potential consequences of the communication task. Low-risk informational functions may permit higher levels of automation, while interactions involving clinical interpretation, emotionally sensitive circumstances, contested evidence, or consequential decisions may require stronger forms of human oversight. This principle is consistent with the responsible-governance logic developed in Chapter 4, in which Human-in-the-Loop mechanisms complement automated systems rather than treating human review as an obstacle to scalability (OECD, 2023).

The cross-case comparison also suggests that stakeholder experience should be evaluated as a multidimensional construct. Traditional digital marketing indicators frequently capture observable behavioral responses such as reach, interaction, click-through behavior, or conversion. AI-enabled commercial architectures can make these indicators increasingly granular and actionable (Sabapathy, 2025). However, behavioral response represents only one dimension of the interaction.

For healthcare and life sciences organizations, a more human-centered assessment could additionally consider whether communication is understandable, whether the provenance of information is sufficiently transparent, whether the technological interaction preserves meaningful human agency, whether the stakeholder perceives the intervention as relevant without experiencing it as intrusive, and whether representational choices are consistent with the lived realities of the people depicted.

These dimensions should not be assumed to improve simply because AI is introduced; they would require explicit measurement and, where appropriate, empirical validation.

This point is methodologically important for the interpretation of the three cases. *Scrolling Therapy* does not prove that integrating rehabilitation with social-media interaction produces superior long-term adherence. *The Most Beautiful Sound* does not demonstrate that sonification systematically increases patient comprehension, hope, or trust. *The Code* does not establish that limiting synthetic representation produces greater consumer trust or improved psychological outcomes. The cases instead provide observable design and governance choices that can be interpreted through the Data + Empathy framework. Their outcomes should therefore not be extended beyond what the available evidence supports.

This methodological restraint strengthens rather than weakens the comparative analysis. The purpose of the cases is to identify patterns of technological role allocation and human-centered design that can subsequently generate more precise questions for empirical research. Future studies could examine, for example, whether utility-oriented AI applications are perceived as more trustworthy than explicitly persuasive applications, whether sensory translation improves comprehension of complex scientific processes, or whether transparent boundaries around synthetic representation affect perceptions of authenticity. Such questions represent testable extensions of the conceptual framework rather than conclusions already established by the cases.

From an organizational perspective, the finding also implies that empathy cannot be delegated exclusively to creative teams or inserted at the final stage of communication design. If empathy concerns the appropriateness of technological intervention, it must influence decisions throughout the AI lifecycle: problem definition, data collection, model selection, content generation, channel orchestration, compliance review, human oversight, deployment, and post-deployment evaluation. This interpretation connects the Data + Empathy framework directly with the Compliance by Design approach developed in Chapter 4.

A pharmaceutical organization applying this principle would therefore evaluate an AI initiative through at least two complementary questions. The first is computational: Can the system perform the intended function accurately, efficiently, and at the required scale? The second is human-centered: Should this function be automated in this context, and does the resulting interaction create meaningful value for the stakeholder? Neither question is sufficient independently.

This produces a broader decision logic for the framework:

Stakeholder Context → Data-Informed Understanding → Technological Capability → Empathetic Evaluation → Human Oversight → Stakeholder Experience

The sequence is important because empathy operates as an evaluative layer rather than as a decorative feature added after technological development. A technically successful intervention may be reconsidered if its deployment conflicts with stakeholder autonomy, representational integrity, privacy expectations, or the sensitivity of the communication context. Conversely, a comparatively simple technological intervention may represent a stronger human-centered solution when it directly addresses a meaningful stakeholder difficulty.

The three cases illustrate these possibilities through different strategic configurations. Eurofarma shows how computational precision can be directed toward reducing friction within a specific patient-related activity. ASCO illustrates how computation can mediate between scientific complexity and human sensory accessibility. Dove demonstrates that human-centered technological governance may also require an organization to establish explicit limits when synthetic generation conflicts with its representational objectives. Across all three cases, the central strategic variable is therefore not technological sophistication but the quality of the relationship established between technological capability and human context.

The third cross-case finding can consequently be summarized as a transition from personalized communication to context-sensitive value creation. Data-driven systems can help organizations understand and respond to stakeholders with increasing precision, but empathy determines whether those capabilities are deployed in ways that remain proportionate to human needs, vulnerabilities, expectations, and representational concerns. Within the Data + Empathy framework, empathy is therefore not positioned in opposition to computational intelligence. It operates as the human-centered governance layer through which computational intelligence acquires strategic and ethical direction.

This interpretation also clarifies one of the central propositions of the thesis: the objective is not to make Artificial Intelligence more human, but to make the organizational use of Artificial Intelligence more human-centered. Such a distinction preserves the analytical value of computational systems while avoiding the assumption that technological anthropomorphism is equivalent to empathy. It also provides the conceptual bridge to the fourth and final cross-case finding. Once empathy is understood as a principle that influences technological boundaries and organizational decision-making, its implementation necessarily becomes a question of governance. The following finding therefore examines how proactive ethical boundary-setting can operate under conditions in which technological development advances more rapidly than formal regulatory adaptation.

4. Proactive Ethical Governance under Conditions of Regulatory Lag

The fourth comparative finding concerns the role of **proactive ethical governance when technological development advances more rapidly than the regulatory frameworks designed to oversee its application**. The three initiatives examined in this chapter operate in substantially different regulatory and organizational contexts, yet their comparison highlights a common strategic issue: the absence of an explicit legal prohibition does not necessarily establish that every technically feasible application of Artificial Intelligence is appropriate for stakeholder communication. Responsible technological deployment therefore requires organizations to determine not only what they are legally permitted to do, but also what forms of automation, representation, and data processing are consistent with their scientific responsibilities, stakeholder relationships, and institutional values.

This distinction reflects the broader problem of **regulatory lag** discussed in Chapter 4. Emerging technologies frequently evolve more rapidly than the legal and ethical structures intended to govern their use, producing what Marchant et al. (2011) describe as a *pacing problem* between technological innovation and regulatory oversight. Artificial Intelligence intensifies this challenge because its applications can develop across multiple organizational functions simultaneously, including content generation, audience segmentation, predictive analytics, automated decision support, conversational

interfaces, and synthetic media. Consequently, compliance with existing rules remains necessary, but formal compliance alone may not provide sufficient guidance for every novel technological configuration.

The pharmaceutical environment makes this distinction particularly important. Communication concerning medicinal products already operates within extensive regulatory constraints. In the European Union, Directive 2001/83/EC establishes specific requirements concerning medicinal-product advertising, including restrictions on advertising prescription-only medicines to the general public and requirements that advertising remain consistent with the authorized characteristics of the medicinal product. In the United States, prescription-drug advertising is similarly subject to specific FDA requirements, including provisions concerning the presentation of information about benefits and risks (Directive 2001/83/EC, 2001; 21 C.F.R. § 202.1). However, the increasing use of AI introduces additional questions concerning algorithmic personalization, synthetic representation, automated content generation, data provenance, and human oversight that may not always be resolved through rules originally developed for earlier communication environments.

Regulatory lag should therefore not be interpreted as regulatory absence. Existing pharmaceutical advertising, privacy, consumer-protection, and data-governance requirements continue to apply where relevant. Rather, the concept identifies a potential **gap between the general obligations established by existing law and the specific operational decisions required by emerging technological applications**. Organizations may know that communication must not be misleading, that personal data must be processed according to applicable legal principles, or that promotional claims require appropriate substantiation, while still facing new questions concerning how much decision-making authority should be delegated to an algorithm, when synthetic representation requires disclosure, or which AI-generated outputs should undergo mandatory human review.

This gap creates a strategic role for proactive governance. Instead of treating regulation as the final stage of technological deployment, organizations can integrate ethical and compliance considerations into the design of AI-enabled processes from the beginning. This logic is consistent with the **Compliance by Design** approach developed in Chapter 4, in which regulatory requirements, evidence standards, data-governance principles, and human accountability are incorporated into technological architectures and organizational workflows rather than applied exclusively as retrospective controls.

The three cases examined in this chapter illustrate different dimensions of this principle. *Scrolling Therapy* provides an example at the level of technological architecture. According to the campaign documentation, facial-movement processing was designed to occur locally on the user's device rather than requiring the external transmission of the relevant facial data (Dentsu Creative, 2023). This design choice should not be interpreted as evidence of formal GDPR certification or as proof that the application satisfies every potentially applicable privacy requirement. Nevertheless, it can be interpreted as a **privacy-oriented architectural decision** that reduces the need for external data transmission and is conceptually consistent with the GDPR principle of data minimization, according to which personal data should be adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed (Regulation (EU) 2016/679, 2016).

The significance of this example extends beyond the specific application. Privacy governance can occur not only through consent notices, privacy policies, or legal documentation, but also through

decisions concerning the technical architecture itself. If a function can be performed locally without transferring potentially sensitive information to external infrastructure, the architecture may reduce certain forms of data exposure by design. This does not eliminate the need for legal assessment, security controls, transparency, or appropriate data governance. It demonstrates instead that **ethical and regulatory considerations can influence technological design before stakeholder data enter broader organizational systems**.

The Most Beautiful Sound illustrates proactive governance through a different problem: the relationship between **scientific evidence and communicative representation**. Sonification transforms measured biological information into an auditory form, creating a representation that may be more perceptually accessible to non-specialist audiences (Grey Health & Wellness, 2023). Yet the communicative power of such translation also creates a responsibility to preserve the distinction between the underlying biological phenomenon and the sensory representation constructed from it.

This distinction is essential because computational translation can increase the experiential accessibility of scientific information without increasing the evidentiary strength of the underlying data. An auditory representation of a biological process does not independently establish therapeutic effectiveness, predict an individual patient's outcome, or substitute clinical interpretation. Proactive governance in this context therefore concerns **epistemic boundaries**: organizations must ensure that the communicative form through which evidence is presented does not encourage conclusions that exceed what the underlying scientific information can support.

This principle has broader implications for AI-mediated pharmaceutical communication. Generative and computational systems can translate complex scientific material into simplified text, visualizations, conversational explanations, personalized educational content, or alternative sensory formats. These capabilities may increase accessibility and operational scalability (Jones et al., 2024; Sabapathy, 2025; Verma et al., 2024). At the same time, every additional layer of algorithmic transformation creates a need to preserve traceability between the generated communication and the evidence from which it originates.

The governance problem is therefore not limited to whether an AI-generated statement is linguistically accurate. Organizations must also consider whether the output preserves the appropriate level of scientific uncertainty, whether relevant contextual information has been omitted, whether the source material remains identifiable, and whether the resulting communication could imply a degree of certainty unsupported by the evidence. In regulated health communication, computational fluency cannot substitute evidentiary integrity.

Dove's *The Code* provides the clearest example among the three cases of **explicit organizational boundary-setting**. Rather than waiting for a comprehensive regulatory framework to determine every acceptable use of generative imagery, Dove publicly established a commitment concerning the use of AI-generated representations of women in its advertising and complemented this position with guidance intended to encourage more representative prompting practices (Dove, 2024; Unilever, 2024). As established earlier, this initiative originates outside the pharmaceutical sector and should not be interpreted as a direct regulatory model for medicinal-product communication. Its relevance lies instead in the governance principle it illustrates.

The principle is that an organization may voluntarily establish technological boundaries that are **more specific than the minimum requirements currently articulated by law** when particular applications conflict with its representational objectives or stakeholder commitments. This does not mean that private corporate policies can replace statutory regulation. Self-regulation has structural limitations, including differences in enforcement, accountability, consistency, and organizational incentives. Nevertheless, voluntary boundary-setting can function as an additional governance layer while formal regulatory systems continue to adapt to emerging technological practices.

Dove's decision is especially relevant because it challenges an assumption that can accompany rapid technological adoption: that organizational innovation should be measured according to the extent to which available capabilities are deployed. From a governance perspective, technological maturity may instead include the capacity to determine **when not to automate**. The existence of a technically feasible generative function does not create an organizational obligation to use it.

This principle can be cautiously transferred to pharmaceutical and life-sciences communication. An organization may determine, for example, that generative AI is appropriate for internal drafting assistance but not for autonomous publication of regulated promotional claims; appropriate for evidence retrieval but not for independent clinical interpretation; appropriate for generating non-human illustrative assets but inappropriate for undisclosed synthetic patient testimonials; or appropriate for supporting content adaptation while requiring human approval before external dissemination. The precise boundaries would necessarily depend on jurisdiction, use case, risk profile, organizational policy, and applicable regulation. The broader principle is that **responsible AI governance requires explicit decisions about both permitted and restricted technological roles**.

This represents an important evolution from a purely reactive model of compliance. Under a reactive approach, technological systems are developed or deployed first and subsequently reviewed to determine whether their outputs violate existing rules. Under a proactive model, governance criteria participate in defining the technological architecture itself.

The distinction can be represented as follows:

Reactive Compliance Logic:

[Technological Capability] → [Deployment] → [Output] → [Regulatory Review] → [Correction if Necessary]

Proactive Governance Logic:

[Stakeholder and Regulatory Context] → [Risk Assessment] → [Permitted Technological Role] → [Governance Controls] → [Deployment] → [Continuous Human Oversight and Evaluation]

The second configuration does not eliminate retrospective review. Rather, it introduces governance earlier in the process and maintains it throughout the technological lifecycle. Such an approach becomes increasingly important as AI systems participate in dynamic processes in which content, recommendations, or decisions may be generated at a scale that makes exclusively manual ex-post control increasingly difficult.

This is where **Human-in-the-Loop governance** becomes particularly relevant. Human oversight should not be interpreted simply as requiring an employee to approve every computational operation

manually. Such an approach could eliminate many of the operational benefits that AI is intended to provide. Instead, the appropriate degree and location of human intervention should correspond to the potential consequences, uncertainty, regulatory sensitivity, and reversibility of the automated function. Responsible AI principles similarly emphasize human oversight, accountability, transparency, and risk-sensitive governance rather than unrestricted technological autonomy (OECD, 2023).

For pharmaceutical communication, this implies that different AI functions may require different governance intensities. An internal system used to organize non-sensitive content may justify a different level of supervision from a system generating externally visible information about treatment benefits and risks. Similarly, an algorithm supporting evidence retrieval may operate under different controls from an automated system capable of modifying patient-facing messages in real time. Human oversight should therefore be risk-proportionate rather than uniform.

This logic reinforces the functional separation developed in Chapter 4 between computational backend capabilities and human-governed stakeholder interaction. Predictive analytics, Retrieval-Augmented Generation, automated pre-screening, Dynamic Creative Optimization, and Next Best Action systems can contribute to evidence retrieval, content selection, compliance support, and communication orchestration (Jones et al., 2024; Sabapathy, 2025; Verma et al., 2024). Their inclusion within the communication architecture, however, does not require that each system independently determine the final form or appropriateness of stakeholder-facing communication.

Instead, the Data + Empathy framework positions governance as the mechanism through which computational capability is translated into organizationally accountable action. **Data** provides analytical precision, evidence retrieval, pattern recognition, and predictive capability. **Empathy** introduces consideration of stakeholder context, vulnerability, autonomy, representation, and communicative appropriateness. **Governance** determines the institutional rules through which the relationship between these dimensions is operationalized.

The relationship can therefore be conceptualized as:

Data → Computational Possibility

Empathy → Human-Centered Appropriateness

Governance → Organizational Permission, Limitation, and Accountability

This three-layer interpretation extends the Data + Empathy framework beyond a communication philosophy and toward an operational governance model. Data without governance may create technically sophisticated systems whose deployment boundaries remain unclear. Empathy without governance may remain an aspirational organizational value without consistent implementation. Governance provides the procedures, responsibilities, controls, and escalation mechanisms through which both computational capability and human-centered principles can influence actual organizational behavior.

Such governance also requires attention to **accountability**. If an AI-generated communication is misleading, inappropriate, discriminatory, or inconsistent with regulatory requirements, responsibility cannot be transferred abstractly to the algorithm. The organization deploying the system remains

responsible for defining its intended use, establishing appropriate controls, determining acceptable levels of risk, and ensuring that human professionals retain meaningful authority where consequential decisions are involved. Responsible AI frameworks similarly position accountability as an organizational requirement rather than a characteristic that can be delegated to automated systems (OECD, 2023).

This principle is particularly relevant to generative AI because probabilistic systems can produce outputs that are fluent and persuasive while still containing inaccurate, unsupported, or contextually inappropriate information. The possibility of such outputs makes source traceability, evidence grounding, pre-deployment testing, human review, and post-deployment monitoring important components of a responsible communication architecture. The objective is not necessarily to eliminate every possible error—an unrealistic standard for either automated or human communication—but to establish processes capable of identifying, preventing, escalating, and correcting material risks.

The three cases provide complementary illustrations of how such governance thinking can operate at different levels:

Eurofarma: Architectural Governance

Privacy-related considerations can influence how and where computational processing occurs.

ASCO:Epistemic Governance

The transformation of scientific information into an experiential format requires clear boundaries between underlying evidence and communicative representation.

Dove: Representational Governance

An organization can establish explicit limits on technological use when synthetic generation conflicts with its objectives concerning authentic human representation.

These categories should not be interpreted as exhaustive or mutually exclusive. Eurofarma also raises representational and clinical questions; ASCO operates within institutional and scientific governance structures; and Dove's initiative has implications beyond representation. Their value lies in showing that proactive governance is **multidimensional**. It can be embedded in system architecture, evidence-management practices, communication design, corporate policies, human-review processes, and decisions to prohibit particular technological uses.

The comparison also demonstrates why regulation and ethics should not be treated as synonymous. Regulation establishes legally enforceable obligations, whereas ethical governance may require organizations to evaluate situations in which the law permits multiple possible courses of action. An organization can therefore remain formally compliant while still facing legitimate questions concerning transparency, autonomy, authenticity, fairness, or the proportionality of technological intervention. Conversely, ethical intentions cannot justify non-compliance with binding legal requirements. The two dimensions are complementary rather than interchangeable.

This distinction becomes particularly important under conditions of regulatory lag. When technological capability develops more rapidly than technology-specific rules, organizations may face a temptation to interpret legal ambiguity as strategic freedom. The Data + Empathy framework proposes a different approach: ambiguity should trigger structured governance rather than

unrestricted experimentation, particularly where communication involves vulnerable stakeholders, health information, therapeutic expectations, or consequential decisions.

Such a position does not advocate technological conservatism. Excessively restrictive governance can prevent beneficial innovation, reduce organizational learning, and delay applications capable of improving accessibility, efficiency, or stakeholder support. The objective is therefore not to minimize AI deployment but to create proportionate governance in which the intensity of controls corresponds to the potential impact of the technological function.

A low-risk internal application may require relatively limited intervention, whereas a patient-facing system capable of producing individualized treatment-related communication would justify substantially stronger evidence controls, transparency requirements, human supervision, and monitoring. Between these extremes lies a spectrum of applications that should be evaluated according to their specific context rather than classified through a binary distinction between “safe” and “unsafe” AI.

This risk-proportionate logic also strengthens the strategic relevance of Compliance by Design. If governance is incorporated at the architecture stage, organizations can potentially avoid the operational inefficiency of repeatedly correcting inappropriate outputs after deployment. Automated pre-screening, controlled evidence repositories, predefined content boundaries, provenance mechanisms, escalation pathways, and human approval requirements can be integrated into the communication process according to the risk profile of the use case. In this configuration, compliance does not simply restrict innovation; it contributes to defining the conditions under which innovation can scale responsibly.

The same principle applies to the measurement of AI performance. Operational metrics such as speed, cost reduction, content volume, engagement, or conversion remain relevant to organizational decision-making (Sabapathy, 2025; Verma et al., 2024). Yet responsible governance requires additional indicators capable of capturing the quality and consequences of technological deployment. Depending on the use case, these may include the frequency and severity of unsupported outputs, human-override rates, evidence-traceability performance, privacy incidents, stakeholder complaints, representational disparities, or the proportion of generated material requiring substantive correction before approval.

Importantly, these indicators are proposed here as governance dimensions rather than empirically validated universal metrics. Their appropriate operationalization would require further research and organizational testing. Nevertheless, incorporating such measures would allow organizations to evaluate AI systems not only according to what they produce, but also according to how reliably, transparently, and responsibly they operate within the intended communication environment.

The cross-case analysis therefore reinforces an important principle: responsible innovation is not defined by maximum technological adoption, but by the capacity to make explicit, evidence-informed, and accountable decisions about technological deployment. In *Scrolling Therapy*, this is reflected in the architecture through which a functional computational capability is implemented. In *The Most Beautiful Sound*, it emerges through the need to preserve the epistemic boundary between scientific evidence and sensory translation. In *The Code*, it becomes visible through deliberate organizational restraint concerning synthetic representation.

These cases do not establish that the respective organizations have resolved every ethical or regulatory challenge associated with Artificial Intelligence. Nor do they empirically validate a universal governance model. Instead, they illustrate different organizational choices that can be interpreted through the Data + Empathy framework and used to identify broader principles for pharmaceutical communication.

The fourth finding can therefore be summarized as a transition from regulatory compliance as a downstream control to responsible governance as an upstream design principle. Under conditions of regulatory lag, pharmaceutical organizations cannot assume that the absence of a technology-specific prohibition is sufficient justification for deployment. They must instead define technological roles, establish risk-proportionate human oversight, preserve evidence integrity, incorporate privacy and transparency considerations, and determine explicit boundaries where automation may conflict with stakeholder needs or institutional responsibilities.

Within the Data + Empathy framework, this produces a final governance logic:

[Technological Possibility] → [Stakeholder and Regulatory Context] → [Data + Empathy Evaluation] → [Risk-Proportionate Governance] → [Permitted Technological Role] → [Human-Governed Deployment] → [Continuous Evaluation]

The four cross-case findings can consequently be understood as sequential components of a single strategic logic. The first establishes where AI should be positioned and how much authority it should receive; the second determines whether technological capability creates substantive stakeholder utility rather than merely increasing communicative capacity; the third defines human-centered value through context-sensitive empathy rather than technological anthropomorphism; and the fourth translates these principles into organizational boundaries, accountability mechanisms, and governance processes.

Taken together, the comparison does not demonstrate that the Data + Empathy framework is empirically superior to alternative models. Rather, it provides an applied interpretation of how the framework can organize strategic decisions across heterogeneous technological contexts. The central proposition emerging from the three cases is therefore not that pharmaceutical organizations should deploy more or less Artificial Intelligence in absolute terms. It is that they should become more deliberate about which capabilities are automated, why they are automated, how they create stakeholder value, where human judgment remains necessary, and under which conditions technological restraint represents the more responsible strategic choice.

5.2 Applying the Data + Empathy Framework to AI-Driven Pharma Communication

The comparative analysis conducted in Section 5.1 suggests that the responsible integration of computational technologies into healthcare communication depends not on the indiscriminate expansion of automated systems, but on deliberate, context-sensitive role allocation (Jones et al., 2024; Sabapathy, 2025; Verma et al., 2024). The cross-case findings established in Section 5.1.4 indicate that computational capability acquires strategic relevance when it is positioned as an enabling infrastructure, when it is evaluated against substantive stakeholder utility rather than capacity-driven persuasion, when it prioritizes human-centered value over technological anthropomorphism, and when it is governed through proactive, upstream design choices under conditions of regulatory lag (Marchant et al., 2011; OECD, 2023).

Building directly upon these applied and theoretical insights, this section operationalizes the core conceptual proposal of this thesis: the **Data + Empathy framework**. Rather than conceptualizing Data + Empathy as an abstract corporate slogan or an unconstrained communication strategy, this research models it as a structured governance philosophy and operational architecture for pharmaceutical digital communication.

Crucially, the framework does not establish an absolute technological barrier between invisible algorithms and visible human interaction, nor does it posit a simplistic dichotomy in which computational visibility is inherently detrimental and human intervention is universally required. Such a binary division would overlook legitimate low-risk applications of automated communication—such as basic navigation support, factual directory retrieval, or administrative scheduling—while failing to reflect the complex operational realities of modern omnichannel healthcare networks (Sabapathy, 2025; Verma et al., 2024).

Instead, this thesis proposes a **risk-proportionate functional asymmetry**. Under this model, computational systems assume a stronger operational role in data-intensive, analytical, retrieval, pattern-recognition, and compliance-support functions, where machine-learning architectures can provide substantial processing capacity and structural consistency (Gao et al., 2023; Sabapathy, 2025; Zheng et al., 2023). Progressively greater human oversight, contextual interpretation, and communicative stewardship are retained as interactions become more clinically consequential, emotionally sensitive, ethically contested, or legally accountable (OECD, 2023; Verma et al., 2024).

This functional asymmetry reflects the tripartite governance logic established in the preceding synthesis:

Data → Computational Possibility

Empathy → Human-Centered Appropriateness


Governance → Organizational Permission, Limitation, and Accountability

Within this architecture, *Data* determines what can be computationally processed, predicted, retrieved, or adapted; *Empathy* introduces an evaluative layer through which the appropriateness of technological intervention is considered in relation to stakeholder circumstances; and *Governance* establishes the procedural controls, verification mechanisms, organizational boundaries, and accountability structures through which computational capability is translated into responsible communication (OECD, 2023; Sabapathy, 2025).

Empathy should therefore not be interpreted as an additional computational module capable of autonomously determining human vulnerability. Rather, it represents a human-centered evaluative principle through which organizations assess whether the proposed technological intervention corresponds appropriately to the informational, functional, emotional, and ethical circumstances of the stakeholder. Governance subsequently converts that evaluation into concrete decisions concerning technological permission, restriction, oversight, and accountability.

The following subsections formalize this governance architecture, examine its cognitive and behavioral foundations, and construct an operational implementation matrix for life sciences organizations.

The Data + Empathy Governance Architecture: Risk-Proportionate Functional Asymmetry and Oversight Gradient



5.2.1 Architectural Configuration: Risk-Proportionate Functional Asymmetry

The architectural configuration of the Data + Empathy framework addresses the operational realities of modern life sciences communication by structuring the relationship between machine capabilities and human responsibilities (Jones et al., 2024; Sabapathy, 2025). Rather than imposing a blanket prohibition on automated stakeholder interfaces, the model establishes a structured functional asymmetry: computational architectures are deployed where statistical pattern recognition, data ingestion, retrieval, and scalable information processing provide clear operational value, while

progressively greater human stewardship is introduced where clinical nuance, emotional sensitivity, ethical interpretation, and legal accountability become more consequential (OECD, 2023; Verma et al., 2024).

The model should therefore be interpreted as a continuum rather than a binary division. The appropriate allocation of computational and human functions depends on the specific use case, the sensitivity of the information involved, the characteristics of the stakeholder, the degree of uncertainty associated with the communication, and the potential consequences of an inaccurate or inappropriate output. In this sense, the architecture operationalizes the principle established in Section 5.1.4: the relevant question is not whether Artificial Intelligence should be visible or invisible in absolute terms, but **what degree of visibility, autonomy, and authority is proportionate to the communicative context**.

The Computational Enablement Layer (Backend Functions)

Within the proposed framework, the computational infrastructure operates primarily as an analytical and operational foundation. It is designed to support data engineering, information retrieval, pattern recognition, and structural compliance tasks across four primary domains.

1. Privacy-Preserving Ingestion and Pattern Recognition

The computational infrastructure can process multi-source interaction information, including professional medical portal navigation trends, scientific-content engagement, and appropriately governed Real-World Evidence (RWE) sources, where their use is lawful and relevant to the defined organizational purpose (Sabapathy, 2025).

Where appropriate to the specific use case, the architecture may incorporate privacy-preserving approaches such as federated learning, localized or edge processing, and differential privacy (Dwork & Roth, 2014; Yang et al., 2019). These approaches can reduce particular forms of centralized data exposure by limiting the amount or granularity of information transferred to centralized infrastructures or by introducing formal privacy protections into aggregate analysis.

Their use, however, should not be interpreted as independently establishing legal compliance. Technical privacy mechanisms complement rather than replace lawful-basis assessments, purpose limitation, transparency requirements, security safeguards, and other applicable data-protection obligations. Within the European context, their implementation should therefore remain consistent with principles such as data minimization under the General Data Protection Regulation (Regulation (EU) 2016/679).

Under the Data + Empathy framework, the purpose of this layer is not to construct increasingly granular individual surveillance profiles simply because such profiling is technically possible. Instead, data processing should remain proportionate to the defined informational or organizational objective. This reflects the broader framework principle that computational possibility must be evaluated against human-centered appropriateness before technological capability is translated into stakeholder intervention.

2. Grounded Information Retrieval through RAG

In regulated healthcare communication, unconstrained generative systems introduce important accuracy and governance challenges because Large Language Models can produce fluent outputs that are nevertheless unsupported or factually inaccurate (Ji et al., 2023). Within the proposed architecture, Retrieval-Augmented Generation (RAG) therefore functions as a grounding mechanism through which generation can be conditioned on controlled organizational evidence repositories rather than relying exclusively on information encoded within model parameters (Gao et al., 2023; Lewis et al., 2020).

Dense retrieval architectures can map incoming informational queries into vector representations and retrieve semantically relevant passages from controlled repositories (Karpukhin et al., 2020). In a pharmaceutical organizational environment, such repositories could include peer-reviewed publications, approved product information, Summaries of Product Characteristics (SmPC), Core Data Sheets (CCDS), regulatory documents, or other internally validated evidence sources according to the specific use case.

Retrieved material can subsequently be ranked and provided as contextual evidence for downstream generation or human decision support. This architecture creates stronger traceability between generated material and the underlying evidence base than an unconstrained open-domain generation process.

Nevertheless, retrieval grounding should be treated as a **risk-mitigation mechanism rather than a guarantee of empirical truth or regulatory compliance**. Retrieval systems may select incomplete, outdated, contextually inappropriate, or insufficiently relevant evidence, while generative models may still misinterpret retrieved passages (Gao et al., 2023). Consequently, evidence grounding must remain accompanied by source validation, repository governance, provenance mechanisms, and risk-proportionate human review.

3. Symbolic Guardrails through Knowledge Graphs

A further proposed component involves coupling retrieval architectures with structured enterprise Knowledge Graphs capable of representing relationships among scientific evidence, product information, safety content, regulatory metadata, and approved communication modules (Noy et al., 2019).

Within the proposed Data + Empathy architecture, predefined Knowledge Graph rules could associate benefit-related content with corresponding safety information, contraindications, evidence references, and relevant regulatory metadata before downstream review. The objective is not to allow the Knowledge Graph to determine legal compliance autonomously, but to create structural dependencies that make omission of related information easier to identify.

For example, when a content module concerning therapeutic benefit is retrieved, predefined relational rules could require the system to retrieve associated safety or risk information before the asset progresses through the communication workflow. Such structural constraints are intended to support more consistent association between benefit-related and safety-related information and can contribute to compliance-oriented content design under applicable pharmaceutical advertising requirements (21 C.F.R. § 202.1; Directive 2001/83/EC).

The Knowledge Graph therefore functions as a **computational guardrail**, not as an autonomous regulatory authority. Final assessment of whether a communication satisfies applicable statutory standards remains dependent on the complete context, jurisdiction, communication format, approved evidence base, and qualified human review.

4. Proposed Compliance-by-Design Pre-Screening

The final computational component proposed within this layer is an automated pre-screening mechanism operating before formal human Medical, Legal, and Regulatory (MLR) evaluation.

Drawing conceptually on the LLM-as-a-Judge paradigm examined by Zheng et al. (2023), this thesis proposes that an independent evaluator model could be adapted to pharmaceutical pre-screening tasks. Such a system would not constitute an autonomous regulatory decision-maker. Instead, it would operate as a decision-support layer calibrated against institutional review criteria, approved terminology, evidence requirements, and predefined communication-risk parameters.

Within such a proposed configuration, the evaluator could flag structural or semantic indicators including missing evidence references, potentially absolute therapeutic language, inconsistencies between claims and retrieved evidence, or inadequate association between benefit and safety information. The purpose would be to prioritize potentially problematic material for subsequent human review rather than to certify the communication as compliant.

Drafts failing predefined screening parameters could be automatically routed for revision and subsequent human evaluation. Importantly, these functions represent a **proposed application of automated evaluation architectures** rather than an empirically validated substitute for pharmaceutical MLR review. Their effectiveness would require domain-specific validation, monitoring, governance, and comparison against qualified human assessment before organizational deployment.

The Human Governance and Stewardship Layer (Frontend Functions)

Complementing computational processing, the human governance layer provides contextual interpretation, relational communication, and definitive organizational oversight. Under the proposed model, the degree of human involvement increases according to the sensitivity, uncertainty, and potential consequences of the communication.

1. Contextual and Qualitative Gatekeeping

Automated systems can support screening against predefined criteria, but many communication decisions require contextual interpretation that cannot be reduced reliably to structural or lexical rules. Medical affairs professionals, compliance specialists, and other qualified reviewers therefore assess qualitative tone, cultural context, stakeholder vulnerability, scientific uncertainty, and ethical nuance (OECD, 2023).

Human evaluation becomes particularly important where technically accurate information may nevertheless be inappropriate in tone, timing, framing, or context. The purpose of the human layer is therefore not merely to identify errors missed by computational systems, but to evaluate dimensions of communicative appropriateness that extend beyond computational accuracy.

2. Relational Dialogue in High-Sensitivity Contexts

As communication becomes more clinically consequential or emotionally sensitive, the framework progressively increases human communicative stewardship. Therapeutic areas involving severe chronic illness, oncology, neurodegenerative conditions, rare diseases, or significant psychological distress may involve questions and experiences whose interpretation requires substantial contextual sensitivity.

In these environments, qualified professionals can provide active listening, contextual clarification, and relational continuity while using computational systems as background decision-support tools where appropriate (Jones et al., 2024; Verma et al., 2024). This does not imply that every automated interface is inappropriate for healthcare communication. Rather, the framework proposes that the degree of autonomous automation should decrease as the potential consequences of misunderstanding, emotional insensitivity, or inappropriate clinical interpretation increase.

3. Data Provenance and Communication Transparency

The human governance layer also contributes to transparent communication of evidence provenance. Information delivered through digital systems should remain traceable, where appropriate, to the scientific or regulatory sources from which it originates.

Providing accessible references to clinical trials, approved product information, regulatory documents, and institutional evidence sources can support informed interpretation and allow stakeholders to distinguish organizational communication from the evidence underlying it (Sabapathy, 2025; Verma et al., 2024).

Transparency also concerns the technological process itself. Where the use of Artificial Intelligence materially affects the stakeholder interaction, organizations should evaluate whether disclosure is necessary or appropriate according to applicable regulation, organizational policy, ethical considerations, and stakeholder expectations. Data provenance within the Data + Empathy framework therefore encompasses both **scientific traceability and technological transparency**.

4. Statutory, Professional, and Organizational Accountability

Responsibility for pharmaceutical communication ultimately remains with the organization and the human professionals responsible for designing, approving, and deploying its communication systems. Regulatory accountability cannot be transferred to an algorithm simply because an automated system participated in content retrieval, generation, evaluation, or delivery.

In the United States, pharmaceutical communication is governed through applicable federal requirements and FDA oversight, while in the European Union Directive 2001/83/EC establishes a common legal framework implemented and enforced through relevant Member State structures. Human gatekeepers and corporate governance processes therefore retain final authority over consequential external dissemination according to the jurisdiction and communication context concerned (Directive 2001/83/EC; U.S. Food and Drug Administration, 2023).

Within the proposed framework, accountability represents the final boundary of automation. Computational systems may support analysis, retrieval, screening, and decision-making, but responsibility for defining appropriate use, establishing controls, approving consequential communication, and responding to failures remains organizationally human.

5.2.2 Cognitive and Behavioral Considerations: Addressing Aversion, Reactance, and Authenticity

The rationale for establishing a risk-proportionate oversight gradient can be further examined through three cognitive and behavioral frameworks: **Algorithm Aversion** (Dietvorst et al., 2015), **Psychological Reactance Theory** (Brehm, 1966), and the **Uncanny Valley of Mind** (Stein & Ohler, 2017). These theories do not establish that front-facing Artificial Intelligence necessarily produces negative stakeholder responses. Rather, they identify psychological mechanisms that can become relevant when automated systems assume visible, evaluative, persuasive, or apparently emotional roles.

Their incorporation into the Data + Empathy framework is therefore interpretative and risk-oriented. The purpose is not to predict universal stakeholder reactions, but to identify conditions under which increasing automation may require stronger transparency, human oversight, or contextual safeguards.

Algorithm Aversion in Consequential Health Contexts

Research on Algorithm Aversion demonstrates that individuals may become reluctant to rely on algorithmic forecasters after observing them make mistakes, even when those algorithms subsequently outperform human forecasters (Dietvorst et al., 2015). The significance of this finding lies not in establishing a universal preference for human judgment, but in demonstrating that visible algorithmic error can affect subsequent willingness to rely on automated recommendations.

The implications for healthcare communication should therefore be drawn cautiously. Dietvorst et al. (2015) did not investigate pharmaceutical conversational systems or establish that patients categorically reject algorithmic support in medical contexts. Nevertheless, the study provides a relevant theoretical basis for considering how observable errors may influence stakeholder reliance when automated systems participate in consequential communication.

This consideration becomes particularly important because generative systems can produce linguistically convincing responses even when the underlying information is inaccurate or insufficiently grounded (Ji et al., 2023). In a low-risk administrative interaction, such an error may be relatively easy to identify and correct. In a communication concerning treatment information, contraindications, adverse effects, or emotionally consequential health circumstances, the implications may be substantially greater.

The Data + Empathy framework addresses this risk not by excluding algorithmic participation, but by calibrating its authority. Computational models can support healthcare professionals and medical communicators through retrieval, evidence synthesis, and decision support, while human oversight increases as the consequences of error become more significant (Jones et al., 2024; Sabapathy, 2025; Verma et al., 2024).

Algorithm Aversion therefore contributes to the framework as a **governance consideration rather than a deterministic behavioral outcome**. It reinforces the importance of determining how much

authority an automated system should receive and how errors, uncertainty, and limitations should be communicated to stakeholders.

The Uncanny Valley of Mind as an Interpretative Analogy

The concept of the *Uncanny Valley of Mind* provides a second theoretical lens for considering automated emotional communication. Stein and Ohler (2017) examined responses to artificial entities perceived as possessing human-like mental and emotional characteristics. Their study was not conducted in pharmaceutical marketing or with contemporary Large Language Model interfaces. Consequently, direct causal conclusions concerning AI-mediated healthcare communication cannot be inferred from their findings.

Nevertheless, the concept provides a useful interpretative analogy for examining the increasing capacity of conversational systems to generate apparently empathetic language. In pharmaceutical communication, the framework can be used cautiously to hypothesize that overtly simulated emotional language may be experienced as incongruent when users remain aware that the apparent emotional response is computationally generated.

This distinction becomes particularly relevant when an artificial system communicates about grief, fear, chronic illness, treatment uncertainty, or other emotionally sensitive circumstances. Linguistic fluency does not independently establish contextual understanding, moral responsibility, or interpersonal empathy. A computational system may generate language that resembles compassionate communication without possessing the lived understanding or accountability associated with human relational interaction.

The Data + Empathy framework therefore distinguishes between **linguistic simulation of empathy and human-centered design**. It does not require automated systems to avoid all conversational interaction, nor does it assume that every synthetic expression of warmth will be negatively perceived. Instead, it proposes that organizations should exercise greater caution when automated communication attempts to simulate emotional understanding in contexts where stakeholder vulnerability is high.

Under this model, computational systems remain particularly valuable for functional, translational, retrieval, and analytical tasks, while progressively greater human involvement is introduced when communication requires psychosocial interpretation, emotional reassurance, or consequential relational judgment. This approach reflects the broader principle that the objective is not necessarily to make Artificial Intelligence appear more human, but to make its organizational deployment more human-centered.

Psychological Reactance and Digital Ad Avoidance

Psychological Reactance Theory provides a third theoretical foundation for the framework. Brehm (1966) conceptualized reactance as a motivational response that can emerge when individuals perceive their behavioral freedom or autonomy to be threatened.

In contemporary digital communication environments, highly personalized targeting can create tension between informational relevance and perceived intrusiveness. AI-driven systems can increase the capacity of organizations to infer preferences, predict potential needs, and adapt communication

accordingly (Sabapathy, 2025; Verma et al., 2024). However, when stakeholders perceive that personalization depends on opaque tracking or excessive use of personal information, the resulting communication may raise concerns regarding privacy, autonomy, and commercial surveillance.

Research on online advertising further indicates that negative attitudes toward advertising and awareness of ad-blocking technologies can be associated with avoidance behaviors (Redondo & Aznar, 2018). Redondo and Aznar's findings should not be interpreted as direct empirical evidence of psychological reactance in pharmaceutical communication; rather, they provide complementary evidence concerning the broader phenomenon of digital advertising avoidance.

Within this context, the Data + Empathy framework is designed to reduce conditions that may contribute to reactance by prioritizing transparent, proportionate, and utility-oriented communication over opaque behavioral targeting. Where personalization is used, the framework requires organizations to consider not only whether the message is algorithmically relevant but also whether the data practices and communication logic underlying that personalization are appropriate to the stakeholder context.

Privacy-preserving architectures, aggregate analysis where appropriate, explicit evidence attribution, and transparent communication regarding institutional involvement can therefore operate as complementary governance mechanisms. Their purpose is not to guarantee the elimination of reactance, but to reduce reliance on communication practices that stakeholders may perceive as unnecessarily intrusive or manipulative.

This distinction directly reinforces the personalization paradox discussed in Chapter 3. Greater predictive precision does not automatically generate greater stakeholder value. Personalization becomes strategically sustainable only when relevance is balanced against privacy expectations, autonomy, transparency, and the substantive utility of the interaction.

Lifecycle Phase	Proposed Computational Backend Functions	Human Governance & Frontend Functions	Regulatory & Compliance Guardrails	Target Behavioral & Ethical Hazard Addressed
1. Insight Harvesting & Needs Identification	Processes appropriately governed population-level interaction metrics and evidence sources; where appropriate, privacy-preserving approaches such as federated learning, localized processing, or differential privacy may be incorporated (Dwork & Roth, 2014; Yang et al., 2019).	Medical affairs and relevant cross-functional teams interpret informational needs, contextualize patterns, and determine whether identified signals justify communication intervention.	GDPR Art. 5(1)(c) on data minimization and Art. 9 where special-category health data are processed, together with other applicable jurisdictional requirements.	Risk of Invasive Profiling: seeks to limit unnecessary individual-level tracking and reduce reliance on intrusive behavioral surveillance.
2. Modular Content Compilation	RAG architectures query controlled enterprise repositories; generative systems may compile preliminary modular drafts supported by retrieved evidence and structured relational constraints (Gao et al., 2023; Lewis et al., 2020; Noy et al., 2019).	Medical writers, scientific experts, and creative professionals assess scientific accuracy, readability, contextual appropriateness, and alignment with stakeholder information needs.	Alignment with approved product information, applicable SmPC content, internal Core Data Sheets, substantiation requirements, and jurisdiction-specific communication rules.	Risk of Unsupported Generation: grounding conditions generated material on controlled evidence sources and facilitates identification of potential hallucinations or unsupported claims (Ji et al., 2023).
3. Regulatory Pre-Screening & MLR Review	Proposed evaluator models screen predefined structural and semantic risk indicators, flag potentially absolute claims, identify missing references or safety associations, and route high-risk drafts for human review (Zheng et al., 2023).	Multidisciplinary Medical, Legal, and Regulatory teams evaluate ambiguous claims, contextual meaning, scientific substantiation, and jurisdiction-specific requirements before formal approval.	Applicable advertising requirements, including 21 C.F.R. § 202.1 and national implementation of Directive 2001/83/EC; responsible-AI principles concerning human oversight (OECD, 2023).	Risk of Compliance Blind Spots: computational pre-screening may reduce repetitive review burdens while preserving qualified human authority over final compliance assessment.
4. Stakeholder Dialogue & Communication Delivery	Recommendation systems may support selection of relevant approved educational resources, channels, or communication options according to appropriately governed contextual information (Sabapathy, 2025).	Qualified medical professionals, medical information specialists, patient navigators, or trained communicators assume progressively greater stewardship as interactions become more sensitive or consequential.	Transparent institutional attribution, evidence provenance, appropriate separation between education and promotion, and applicable disclosure requirements.	Risk of Perceived Artificiality & Reactance: limits reliance on synthetic emotional simulation and supports more transparent human stewardship in sensitive interactions (Stein & Ohler, 2017).
5. Monitoring & Pharmacovigilance	NLP-based systems may support screening of company-managed digital channels for potential adverse-event mentions or product-quality complaints requiring further assessment.	Qualified pharmacovigilance personnel evaluate flagged information, assess reportability and available clinical information, conduct required follow-up, and manage applicable regulatory submissions.	Applicable postmarketing safety-reporting requirements, including 21 C.F.R. § 314.80 and relevant European Good Pharmacovigilance Practices, including GVP Module VI.	Risk of Safety Reporting Latency: computational screening can support earlier identification and escalation of potentially reportable safety information while preserving qualified human assessment.

5.2.3 Strategic and Operational Implementation Matrix

To translate the Data + Empathy framework into organizational practice, life sciences enterprises must define clear workflows, operational responsibilities, technological boundaries, and regulatory control points. The proposed operational model is structured across five core phases of the pharmaceutical communication lifecycle:

1. insight harvesting and needs identification;
2. modular content compilation;
3. regulatory pre-screening and MLR review;
4. stakeholder dialogue and communication delivery; and

5. monitoring and pharmacovigilance.

These phases should not be interpreted as a universally standardized pharmaceutical workflow. Rather, they provide an organizational model through which the principles of risk-proportionate automation, human-centered evaluation, evidence traceability, and human accountability can be translated into operational responsibilities.

Table 5.2 Operational Matrix of the Data + Empathy Governance Model

Analytical Commentary on Organizational Implementation

The operational matrix presented in Table 5.2 indicates that implementing the Data + Empathy framework requires more than the acquisition of Artificial Intelligence technologies. It requires adjustments to internal workflows, governance responsibilities, evidence infrastructures, and cross-functional collaboration within pharmaceutical organizations.

Digital transformation can increase the speed and granularity with which organizations generate, adapt, and distribute communication, creating pressure on organizational review structures originally developed for comparatively static communication environments (Jones et al., 2024; Sabapathy, 2025). A risk-proportionate governance model therefore requires technological acceleration to be accompanied by corresponding changes in organizational coordination.

Three implementation priorities emerge from the proposed matrix.

1. Restructuring the Medical, Legal, and Regulatory Workflow

Medical, Legal, and Regulatory review represents a critical organizational control point because pharmaceutical communication must reconcile commercial objectives, scientific substantiation, legal requirements, and patient safety. As content becomes increasingly modular, adaptive, and digitally distributed, exclusively manual review processes may face increasing scalability challenges (Jones et al., 2024; Sabapathy, 2025).

The Data + Empathy framework proposes automated pre-screening as a supporting layer before formal human review. Routine or partially structured tasks—such as identifying missing references, flagging potentially absolute therapeutic language, detecting inconsistencies with controlled terminology, or identifying missing associations between benefit and safety information—could be supported computationally.

This technological pre-screening should not be interpreted as automated regulatory approval. Rather, its purpose is to help prioritize human attention. Multidisciplinary MLR professionals retain responsibility for evaluating contextual meaning, scientific substantiation, ambiguous claims, jurisdictional differences, and ethical considerations that cannot be reduced reliably to predefined automated criteria (OECD, 2023; Sabapathy, 2025).

Under this configuration, automation does not displace human expertise. It reallocates repetitive computational tasks while preserving human authority over decisions requiring interpretation and accountability. The organizational value of Artificial Intelligence therefore derives not simply from reducing review time, but from enabling specialized human attention to be concentrated on the dimensions of communication where it is most consequential.

2. Operationalizing Data Provenance and Traceability

A second organizational requirement is the systematic management of **data provenance**. As generative and retrieval-based systems participate more extensively in communication workflows, organizations must preserve traceability between generated statements and the evidence from which those statements originate.

Under the Data + Empathy model, content compiled through backend retrieval architectures should retain structured metadata linking relevant statements to controlled source material, such as peer-reviewed clinical studies, approved product information, regulatory documentation, or institutional evidence repositories (Gao et al., 2023; Karpukhin et al., 2020).

Such traceability performs several functions. Internally, it enables reviewers to verify whether generated statements accurately reflect their source material. Externally, where appropriate, accessible references can allow stakeholders to distinguish organizational interpretation from primary evidence. From a governance perspective, provenance also supports auditing by creating a reconstructable relationship between input evidence, computational processing, human review, and final communication.

Data provenance therefore represents a connection point between the two dimensions of the framework. **Data** provides the evidentiary infrastructure, while **Empathy** requires that this evidence be communicated in a manner that respects the stakeholder's ability to understand, interrogate, and contextualize the information received.

3. Managing Organizational Culture and Change

The practical implementation of the framework ultimately depends on organizational culture. Life sciences leadership must navigate between two counterproductive tendencies: an uncritical expansion of front-facing generative AI without risk-proportionate oversight, which may introduce accuracy, trust, ethical, and compliance concerns; and an excessively cautious resistance to computational adoption, which may limit organizational agility and prevent beneficial technological capabilities from being used appropriately (Marchant et al., 2011; Sabapathy, 2025).

The Data + Empathy framework does not resolve this tension by selecting one extreme. Instead, it establishes a decision architecture through which technological adoption can be calibrated according to purpose, stakeholder context, evidentiary requirements, and risk.

This requires organizational teams to move beyond a binary question—*Should Artificial Intelligence be used?*—toward a more differentiated set of questions: Which function should be automated? What stakeholder problem does automation address? What evidence should constrain the system? What level of autonomy is proportionate to the risk involved? Where must human judgment remain decisive? How will the organization detect and respond to failures after deployment?

These questions transform AI governance from a defensive compliance activity into an organizational capability. Regulatory, medical, commercial, technological, and ethical considerations become integrated earlier in the communication lifecycle rather than addressed independently after technological development has already occurred.

Within this model, Artificial Intelligence can be understood as an enabling infrastructure for data processing, evidence retrieval, pattern recognition, compliance support, and operational coordination. Human professionals retain progressively greater responsibility for contextual judgment, emotional sensitivity, scientific interpretation, and definitive accountability as communication becomes more consequential (Jones et al., 2024; OECD, 2023; Verma et al., 2024).

The resulting architecture therefore avoids both technological maximalism and categorical technological rejection. It recognizes that some stakeholder interactions may appropriately accommodate substantial automation, while others require stronger forms of human stewardship. The governing principle is **proportionality**: the distribution of computational and human responsibilities should correspond to the sensitivity, uncertainty, regulatory exposure, and potential human consequences of the communication task.

In this sense, the operational contribution of the Data + Empathy framework can be summarized through a final sequence:

Stakeholder Need → Data-Informed Understanding → Computational Enablement → Human-Centered Appropriateness Assessment → Risk-Proportionate Governance → Human-Governed Deployment → Continuous Evaluation

This sequence consolidates the findings developed throughout Sections 5.1 and 5.2. Artificial Intelligence is not positioned as an autonomous substitute for pharmaceutical communication professionals, nor is human involvement treated as a universal requirement at every technological touchpoint. Instead, the framework assigns each component a function proportionate to its capabilities and limitations.

Data creates computational possibility. Empathy evaluates that possibility in relation to human context. Governance determines whether, where, and under what conditions that possibility should become organizational action.

Through this integration, the Data + Empathy framework is proposed as a structured approach for reconciling technological scalability with stakeholder utility, scientific integrity, human-centered communication, and accountable governance in AI-driven pharmaceutical communication. Its value at this stage remains conceptual and applied rather than empirically validated; the framework therefore provides a basis for organizational implementation and future empirical testing rather than a claim of demonstrated superiority over alternative governance models.

5.3 Field Research: Generation Z Perceptions of AI-Mediated Healthcare Communication

The conceptual framework developed throughout this thesis has examined how Artificial Intelligence may reshape pharmaceutical and healthcare communication through increasingly sophisticated forms of personalization, automation, and computational decision support. The comparative case analysis presented in Section 5.1 and the operationalization of the Data + Empathy framework developed in Section 5.2 further suggest that the strategic value of AI cannot be assessed exclusively in terms of technological capability. Its application must also be considered in relation to stakeholder expectations, perceived utility, transparency, privacy, human oversight, and the sensitivity of the communicative context.

To complement these conceptual and case-based perspectives with direct stakeholder evidence, an exploratory field study was conducted among adult members of Generation Z. The study investigates how younger individuals perceive the use of Artificial Intelligence in healthcare and pharmaceutical communication, with particular attention to the conditions under which AI-mediated interactions are considered acceptable, trustworthy, useful, or potentially intrusive. Rather than assuming that younger generations are inherently receptive to automated communication because of their familiarity with digital technologies, the research examines whether frequent exposure to AI translates into willingness to delegate health-related communicative functions to computational systems.

This distinction is particularly relevant within the context of pharmaceutical digital communication. The increasing capacity to collect and process behavioral data, generate personalized content, automate interactions, and support professional decision-making creates opportunities for more responsive stakeholder engagement. At the same time, the theoretical analysis developed in the preceding chapters has highlighted potential tensions involving privacy, algorithmic transparency, evidence provenance, perceived authenticity, human responsibility, and the appropriate boundaries of automation. Understanding how these tensions are perceived by potential recipients of AI-mediated communication therefore represents an important complement to the organizational and technological perspective developed earlier in the thesis.

The field research focuses on five interconnected dimensions derived from the theoretical framework: technological familiarity and functional acceptance; human involvement and technological role allocation; trust, transparency, and scientific provenance; personalization and privacy; and perceptions of synthetic empathy, algorithmic error, and human responsibility. These dimensions were translated into concrete scenarios and attitudinal questions in order to investigate stakeholder perceptions without requiring respondents to possess technical knowledge of AI architectures or pharmaceutical regulation.

The purpose of the field study is exploratory rather than confirmatory. The collected responses are not intended to provide statistically representative evidence of Generation Z as a whole, nor do they constitute formal empirical validation of the Data + Empathy framework. Instead, they provide an additional empirical perspective through which the conceptual propositions developed throughout the thesis can be examined, challenged, and refined. Patterns that converge with the framework are therefore interpreted as evidence of consistency within the surveyed sample, while divergent or ambiguous responses remain analytically relevant rather than being treated as exceptions.

In this sense, the field research introduces a stakeholder-level perspective into the broader analytical progression of Chapter 5. Whereas the case studies examined how organizations have allocated technological capabilities in real-world communication initiatives, and Section 5.2 translated these observations into an operational governance architecture, the survey investigates how potential recipients perceive comparable questions of technological delegation, transparency, personalization, and human involvement. The combination of these perspectives enables the Data + Empathy framework to be examined from organizational, technological, and stakeholder viewpoints without conflating conceptual interpretation with causal validation.

5.3.1 Research Design and Data Collection

The field research was conducted through a structured online questionnaire designed specifically for this study and administered through Google Forms. The questionnaire was selected as the primary data-collection instrument because it enabled the standardized collection of attitudes toward AI-mediated healthcare communication while maintaining a short completion time and allowing respondents to elaborate on selected aspects of their perceptions through an optional open-ended question.

The questionnaire was administered in Italian to facilitate comprehension among the target population and reduce the possibility that linguistic complexity could influence responses to concepts such as trust, privacy, personalization, empathy, and human responsibility. Although the thesis is written in English, the use of the respondents' primary or familiar language was considered more appropriate for the collection of attitudinal data. Technical concepts central to the theoretical framework were consequently translated into accessible scenarios rather than presented through specialist terminology. Respondents were not, for example, asked directly about concepts such as Algorithm Aversion, functional asymmetry, or the Personalization Paradox; instead, these theoretical dimensions were operationalized through questions concerning reactions to AI errors, preferred human–AI configurations, and behavioral-data-driven personalization.

Participation was voluntary and anonymous. Before completing the questionnaire, participants were informed that the research formed part of a Master's thesis concerning Artificial Intelligence in healthcare and pharmaceutical communication and that the responses would be analyzed exclusively for academic purposes. The questionnaire did not request names, email addresses, diagnoses, medication histories, or other directly identifying or health-related personal information.

Data collection produced **152 submitted questionnaires**. The primary analytical population was defined as adult respondents aged between **18 and 29 years**. Five respondents reported being aged 30 years or older and were therefore excluded from the primary analysis in accordance with the predefined age criterion. The resulting analytical sample consisted of **147 respondents (N = 147)**.

Within the final sample, **11 respondents (7.5%) were aged 18–20, 63 (42.9%) were aged 21–23, 64 (43.5%) were aged 24–26, and 9 (6.1%) were aged 27–29**. The sample was therefore concentrated primarily in the 21–26 age range, which accounted for 86.4% of eligible participants. This distribution should be considered when interpreting the findings, as the results reflect the perceptions of the surveyed population rather than an evenly distributed representation of all adult age groups within Generation Z.

Participants were recruited digitally through personal and online networks, resulting in a non-probability convenience sampling strategy with elements of snowball dissemination, as recipients were also able to share the questionnaire within their own networks. This approach enabled rapid access to the target population but does not provide each member of Generation Z with an equal probability of selection. The resulting sample should therefore not be interpreted as statistically representative of the broader Italian or international Generation Z population.

The sampling limitation is particularly relevant because individuals who chose to participate may differ systematically from those who did not. Moreover, recruitment through digitally connected networks may contribute to comparatively high levels of familiarity with Artificial Intelligence within

the sample. Rather than attempting to eliminate these characteristics retrospectively, the analysis treats them transparently as part of the empirical context in which the findings were generated.

The final dataset was cleaned according to the predefined age criterion before analysis. Quantitative responses were subsequently examined through descriptive statistics and selected exploratory associations, while the optional open-ended responses were analyzed separately through thematic coding. This combination enabled the study to preserve the comparability offered by standardized survey questions while incorporating a limited qualitative component capable of capturing respondents' own definitions of appropriate human responsibility.

5.3.2 Survey Instrument and Analytical Strategy

The questionnaire was intentionally designed as a **short-form exploratory instrument**, with an estimated completion time of approximately two to three minutes. The decision to limit respondent burden reflected a practical methodological consideration: a concise questionnaire was expected to facilitate completion and dissemination while preserving sufficient coverage of the constructs most directly relevant to the research question.

The instrument consisted predominantly of closed-ended questions using categorical response options and five-point rating scales, supplemented by an optional open-ended question. Rather than attempting to measure every theoretical construct discussed throughout the thesis, the survey concentrated on dimensions that could be translated meaningfully into stakeholder perceptions and concrete healthcare communication scenarios.

The first analytical dimension examined **AI familiarity and baseline acceptance**. Respondents indicated how frequently they used AI-based tools and how comfortable they would feel receiving health-related information from an AI system. These variables were analyzed separately and subsequently compared in order to investigate whether greater everyday exposure to Artificial Intelligence was associated with greater comfort in a healthcare communication context. Because the variables are ordinal, their exploratory association was examined using **Spearman's rank-order correlation** rather than assuming interval-level measurement.

The second dimension addressed **functional role allocation and human oversight**. Respondents evaluated the acceptability of AI for healthcare-related functions, indicated their preferred form of interaction in the case of a serious or emotionally sensitive health issue, and identified what they considered the most appropriate overall role for AI in healthcare communication. These questions were designed to determine whether attitudes toward AI vary according to the function assigned to the technology and the sensitivity of the interaction. The analytical emphasis therefore concerns technological role allocation rather than a binary distinction between acceptance and rejection of Artificial Intelligence.

The third dimension examined **trust, transparency, and scientific provenance**. Participants evaluated their trust in AI-mediated health information when accompanied by verifiable scientific sources and indicated the importance they attributed to knowing whether healthcare content had been generated or modified through AI. These questions provide an empirical perspective on the evidence-grounding and transparency principles developed in Chapters 3 and 4 and operationalized within the Data + Empathy architecture in Section 5.2.

The fourth dimension investigated the **relationship between personalization and privacy**. Participants were asked how comfortable they would feel if a pharmaceutical or healthcare organization used information concerning their online behavior to personalize health-related content. They were subsequently asked whether an explicit explanation of which data were being used and for what purpose would alter their level of comfort. These questions were designed to explore the tension between personalization and perceived intrusiveness discussed in relation to the Personalization Paradox.

Importantly, the questionnaire does not constitute an experimental manipulation of transparency. Respondents were asked to report how their attitude *would* change under a hypothetical condition rather than being randomly exposed to different communication treatments. Consequently, the findings cannot establish that transparency causally increases acceptance of behavioral personalization. They instead provide evidence concerning respondents' **self-reported expectations regarding the role of transparency**.

The fifth dimension examined **synthetic empathy and sensitivity to algorithmic error**. Respondents evaluated a hypothetical empathetic statement delivered by a healthcare chatbot and indicated how an error in AI-generated health information would affect their subsequent trust. The first scenario was intended to investigate whether human-like emotional language is necessarily experienced as reassuring when generated by an artificial system, while the second provides an exploratory stakeholder perspective on the fragility of trust following observable algorithmic error.

These questions can subsequently be interpreted in relation to the theoretical literature on Algorithm Aversion and the Uncanny Valley of Mind. However, the questionnaire was not designed as an experimental replication of Dietvorst et al. (2015) or Stein and Ohler (2017). Any correspondence between the observed responses and these theoretical models is therefore treated as interpretative consistency rather than direct empirical confirmation.

Finally, participants were invited to respond in their own words to the question of **which aspects of healthcare communication should remain under human responsibility when Artificial Intelligence is used**. This optional item generated a qualitative dataset that complements the standardized responses. The answers were examined through iterative thematic coding, allowing recurring concepts to be grouped into broader categories while preserving the possibility of identifying perspectives not anticipated by the theoretical framework.

The quantitative analysis primarily employs **frequencies, percentages, medians, and means for descriptive purposes**, alongside selected cross-variable comparisons. Where theoretically relevant, ordinal associations are explored through Spearman's rank-order correlation. Given the convenience-sampling strategy and exploratory research design, statistical significance is not interpreted as evidence of population-level generalizability. Its role is instead to identify patterns within the observed sample that warrant theoretical interpretation.

The qualitative analysis follows a complementary logic. Individual responses may contain multiple concepts and are therefore not necessarily assigned to mutually exclusive categories. Thematic coding focuses on recurring references to areas such as **clinical judgment and final decision-making, relational and emotional responsibility, professional oversight and verification, accountability, privacy, transparency, and patient autonomy**. Particular attention is paid to both convergent and

divergent responses in order to avoid coding the material exclusively around propositions already contained within the Data + Empathy framework.

Taken together, the analytical strategy combines **descriptive quantitative evidence, exploratory relational analysis, and a limited qualitative thematic component**. The purpose is not to determine whether Generation Z simply accepts or rejects Artificial Intelligence in healthcare. Instead, the analysis investigates a more specific set of questions: *which technological functions respondents appear willing to delegate, which conditions influence their stated trust and comfort, how they perceive data-driven personalization and synthetic empathy, and where they believe human judgment and responsibility should remain central*.

5.3.3 Quantitative Findings: Technological Role Allocation, Trust, and Privacy Dynamics

The quantitative analysis was conducted on the final analytical sample of **147 respondents aged between 18 and 29 years**. The findings are organized according to the five analytical dimensions defined in Section 5.3.2: technological familiarity and baseline acceptance; functional role allocation and human oversight; trust, transparency, and scientific provenance; personalization and privacy; and perceptions of synthetic empathy and algorithmic error. Rather than treating individual questionnaire items as isolated indicators, the analysis examines their distributions and selected relationships in order to identify broader patterns in respondents' attitudes toward AI-mediated healthcare communication. Overall, the results reveal a pattern of **conditional rather than unconditional acceptance**, in which high familiarity with Artificial Intelligence coexists with differentiated expectations concerning technological autonomy, evidence transparency, behavioral data use, and human involvement.

Table 5.3. Descriptive Statistics of Survey Rating Scales (N = 147)

Survey Construct / Item Description	Mean	SD	Median	Low Ratings (1–2)	Neutral (3)	High Ratings (4–5)
Comfort Receiving AI Health Information	3.03	0.92	3.0	45 (30.6%)	57 (38.8%)	45 (30.6%)
Trust with Verifiable Scientific Sources	3.59	0.87	4.0	13 (8.8%)	59 (40.1%)	75 (51.0%)
Importance of AI Content Labeling	4.37	1.03	5.0	11 (7.5%)	15 (10.2%)	121 (82.3%)
Comfort with Behavioral Personalization	2.19	1.08	2.0	89 (60.5%)	42 (28.6%)	16 (10.9%)

Note. All items were measured on five-point rating scales. Scale anchors differed according to the specific questionnaire item. Low ratings represent scores 1 and 2; high ratings represent scores 4 and 5.

1. Technological Familiarity and Baseline Acceptance

The surveyed sample displays a high level of everyday familiarity with Artificial Intelligence. A total of **37.4% (n = 55)** of respondents reported using AI-based tools every day, while an additional **39.5% (n = 58)** reported using them often. Combined, **76.9% (n = 113)** of participants therefore reported frequent or daily AI use. Occasional use accounted for 17.0% (n = 25), while rare or non-use represented the remaining 6.1% (n = 9).

This high level of technological familiarity, however, coexists with considerably more moderate attitudes toward receiving health-related information from an AI system. As shown in Table 5.3, responses were centered around the midpoint of the five-point scale (**M = 3.03, SD = 0.92, Mdn = 3.0**). Exactly **30.6% (n = 45)** selected the two lowest levels of comfort, **38.8% (n = 57)** selected the midpoint, and **30.6% (n = 45)** selected the two highest levels. The symmetrical distribution suggests that familiarity with AI in everyday contexts should not automatically be interpreted as equivalent to acceptance of its involvement in health communication.

An exploratory Spearman rank-order correlation identified a moderate positive association between frequency of AI use and stated comfort with receiving AI-mediated health information (**r_s = .347, p < .001**). Within the surveyed sample, respondents who reported more frequent AI use therefore tended to express greater comfort with its application to health information. Nevertheless, the moderate magnitude of the association indicates that technological familiarity accounts for only part of the variation in respondents' attitudes. Regular interaction with AI does not eliminate reservations when the technology is transferred into a more sensitive informational domain.

The finding introduces an important distinction for the subsequent analysis. Rather than displaying a uniform attitude toward Artificial Intelligence, respondents appear to evaluate its acceptability in relation to the specific function and context in which it is deployed. This becomes particularly evident when AI moves from general informational assistance toward activities involving greater clinical, emotional, or decisional sensitivity.

2. Functional Role Allocation and Interaction Boundaries

The responses concerning specific healthcare-related activities reveal substantial variation in the functions for which participants considered AI acceptable. **General health information retrieval** was selected by 26.5% (n = 39), followed by **administrative bookings and scheduling** at 24.5% (n = 36), **clarification of medical terminology** at 19.0% (n = 28), and **support to healthcare professionals** at 14.3% (n = 21). Lower levels of acceptance were recorded for information concerning medicinal products (6.8%, n = 10), highly personalized health recommendations (2.7%, n = 4), emotional or psychological support (2.0%, n = 3), and autonomous diagnosis or therapeutic choices (1.4%, n = 2).

The distribution indicates that respondents were comparatively more receptive to AI when it was associated with administrative, informational, explanatory, or professional-support functions, while substantially fewer selected activities involving individualized recommendations, emotional support, diagnosis, or therapeutic decision-making. The result therefore suggests that attitudes toward AI cannot be adequately represented through a single measure of technological acceptance: the **function assigned to the system appears to matter considerably**.

This differentiation becomes particularly pronounced in a higher-sensitivity scenario. When respondents were asked whom they would prefer to interact with in the case of a serious or emotionally sensitive health issue, **82.3% (n = 121)** selected a healthcare professional, while **10.9% (n = 16)** preferred a healthcare professional supported by AI. A further **4.8% (n = 7)** selected an AI system supervised by a healthcare professional, and 2.0% (n = 3) expressed no preference. **No respondent selected a fully autonomous AI system.**

The distribution indicates a strong preference for human involvement as the sensitivity of the interaction increases. Importantly, this should not be interpreted as evidence of a generalized rejection of AI. The second- and third-most selected configurations explicitly incorporate Artificial Intelligence, but retain professional participation or supervision. The distinction therefore appears to concern the **degree of technological autonomy and communicative authority**, rather than the mere presence or absence of AI.

A similar pattern emerges when respondents were asked to identify the most appropriate overall role for Artificial Intelligence in healthcare communication. **60.5% (n = 89)** indicated that AI should operate primarily *behind the scenes in support of healthcare professionals*, while **19.0% (n = 28)** considered direct interaction appropriate only for simple activities. A further **17.0% (n = 25)** preferred a very limited role for AI. By contrast, only **1.4% (n = 2)** considered AI and human communicators to have equivalent roles, **0.7% (n = 1)** supported direct AI interaction in complex situations, and **1.4% (n = 2)** remained undecided.

Taken together, **79.6% of respondents selected either a primarily supportive backend role or direct interaction limited to simple activities.** Within the surveyed sample, this distribution is consistent with the principle of **risk-proportionate functional asymmetry** developed in Section 5.2: increasing technological capability does not necessarily correspond with willingness to delegate equivalent levels of communicative authority. The survey does not validate this conceptual model, but the observed preferences provide stakeholder-level evidence consistent with the proposition that the appropriate degree of automation may vary according to functional complexity and contextual sensitivity.

3. Trust, Transparency, and Scientific Evidence Grounding

A third set of findings concerns the conditions under which respondents evaluate AI-mediated health information as trustworthy. When participants were asked how much they would trust health information generated by AI **if accompanied by verifiable scientific sources**, responses were predominantly neutral to positive (**M = 3.59, SD = 0.87, Mdn = 4.0**). More than half of the sample, **51.0% (n = 75)**, selected one of the two highest levels of trust, while 40.1% (n = 59) selected the midpoint. Only 8.8% (n = 13) selected a low trust rating, and no respondent selected the minimum value.

These results cannot be interpreted as evidence that the inclusion of scientific sources *causes* an increase in trust. The questionnaire did not experimentally compare identical AI-generated messages presented with and without scientific references, and the general comfort item reported previously measures a different construct. Nevertheless, the response distribution indicates that **verifiable scientific provenance is positively regarded by a substantial proportion of the sample** and represents a relevant condition in respondents' evaluation of AI-mediated health information.

An even clearer pattern emerges regarding disclosure of AI involvement. When participants were asked how important it was to know whether healthcare content had been created or modified using Artificial Intelligence, the mean response was **4.37 out of 5 (SD = 1.03, Mdn = 5.0)**. A total of **64.6% (n = 95)** selected the maximum rating of 5, while an additional 17.7% (n = 26) selected 4. Consequently, **82.3% (n = 121)** assigned high importance to knowing whether AI had been involved in the creation or modification of healthcare content. Only 7.5% (n = 11) selected one of the two lowest ratings.

Within this sample, transparency concerning technological provenance therefore emerges as a particularly salient stakeholder expectation. The result does not establish that disclosure automatically produces trust, nor does it determine the specific form that effective disclosure should take. It does, however, indicate that respondents generally consider knowledge of AI involvement relevant when evaluating health-related communication.

Taken together, the findings concerning scientific provenance and AI disclosure suggest that respondents' evaluation of automated health communication extends beyond the informational output itself. **How the information was produced, whether its evidentiary basis can be verified, and whether technological involvement is made visible appear to constitute relevant dimensions of perceived credibility and transparency.** These results provide an empirical basis for subsequently discussing the provenance and traceability principles incorporated into the Data + Empathy framework, without implying that the survey directly tested any specific technical architecture such as Retrieval-Augmented Generation or Knowledge Graphs.

4. Personalization, Privacy, and Procedural Transparency

The strongest resistance observed in the quantitative data concerns the use of behavioral information for personalized health communication. When respondents were asked how comfortable they would feel if a pharmaceutical or healthcare organization used information concerning their online behavior to personalize health-related content, responses were concentrated toward the lower end of the scale (**M = 2.19, SD = 1.08, Mdn = 2.0**).

A total of **34.0% (n = 50)** selected the minimum level of comfort and **26.5% (n = 39)** selected a rating of 2. Consequently, **60.5% (n = 89)** expressed low comfort with behavioral personalization. A further 28.6% (n = 42) selected the midpoint, while only **10.9% (n = 16)** selected one of the two highest levels of comfort.

This distribution is particularly noteworthy when considered alongside the comparatively high frequency of AI use observed within the sample. Respondents who are broadly familiar with Artificial Intelligence do not necessarily appear equally receptive to the collection or use of behavioral information for personalized health communication. The result therefore reinforces the importance of distinguishing **technological familiarity from data-use acceptance**.

Participants were subsequently asked whether their perception would change if the organization clearly explained **which data were being used and for what purpose**. In response, **55.1% (n = 81)** indicated that they would feel slightly more comfortable and **12.2% (n = 18)** that they would feel much more comfortable. Taken together, **67.3% (n = 99)** reported that greater transparency would improve their stated level of comfort to some degree. Conversely, 23.1% (n = 34) indicated that

disclosure would make no difference, 3.4% (n = 5) would feel less comfortable, and 6.1% (n = 9) were uncertain.

These responses should not be interpreted as demonstrating that transparency causally resolves privacy concerns. The item measures participants' self-reported reaction to a hypothetical disclosure rather than their observed behavior under experimental conditions. Moreover, almost one quarter of respondents indicated that greater transparency would make no difference, demonstrating that disclosure alone does not eliminate resistance to behavioral personalization.

Instead, the two items reveal a more nuanced **privacy–personalization tension**. Respondents display substantial discomfort with the use of online behavioral information for health-related personalization, while many simultaneously report that explicit explanations concerning data use and purpose could make the practice more acceptable. Within the context of the Personalization Paradox discussed earlier in the thesis, these findings suggest that the perceived value of personalization cannot be separated from the conditions under which personal information is collected, processed, and communicated to stakeholders.

For pharmaceutical and healthcare communication, the implication is therefore not that personalization should be abandoned, nor that transparency automatically legitimizes behavioral tracking. Rather, the findings indicate that **technical capacity for increasingly granular personalization and stakeholder willingness to accept such personalization are distinct considerations**. Data-driven relevance must consequently be evaluated alongside transparency, perceived intrusiveness, purpose limitation, and stakeholder expectations concerning informational autonomy.

5. Synthetic Empathy and the Fragility of Algorithmic Trust

The final quantitative dimension examines respondents' perceptions of simulated empathy and their anticipated reaction to algorithmic error. Participants were presented with a hypothetical healthcare chatbot stating, *“I understand how difficult this situation might be for you”* and were asked how they would perceive the interaction.

Responses indicate considerable skepticism toward this form of simulated empathetic language. **45.6% (n = 67)** described the statement as either fairly or very artificial or uncomfortable, including 27.9% (n = 41) selecting the less intense negative response and 17.7% (n = 26) selecting the strongest negative response. A further **29.3% (n = 43)** perceived the statement neutrally, while 10.9% (n = 16) indicated that their reaction would depend on the context. Only **14.2% (n = 21)** considered the statement reassuring, including 12.2% (n = 18) who found it moderately reassuring and 2.0% (n = 3) who found it highly reassuring.

Within the surveyed sample, linguistic simulation of empathy therefore did not consistently translate into perceived reassurance. This finding should not be generalized to all conversational AI interactions, nor does it demonstrate that Artificial Intelligence is incapable of participating appropriately in emotionally sensitive communication. The response concerns a single hypothetical statement presented without a broader conversational context. Nevertheless, the distribution suggests that **human-like language and human-centered communication should not be treated as equivalent concepts**.

The pattern can be interpreted cautiously alongside the *Uncanny Valley of Mind* discussed by Stein and Ohler (2017), insofar as the attribution or simulation of human-like mental and emotional characteristics may produce ambivalent responses. However, the survey does not constitute a test of the Uncanny Valley of Mind, and the theoretical model was not developed specifically to explain AI-mediated healthcare communication. Its relevance here is therefore interpretative rather than confirmatory.

An even stronger response emerged when participants were asked how an observable error in AI-generated health information would affect their subsequent trust. **51.7% (n = 76)** reported that their trust would decrease substantially, **19.0% (n = 28)** that it would decrease slightly, and **7.5% (n = 11)** that they would stop using the system entirely. Consequently, **78.2% (n = 115)** anticipated some explicit reduction in trust or discontinuation of use following an error. A further 13.6% (n = 20) indicated that their reaction would depend on the severity of the error, while only 8.2% (n = 12) stated that their trust would remain unchanged.

The result indicates considerable sensitivity to errors in a health-information context. It can be interpreted as broadly consistent with the phenomenon of **Algorithm Aversion** described by Dietvorst et al. (2015), whereby observable algorithmic errors may reduce willingness to rely on automated systems. However, the present questionnaire did not replicate the experimental conditions of Dietvorst et al. and therefore cannot establish the presence or magnitude of Algorithm Aversion as a causal mechanism within the sample.

Respondents' stated concerns provide further context for this result. **Incorrect information was the most frequently selected primary concern (39.5%, n = 58)**, followed by **the absence of the human element (26.5%, n = 39)**. Legal responsibility in the event of an error accounted for 11.6% (n = 17), privacy concerns for 8.8% (n = 13), and algorithmic bias for 6.8% (n = 10), with the remaining responses distributed across other available categories.

Taken together, the findings concerning simulated empathy, algorithmic error, and perceived risks reinforce the broader pattern emerging throughout the quantitative analysis. Respondents do not appear to evaluate AI-mediated healthcare communication solely according to whether the technology can technically perform a task. Their attitudes vary according to **the sensitivity of the function, the degree of technological autonomy, the availability of human oversight, the transparency of data and content provenance, and the perceived consequences of technological failure**.

The quantitative findings therefore depict a surveyed population characterized by substantial familiarity with Artificial Intelligence but selective acceptance of its application to healthcare communication. Frequent AI use is positively associated with greater comfort, yet this openness coexists with strong preferences for human involvement in sensitive situations, high expectations regarding AI disclosure and scientific provenance, resistance to opaque behavioral personalization, and pronounced sensitivity to algorithmic error. Rather than supporting a binary interpretation of younger respondents as either enthusiastic adopters or technological skeptics, the results point toward a more differentiated pattern of **context-dependent and conditional acceptance**.

5.3.4 Qualitative Findings: Human Responsibility in AI-Mediated Healthcare Communication

To complement the standardized quantitative measures, the questionnaire included an optional open-ended question asking respondents: *“In your opinion, what should necessarily remain under human responsibility when AI is used in healthcare communication?”* A total of **68 eligible participants (n = 68)** provided a substantive written response. These answers, translated here from the original Italian, introduced a qualitative component into the field research, allowing respondents to define the boundaries of appropriate AI involvement in their own words rather than selecting exclusively from predetermined categories.

The responses were examined through an iterative thematic coding process. Because individual answers frequently referred to more than one dimension of human responsibility, the identified themes were treated as **non-mutually exclusive**. The purpose of the analysis was not to quantify qualitative responses as if they represented independent survey variables, but to identify recurrent patterns, points of convergence, and distinctive perspectives capable of contextualizing the quantitative findings. Four broad themes emerged from the material: **clinical judgment and therapeutic decision-making; relational and emotional responsibility; human oversight and final accountability; and informational integrity and commercial boundaries**.

1. Clinical Judgment and Therapeutic Decision-Making

The most recurrent theme concerned the preservation of **clinical judgment and consequential healthcare decisions under human responsibility**. Respondents repeatedly referred to diagnosis, treatment selection, prescription, prognosis, and the final evaluation of an individual patient's condition as activities for which professional involvement should remain central. The emphasis was therefore not necessarily on excluding AI from clinical processes, but on limiting the extent to which computational assistance should become autonomous decision-making.

Several responses articulated this distinction explicitly. One participant emphasized:

“The final choice of diagnosis and targeted therapies... the diagnosis above all, because only the doctor knows and directly sees the patient.” (Respondent 2)

Another respondent similarly distinguished between informational assistance and final clinical authority:

“The final evaluation and prescription of medications; AI should be strictly limited to informational support or explaining complex terms.” (Respondent 32)

A third participant highlighted the potential consequences of interpreting symptom-based AI responses without professional mediation:

“When a person describes symptoms to an AI, they might receive responses suggesting severe pathologies, generating unnecessary anxiety. Diagnosis should remain exclusively with the visiting clinician.” (Respondent 61)

These responses are consistent with the quantitative findings reported in Section 5.3.3, where **82.3% of respondents preferred direct interaction with a healthcare professional in the case of a serious or emotionally sensitive health issue**, while no participant selected a fully autonomous AI system. The qualitative material provides additional context for this distribution: participants frequently

differentiated between the capacity of AI to provide informational or analytical support and the authority required to make consequential judgments concerning an individual patient.

This distinction is particularly relevant to the concept of functional role allocation developed earlier in the thesis. Within these responses, AI is not uniformly rejected; rather, its acceptable role is frequently described as **supportive rather than substitutive**. Human responsibility becomes particularly salient when communication intersects with diagnosis, treatment, prescription, or other decisions perceived as having direct consequences for the individual.

2. Relational and Emotional Responsibility

A second prominent theme concerns the **relational dimension of healthcare**. Respondents repeatedly referred to empathy, listening, direct contact, psychological support, communication during difficult circumstances, and the doctor–patient relationship as areas in which human involvement remains particularly important.

One respondent summarized this expectation through an explicitly interpersonal formulation:

“Humanity must be received with specific human faces and genuine smiles!” (Respondent 45)

Another participant emphasized the importance of the relationship itself:

“The doctor-patient relationship, especially in difficult circumstances; AI frequently relies on pre-formulated templates that maintain an inherent communicative coldness.” (Respondent 30)

A more extensive response connected clinical knowledge with relational capabilities:

“Taking charge of a patient involves knowing their clinical and personal history, the diagnosis, and above all the treatment, combined with an ability to listen and provide empathy that only a human healthcare professional can offer.” (Respondent 17)

These statements should be interpreted as respondents' perceptions rather than as evidence that AI systems are inherently incapable of participating in emotionally sensitive interactions. Nevertheless, they provide important qualitative context for the quantitative finding that **45.6% of participants perceived the hypothetical empathetic chatbot statement as artificial or uncomfortable, compared with 14.2% who perceived it as reassuring**.

Taken together, the two forms of evidence suggest that, within the surveyed sample, **linguistic simulation of empathy was not consistently perceived as equivalent to interpersonal empathy**. Participants' open-ended responses frequently located empathetic value in listening, contextual understanding, personal presence, and relational continuity rather than in the linguistic form of the message alone.

This distinction reinforces an important conceptual separation developed throughout the thesis: **human-like communication and human-centered communication are not necessarily synonymous**. An AI system may generate language that resembles empathetic human expression, while stakeholders may evaluate the broader interaction according to additional characteristics such as who assumes responsibility, whether their individual circumstances are understood, and whether access to a human professional remains available.

3. Human Oversight, Verification, and Final Accountability

A third recurring theme concerns **oversight and final responsibility**. Several respondents explicitly described AI as a potentially useful support instrument while maintaining that a human professional should verify its outputs, retain decision-making authority, or assume final responsibility for consequential communication. The emphasis therefore shifts from whether AI should be used to **who remains accountable when it is used**.

One of the most concise responses captured this principle directly:

“The final word, the ultimate signature must always belong to a human brain.” (Respondent 3)

Another participant articulated a similar division of responsibility:

“AI can be an excellent support tool, but clinical decisions, professional liability—such as diagnosis, prognosis, and patient communication—must remain human responsibilities.” (Respondent 37)

A third respondent used a particularly illustrative analogy:

“I believe AI is a tool much like any other, but it must be managed with absolute awareness. Exactly like a scalpel: used correctly it saves lives, used improperly it causes severe damage.” (Respondent 69)

The importance of these statements lies in their rejection of a simple human-versus-machine dichotomy. The technology itself is not necessarily characterized as undesirable. Instead, respondents frequently frame the central issue as one of **governance, supervision, and responsibility for outcomes**.

The metaphor of the scalpel is particularly illustrative in this respect. Rather than attributing autonomous moral value to the technology, it places emphasis on the conditions under which a powerful instrument is deployed and on the actor responsible for its use. This perspective closely parallels the governance problem addressed in Sections 5.1 and 5.2: technological capability does not independently determine the appropriate degree of organizational delegation.

The qualitative material therefore complements the quantitative preference for AI operating primarily behind the scenes or within limited functional domains. Together, the findings suggest that respondents are often willing to envisage a role for computational support while simultaneously expecting **identifiable human authority to remain available for verification, escalation, and consequential decision-making**.

4. Informational Integrity and Commercial Boundaries

A smaller but analytically distinctive group of responses extends the discussion beyond clinical and relational responsibility toward the **integrity of healthcare information itself**. These participants raised concerns about whether AI-generated communication might prioritize conversational agreement, reassurance, or commercial objectives over informational accuracy and stakeholder interests.

One respondent observed:

“I think AI often tends to give responses designed to ‘please’ the user. This might be acceptable if you ask how you look with bangs, but not in health communication where answers must be entirely objective.” (Respondent 46)

Another participant summarized the issue from an organizational perspective:

“Transparency: putting health first, rather than the commercial product.” (Respondent 60)

These responses should not be interpreted as empirical evidence that the AI systems encountered by respondents systematically produce agreeable or commercially biased information. Nor does the questionnaire establish the technical mechanisms responsible for such perceptions. Their analytical significance instead lies in the **stakeholder concern they express**: health communication is expected to preserve informational integrity even where this conflicts with conversational convenience or commercial objectives.

This theme connects qualitative perceptions with several quantitative findings. Incorrect information represented the most frequently selected primary concern regarding AI-mediated healthcare communication, while respondents also attributed high importance to verifiable scientific sources and disclosure of AI involvement. Across these different questionnaire components, **informational reliability, provenance, and transparency repeatedly emerge as relevant conditions of technological acceptance**.

For pharmaceutical communication specifically, this observation is significant because AI-mediated interactions may occur within organizations that simultaneously pursue commercial, informational, and public-health objectives. The qualitative responses suggest that some participants expect clear boundaries between these functions and attach importance to the preservation of scientific integrity when AI participates in the communication process.

Cross-Thematic Interpretation

Considered collectively, the qualitative responses do not reveal a simple preference for human communication over technological communication. Instead, they describe a more differentiated allocation of responsibilities. Participants frequently envisage Artificial Intelligence as capable of supporting information retrieval, explanation, routine processes, or professional activity, while identifying **clinical judgment, emotionally sensitive interaction, verification, and final accountability** as areas in which human involvement should remain particularly prominent.

The relationship between the themes is also important. Clinical responsibility and relational responsibility frequently overlap: respondents may simultaneously expect a professional to make the final clinical decision and to communicate that decision in a manner sensitive to the patient's individual circumstances. Similarly, informational integrity is connected to accountability, because verification becomes meaningful only when an identifiable actor remains responsible for deciding whether AI-generated information is sufficiently reliable and appropriate for communication.

The open-ended responses therefore add explanatory depth to the quantitative pattern of **conditional AI acceptance** identified in Section 5.3.3. The preference for human involvement in sensitive contexts is not expressed exclusively as resistance to technology. Instead, participants repeatedly

articulate specific reasons for retaining human participation, including clinical contextualization, emotional support, verification of information, and responsibility for consequential outcomes.

Importantly, this convergence should not be interpreted as formal validation of the Data + Empathy framework. The open-ended responses were not designed to reproduce the framework's terminology, and thematic coding necessarily involves an interpretative component. Nevertheless, the independently expressed distinctions between **computational assistance and human responsibility** provide stakeholder-level evidence that is consistent with several of the conceptual boundaries developed earlier in the thesis.

The qualitative analysis therefore contributes a dimension that cannot be captured fully through percentages or rating scales. While the quantitative results identify **where** respondents tend to draw boundaries around AI involvement, the open-ended responses provide insight into **why** those boundaries may matter to them. Across the responses, the recurring concern is not simply whether Artificial Intelligence is present in healthcare communication, but whether its role preserves clinical responsibility, interpersonal sensitivity, informational integrity, and identifiable human accountability.

5.3.5 Integrated Interpretation: Empirical Triangulation through the Data + Empathy Framework

The integration of the quantitative and qualitative findings provides a more comprehensive interpretation of how respondents evaluate Artificial Intelligence in healthcare communication. Considered together, the two components reveal neither unconditional enthusiasm for AI nor categorical resistance to its use. Instead, the findings indicate a pattern of **conditional technological acceptance**, in which respondents appear comparatively receptive to computational assistance when its function is clearly delimited, while expressing stronger expectations for transparency, human involvement, and accountability as the informational, clinical, or emotional sensitivity of the interaction increases.

This pattern is particularly relevant to the Data + Empathy framework developed throughout the thesis. The framework does not assume that computational capability and human-centered communication constitute competing alternatives. Rather, it proposes that their relationship should be structured according to the function performed by the technology, the characteristics of the stakeholder context, and the potential consequences of the interaction. The field research does not empirically validate this framework as a predictive model. Nevertheless, several findings **converge with its underlying propositions**, while others contribute additional stakeholder-level considerations that refine its practical interpretation.

1. Data: Computational Capability Requires Evidentiary and Informational Boundaries

The **Data** dimension of the framework concerns the capacity of computational systems to retrieve, process, organize, predict, and adapt information. The survey findings indicate that respondents do not necessarily oppose these capabilities in healthcare communication. Their acceptance, however, appears to vary substantially according to the function assigned to the technology.

This distinction is visible in the activities respondents considered most appropriate for AI. General health information retrieval accounted for **26.5% (n = 39)** of selections, administrative scheduling

for **24.5% (n = 36)**, clarification of medical terminology for **19.0% (n = 28)**, and support to healthcare professionals for **14.3% (n = 21)**. By contrast, highly personalized recommendations, emotional support, and autonomous diagnostic or therapeutic activities received markedly lower levels of acceptance.

The qualitative findings provide additional context for this distribution. Participants frequently distinguished between AI as an **informational or analytical resource** and AI as an autonomous source of clinical judgment. Diagnosis, treatment selection, prescription, and final clinical evaluation were repeatedly identified as activities for which respondents expected human responsibility to remain central. The resulting distinction is therefore not simply between technology and human intervention, but between **computational assistance and consequential authority**.

The importance assigned to scientific provenance further qualifies the Data dimension. When AI-generated health information was described as being accompanied by verifiable scientific sources, **51.0% (n = 75)** of respondents reported high levels of trust, while only 8.8% (n = 13) selected low trust ratings. As discussed in Section 5.3.3, this finding cannot establish that scientific references causally increase trust because the questionnaire did not experimentally manipulate source availability. It nevertheless indicates that the **evidentiary basis of AI-mediated information constitutes a relevant stakeholder consideration**.

This finding is compatible with the architectural principles developed in Section 5.2, where technologies such as Retrieval-Augmented Generation and structured knowledge infrastructures were proposed as possible mechanisms for improving evidence grounding and traceability. Importantly, however, the questionnaire did **not test these technologies directly**. The empirical result therefore supports the importance of verifiable scientific provenance as a stakeholder expectation, rather than validating any specific technical solution for achieving it.

The field research consequently refines the Data component of the framework. Computational capability alone does not define stakeholder value. What appears equally relevant is whether the information processed or generated through computational systems can be connected to an identifiable evidentiary basis and whether the authority assigned to those systems remains proportionate to the consequences of the task.

2. Empathy: Human-Centeredness Extends Beyond Simulated Emotional Language

The **Empathy** dimension receives particularly important clarification from the combination of quantitative and qualitative findings. As defined in the Data + Empathy framework, empathy is not treated as a computational module or as the capacity of an AI system to imitate emotionally expressive language. It functions instead as a human-centered evaluative principle through which technological deployment is assessed in relation to stakeholder needs, vulnerability, context, and experience.

The survey findings are consistent with this distinction. When respondents evaluated the hypothetical empathetic statement produced by a healthcare chatbot, **45.6% (n = 67)** perceived it as artificial or uncomfortable, compared with **14.2% (n = 21)** who regarded it as reassuring. These responses do not establish that AI-generated empathetic communication is inherently ineffective. They do, however, indicate that the presence of empathetic linguistic markers alone was insufficient to produce a reassuring perception for a substantial proportion of the surveyed sample.

The qualitative responses reinforce this distinction by locating empathy in a broader relational context. Participants repeatedly referred to listening, personal history, psychological support, direct contact, and the doctor–patient relationship when describing what should remain under human responsibility. In these responses, empathy is not reduced to the semantic content of an isolated sentence. Instead, it is associated with **contextual understanding, relational continuity, and responsibility toward the person receiving the communication.**

The integration of these findings therefore supports an important distinction already introduced in Section 5.1.4: **human-like communication is not necessarily equivalent to human-centered communication.** An AI system may reproduce the linguistic characteristics of reassurance or concern without this automatically generating the stakeholder experience associated with interpersonal empathy. Conversely, an AI system that performs a narrowly functional task may contribute to a human-centered objective without attempting to imitate human emotional behavior.

This interpretation is illustrated by the case studies examined earlier in the chapter. Scrolling Therapy associated computational capability with the reduction of friction in a rehabilitation activity; The Most Beautiful Sound used computational translation to make complex biological information perceptible through another sensory modality; and The Code introduced explicit limits on synthetic representation where technological deployment could conflict with representational objectives. These cases differ substantially from the survey scenarios, and they cannot be treated as empirical equivalents. Their comparative relevance lies instead in illustrating how **human-centered value may emerge from the design and governance of technological roles rather than from anthropomorphic technological behavior itself.**

Within the Data + Empathy framework, empathy should therefore be interpreted not as something that an organization simply “adds” to AI-generated content, but as a criterion influencing **why a technology is deployed, what function it performs, how stakeholders experience it, and where technological intervention should stop.**

3. Governance: Transparency, Privacy, and Human Accountability as Conditions of Deployment

The integration of the findings also strengthens the importance of the **Governance** dimension introduced in Section 5.2. Although the original conceptual framework is structured around Data and Empathy, its operationalization demonstrated that these dimensions require an organizational mechanism capable of translating computational possibility and human-centered appropriateness into decisions concerning technological deployment.

Transparency emerges as one of the clearest empirical themes in this respect. **82.3% (n = 121)** of respondents attributed high importance to knowing whether healthcare content had been created or modified using Artificial Intelligence. The survey does not establish that disclosure necessarily produces trust, but it indicates that technological provenance itself constitutes relevant information for a substantial majority of the surveyed sample.

A similar relationship emerges in the context of behavioral personalization. **60.5% (n = 89)** expressed low comfort with the use of online behavioral information to personalize health-related content. At the same time, **67.3% (n = 99)** stated that an explanation of which data were being used and for what

purpose would make them at least somewhat more comfortable. These results should not be interpreted as demonstrating that transparency eliminates privacy concerns. Rather, they indicate that **stakeholder evaluation of personalization depends not only on the relevance of the resulting message but also on the perceived legitimacy and intelligibility of the underlying data-processing process.**

The findings therefore provide stakeholder-level context for the Personalization Paradox discussed in Chapter 3. Greater computational precision may enable increasingly individualized communication, yet the same data practices required to achieve that precision may generate discomfort when stakeholders perceive them as intrusive or insufficiently transparent. The field research does not resolve this tension, but it demonstrates its practical relevance within the surveyed sample.

Human accountability constitutes the second major governance theme. In serious or emotionally sensitive scenarios, **82.3% (n = 121)** preferred direct interaction with a healthcare professional, while no respondent selected a fully autonomous AI system. Similarly, **60.5% (n = 89)** identified a backend support function as the most appropriate general role for AI in healthcare communication. The qualitative responses repeatedly reinforced this preference through references to verification, professional responsibility, final decision-making, and the need for an identifiable human actor to retain authority over consequential outcomes.

These findings are consistent with the **risk-proportionate Human-in-the-Loop architecture** proposed in Section 5.2. They should not, however, be interpreted as evidence that every AI-mediated healthcare interaction requires identical levels of human supervision. The quantitative results themselves indicate greater acceptance of direct AI involvement in lower-risk informational and administrative activities. The empirical pattern therefore supports a more differentiated principle: **the appropriate level of human involvement appears to increase as perceived clinical, informational, or emotional consequences become more significant.**

4. From Backend Enablement to Risk-Proportionate Functional Asymmetry

When the different empirical dimensions are considered together, the frequently used distinction between AI operating in the “backend” and humans remaining in the “frontend” requires further refinement. The preference expressed by **60.5% of respondents for AI operating behind the scenes in support of healthcare professionals** provides a clear indication that this configuration is acceptable to a substantial proportion of the sample. However, the remaining results demonstrate that technological visibility alone cannot determine appropriate role allocation.

Respondents also expressed acceptance of direct AI interaction for certain simple functions, while their reservations increased when the technology was associated with diagnosis, highly personalized recommendations, emotional support, or other sensitive activities. The more appropriate conceptual distinction is therefore not simply **invisible AI versus visible humans**, but **risk-proportionate functional asymmetry.**

Under this interpretation, AI may assume comparatively greater operational autonomy where tasks are standardized, low-risk, informational, administrative, or easily reversible. Human involvement becomes progressively more important where communication requires clinical interpretation, emotionally sensitive dialogue, management of uncertainty, or consequential decision-making. This

does not establish a universal threshold between automated and human functions; rather, it provides an organizational principle through which technological roles can be evaluated contextually.

The quantitative and qualitative findings converge particularly clearly on this point. Participants do not appear to reject computational assistance itself. Instead, they repeatedly differentiate between **assistance and authority**, between **information and judgment**, and between **linguistic simulation and interpersonal responsibility**. These distinctions transform the Data + Empathy framework from a simple division of technological and human functions into a more flexible model of role allocation.

The operational logic can therefore be represented as:

Stakeholder Context → Data-Informed Understanding → Assessment of Functional and Contextual Risk → Appropriate Technological Role → Human Oversight Proportionate to Sensitivity → Stakeholder Interaction → Continuous Evaluation

Within this sequence, AI is neither categorically confined to an invisible backend nor assumed to be appropriate as an autonomous frontend communicator. Its degree of visibility and authority depends on the relationship between the function performed, the sensitivity of the context, the evidentiary requirements of the communication, and the consequences of potential failure.

5. Empirical Refinement of the Data + Empathy Framework

The combined field research ultimately contributes to the thesis by **refining rather than validating** the Data + Empathy framework. The case studies in Section 5.1 illustrated different configurations of computational utility, translational accessibility, and ethical boundary-setting. Section 5.2 subsequently translated these observations into an operational architecture based on Data, Empathy, and Governance. The survey adds a third perspective by examining how younger adult respondents evaluate selected manifestations of AI-mediated healthcare communication.

Across these three analytical levels, a common principle emerges. **Technological capability does not independently determine appropriate technological deployment.** Computational systems may provide substantial advantages in information retrieval, data processing, personalization, administrative support, and professional decision support. However, their communicative legitimacy depends on additional considerations concerning evidence, transparency, privacy, stakeholder vulnerability, human accountability, and contextual sensitivity.

The field research also complicates any simplistic characterization of younger audiences as naturally predisposed toward AI-mediated communication. Within the surveyed sample, **76.9% reported frequent or daily AI use**, and greater frequency of AI use was moderately associated with greater comfort receiving AI-mediated health information ($r_s = .347, p < .001$). Yet this familiarity coexisted with low comfort regarding behavioral personalization, strong expectations concerning AI disclosure, pronounced sensitivity to algorithmic error, and clear preferences for human involvement in consequential situations.

The resulting pattern can be summarized as:

High AI Familiarity → Conditional Acceptance → Evidentiary and Transparency Expectations → Contextual Boundaries on Personalization and Automation → Increasing Human Responsibility as Sensitivity Rises

This sequence should not be interpreted as a causal model established by the survey. It represents an **integrative interpretation of the descriptive and exploratory findings** and provides a conceptual summary of the relationships observed within the sample.

The empirical analysis also reinforces the distinction between **personalization and human-centeredness**. A communication system may be highly individualized from a computational perspective while simultaneously being perceived as intrusive if the underlying data practices are unclear or unacceptable. Conversely, a comparatively simple technological function may generate stakeholder value when it addresses a concrete need without unnecessarily expanding technological authority. Within the Data + Empathy framework, personalization should therefore not be evaluated solely according to the precision with which a message can be adapted to an individual, but also according to whether that adaptation remains understandable, proportionate, transparent, and appropriate to the stakeholder context.

The framework can consequently be synthesized through three interdependent questions:

Data: *What can the system know, retrieve, predict, or adapt?*

Empathy: *What should be communicated, how should the stakeholder experience the interaction, and where should technological intervention stop?*

Governance: *Under what conditions should the organization permit, limit, supervise, or reject the proposed technological application?*

The field research does not provide definitive answers to these questions, nor can findings derived from a non-probability sample of 147 respondents be generalized to Generation Z as a whole. Its contribution is instead exploratory: it demonstrates that the boundaries surrounding AI-mediated healthcare communication are **multidimensional and context-dependent**, and that technological familiarity alone is insufficient to explain stakeholder acceptance.

Accordingly, the Data + Empathy framework should not be understood as an argument for either maximum automation or systematic technological restraint. It proposes a model of **appropriate technological deployment**, in which computational capability is evaluated against stakeholder utility, evidentiary integrity, privacy expectations, contextual sensitivity, and human accountability.

The combined evidence developed throughout Chapter 5 therefore suggests a transition from a technology-centered question—“*How much of healthcare communication can Artificial Intelligence automate?*”—toward a stakeholder-centered and governance-oriented question:

“Which functions should Artificial Intelligence perform, under what conditions, and where does human responsibility create value that technological capability alone cannot provide?”

This shift represents the central empirical and strategic contribution of the chapter. Rather than positioning Data and Empathy as opposing forces, the analysis indicates that responsible AI-mediated pharmaceutical communication depends on their coordinated governance: **Data establishes computational possibility; Empathy evaluates that possibility in relation to human context; and Governance determines whether, where, and under what conditions it should become organizational action.**

Chapter 6: Conclusions and Future Outlook

The rapid integration of Artificial Intelligence into pharmaceutical and healthcare communication is transforming not only the technological infrastructure through which information is produced and distributed, but also the strategic relationship between organizations, healthcare professionals, patients, and increasingly digitally mediated audiences. Throughout this thesis, this transformation has been examined as a multidimensional process in which technological capability intersects with regulatory constraints, psychological responses, privacy expectations, scientific accountability, and the need to preserve meaningful human responsibility.

The analysis has demonstrated that the central challenge facing pharmaceutical organizations is not simply whether Artificial Intelligence should be adopted. AI-driven technologies already provide substantial capabilities across content generation, information retrieval, predictive modeling, audience segmentation, compliance support, and communication orchestration. The more consequential strategic question concerns the role that such technologies should be allowed to assume within highly sensitive communication environments.

This distinction becomes particularly important in the life sciences sector. Greater computational precision does not automatically generate greater stakeholder value, just as greater personalization does not necessarily produce more human-centered communication. AI may enhance relevance, accelerate workflows, improve access to information, and support professional decision-making. At the same time, the same technological infrastructures can introduce risks related to opacity, excessive behavioral profiling, misinformation, automation bias, inappropriate emotional simulation, and the erosion of identifiable human accountability.

The thesis has therefore progressively moved away from a binary interpretation in which technological innovation and human-centered communication are treated as opposing alternatives. Instead, the findings support a model of differentiated technological role allocation in which the degree of automation, visibility, and human oversight should vary according to the function performed, the sensitivity of the communicative context, the strength of the available evidence, the regulatory environment, and the consequences of potential failure.

Within this perspective, the Data + Empathy framework developed throughout the research provides an interpretative and operational structure for evaluating AI-mediated pharmaceutical communication. Its central proposition is not that Artificial Intelligence should always remain invisible, nor that every stakeholder interaction should remain exclusively human. Rather, technological deployment should remain proportionate to contextual risk and demonstrable stakeholder utility. Data establishes computational possibility; Empathy evaluates that possibility against human needs and circumstances; Governance determines whether, where, and under what conditions that possibility should become organizational action.

6.1 Synthesis of Findings and Response to the Research Questions

The first research question asked: **How is AI enabling the transition from mass marketing to precision communication in pharmaceutical digital advertising?**

The analysis conducted in Chapters 2 and 3 indicates that Artificial Intelligence contributes to this transition by altering both the informational foundations and the operational architecture of

pharmaceutical communication. Traditional broadcast models relied on comparatively static audience segmentation, standardized content assets, and largely linear communication flows. AI-mediated infrastructures instead enable organizations to process heterogeneous streams of behavioral and contextual data, dynamically adapt content, identify emerging informational needs, and coordinate communication across multiple channels.

Technologies such as predictive behavioral modeling, Next Best Action systems, Retrieval-Augmented Generation, modular content architectures, and Dynamic Creative Optimization illustrate different dimensions of this transformation. Their combined strategic relevance lies in the transition from generalized communication toward increasingly contextual forms of interaction in which message, channel, timing, and informational depth can potentially be adapted to stakeholder circumstances.

However, the findings also demonstrate that precision communication should not be equated exclusively with increasing personalization intensity. The Personalization Paradox examined throughout the thesis shows that the same data infrastructures that enable greater relevance can also increase perceived intrusiveness when stakeholders do not understand how their information is collected, processed, or used. Precision therefore becomes strategically sustainable only when accompanied by appropriate transparency, evidentiary grounding, privacy safeguards, and contextual restraint.

RQ1 can consequently be answered through a distinction between technological precision and responsible precision. Artificial Intelligence enables the transition away from mass communication by expanding the organization's capacity to understand and respond to stakeholder context. Yet the strategic value of this transformation depends on whether computational precision is translated into meaningful utility rather than simply into greater persuasive capability.

The second research question asked: **How do pharmaceutical companies navigate regulatory “gray areas,” including borderline contexts and platform-based self-regulation?**

The regulatory analysis demonstrated that AI-mediated pharmaceutical communication operates within a fragmented governance environment. National regulatory systems remain essential in establishing legally binding boundaries for health-related communication, but their geographically bounded and procedurally structured enforcement mechanisms increasingly interact with digital platforms characterized by speed, decentralization, and cross-jurisdictional distribution.

The examination of unbranded disease-awareness communication, borderline health products, influencer-mediated health narratives, and the wider semaglutide communication environment illustrates how digital media can complicate the conventional distinction between advertising, education, organic communication, and independent user-generated content. At the same time, platform self-regulation cannot independently resolve these challenges. Commercial moderation systems operate at scale, but they are not equivalent to pharmaceutical regulatory assessment and may be poorly suited to evaluating scientific substantiation, contextual claims, implicit promotion, or the appropriateness of health-related persuasion.

The research therefore supports a layered rather than substitutive governance model. Statutory regulation, platform moderation, internal Compliance by Design mechanisms, and Human-in-the-

Loop oversight perform different functions and should not be expected to replace one another. Pharmaceutical organizations consequently require internal governance capabilities capable of operating before communication reaches the public domain.

Within this framework, automated pre-screening, evidence-grounded generation, structured content architectures, and traceable information sources can contribute to regulatory alignment. Nevertheless, automated verification remains supportive rather than definitive. Final responsibility for consequential communication must remain identifiable within the organization, particularly when scientific interpretation, legal ambiguity, or patient vulnerability exceeds what computational evaluation can reliably resolve.

RQ2 therefore suggests that regulatory gray areas cannot be sustainably navigated through regulatory avoidance or technological circumvention. They require proactive organizational governance capable of integrating legal, medical, technological, and ethical considerations earlier in the communication lifecycle.

The third research question asked: **What perceptions and trust dynamics characterize AI-mediated health communication among the Generation Z respondents examined in this research?**

The exploratory field research produces a more differentiated answer than either technological enthusiasm or generalized AI skepticism. The final analytical sample consisted of 147 respondents aged between 18 and 29 years. Despite high levels of technological familiarity, with 76.9% reporting frequent or daily AI use, comfort with receiving AI-mediated health information remained centered around the midpoint of the scale. More frequent AI use was moderately associated with greater comfort, but the relationship did not eliminate substantial variation in attitudes.

This finding is significant because it demonstrates that technological familiarity and healthcare acceptance should not be treated as interchangeable. Respondents differentiated strongly according to the function performed by the technology. AI was comparatively more acceptable for informational, administrative, explanatory, and professional-support activities, while considerably fewer participants selected highly personalized recommendations, emotional support, autonomous diagnosis, or therapeutic decision-making as appropriate functions. In a serious or emotionally sensitive health scenario, 82.3% preferred direct interaction with a healthcare professional, and no participant selected a fully autonomous AI system.

Trust was also strongly connected to informational conditions. Verifiable scientific sources were associated with comparatively higher stated trust, while AI disclosure emerged as particularly important: 82.3% of respondents assigned high importance to knowing whether health-related content had been generated or modified through Artificial Intelligence. At the same time, behavioral personalization generated substantial reservations, with 60.5% reporting low comfort regarding the use of online behavior to personalize health-related content. The fact that 67.3% stated that greater explanation of data use would make them at least somewhat more comfortable further suggests that procedural transparency is relevant, although the non-experimental design does not establish a causal effect.

The field research also demonstrated the fragility of trust following technological error. A substantial majority of respondents anticipated some reduction in trust or discontinuation of use following an observable error in AI-generated health information, while incorrect information emerged as the most frequently selected primary concern. Qualitative responses reinforced this pattern by repeatedly assigning diagnostic judgment, therapeutic decisions, emotional responsibility, final accountability, and informational integrity to human professionals.

Taken together, these findings suggest that the surveyed respondents display **conditional acceptance** of AI-mediated healthcare communication. Familiarity with Artificial Intelligence does not imply willingness to delegate every communicative or decisional function to it. Acceptance appears to vary according to functional complexity, perceived risk, evidence availability, data practices, transparency, and the continued presence of human responsibility.

The fourth research question asked: **How can human-centered evaluation and data-driven precision be integrated through the Data + Empathy framework to support responsible and credible AI-mediated pharmaceutical communication?**

The findings require a cautious response to this question. The present research did not longitudinally measure brand credibility or long-term stakeholder relationships and therefore cannot demonstrate that the Data + Empathy framework causally increases either outcome. The contribution of the thesis lies instead in identifying organizational and communicative conditions that may support credible and responsible stakeholder relationships.

The comparative cases demonstrated that computational technologies can create stakeholder value through different configurations. Scrolling Therapy illustrates functional utility, in which computational capability is directed toward reducing friction within a rehabilitation experience. The Most Beautiful Sound illustrates translational accessibility, in which complex biological information is transformed into a sensory representation. The Code provides an adjacent-sector example of ethical boundary-setting, illustrating that responsible technological strategy may also involve deliberately limiting automation where organizational values prioritize authentic human representation.

These cases, together with the field research, suggest that credibility is unlikely to emerge simply from either greater AI deployment or greater human visibility. More relevant conditions include clear functional purpose, scientific traceability, transparency, respect for privacy, appropriate technological restraint, identifiable accountability, and human involvement proportionate to the sensitivity of the interaction.

The Data + Empathy framework therefore does not claim to guarantee trust. Instead, it provides a structured model through which organizations can evaluate whether technological deployment is likely to remain compatible with stakeholder needs and responsible communication. In this sense, its potential contribution to credibility lies in shifting pharmaceutical AI strategy from **automation maximization** toward **appropriate technological role allocation**.

6.2 Theoretical Contribution: The Data + Empathy Framework

The primary conceptual contribution of this thesis is the progressive development of the Data + Empathy framework as an interpretative model for AI-mediated pharmaceutical communication.

The framework initially emerged from the tension between two capabilities that are frequently discussed independently. On one side, Artificial Intelligence provides unprecedented capacity for data processing, content adaptation, predictive modeling, and communication scalability. On the other, healthcare communication remains dependent on contextual judgment, scientific responsibility, emotional sensitivity, and interpersonal accountability.

A simplistic resolution would separate these functions completely, placing Artificial Intelligence within an invisible computational backend while reserving the visible communication interface for human actors. Although this configuration remains relevant in many high-risk contexts, the empirical findings indicate that such a division is too rigid to account for the variety of stakeholder attitudes observed in practice.

Respondents did not uniformly reject direct interaction with Artificial Intelligence. Comparatively greater acceptance emerged when tasks were simple, informational, administrative, or supportive. Reservations became substantially more pronounced as AI approached diagnosis, therapy, highly individualized recommendations, emotional support, or other areas associated with greater human consequences. The conceptual framework therefore evolved toward **risk-proportionate functional asymmetry**.

Under this principle, Artificial Intelligence and human professionals do not occupy fixed technological territories. Their responsibilities are asymmetrical because they possess different capabilities and limitations, but the precise allocation of functions depends on context. Automation may expand where tasks are standardized, evidence-grounded, reversible, and comparatively low-risk. Human involvement should become progressively stronger where uncertainty, emotional vulnerability, clinical interpretation, accountability, or irreversible consequences increase.

This evolution produces a second theoretical contribution: the distinction between **human-like communication and human-centered communication**. Human-centeredness does not require an algorithm to imitate emotional language, simulate intimacy, or anthropomorphize its interaction. The qualitative and quantitative findings instead support an interpretation of empathy as a design and governance principle.

Within the framework, Empathy does not constitute another computational module. It represents the evaluative process through which organizations consider how technological intervention will be experienced by the stakeholder, whether the system addresses an actual need, whether the interaction preserves agency and dignity, and whether there are contexts in which technological intervention should be reduced or rejected entirely.

A third contribution concerns the distinction between **personalization and human-centeredness**. Computational systems may generate extremely individualized communication while still producing perceptions of intrusion or surveillance. Conversely, a comparatively simple digital intervention may create substantial value when it addresses a clearly defined stakeholder problem without requiring disproportionate data extraction or technological authority.

This distinction reframes personalization from an end in itself into a conditional strategic instrument. The relevant question becomes not how precisely a pharmaceutical organization can target an

individual, but whether the resulting adaptation is useful, understandable, proportionate, scientifically grounded, and ethically appropriate.

The final conceptual architecture can therefore be expressed through three interdependent dimensions:

Data concerns what computational systems can retrieve, process, infer, predict, or adapt.

Empathy concerns what should be communicated, how the interaction should be experienced, which human need technological intervention addresses, and where technological involvement should stop.

Governance concerns the conditions under which the organization permits, constrains, supervises, audits, or rejects the proposed technological application.

The resulting framework can be summarized through the following operational sequence:

Stakeholder Need → Data-Informed Understanding → Computational Enablement → Human-Centered Appropriateness Assessment → Risk-Proportionate Governance → Human-Governed Deployment → Continuous Evaluation

The value of this model lies not in claiming a universal formula for responsible Artificial Intelligence. Rather, it provides a decision architecture through which technological capability can be evaluated against human context, evidentiary standards, regulatory exposure, and organizational accountability.

Accordingly, Data + Empathy should be interpreted as a **conceptual and applied contribution rather than an empirically validated theory**. The case studies illustrate its relevance across different technological configurations, and the field research provides exploratory stakeholder-level evidence that refines several of its assumptions. Further empirical testing remains necessary before claims can be made regarding its comparative effectiveness relative to alternative governance or communication models.

6.3 Managerial and Governance Implications for Pharmaceutical Digital Communication

The findings developed throughout the thesis generate several implications for pharmaceutical organizations considering the expansion of AI-mediated communication.

The first concerns the starting point of technological implementation. AI strategy should begin with an identifiable stakeholder or organizational problem rather than with technological capability itself. The ability to generate more content, collect more behavioral data, or automate additional interactions does not independently justify deployment. Organizations should first establish what problem the technology is intended to address, who is affected by the intervention, and whether automation represents an appropriate mechanism for creating value.

This principle shifts pharmaceutical communication from a capacity-oriented toward a utility-oriented model. In a capacity-oriented approach, Artificial Intelligence is primarily evaluated according to how much additional communication, personalization, or operational efficiency it enables. In a utility-oriented approach, the primary criterion becomes whether that capability improves access to information, reduces unnecessary friction, supports professional activity, enhances comprehension, or addresses another demonstrable stakeholder need.

The second implication concerns technological role allocation. Organizations should avoid developing a single enterprise-wide assumption regarding the appropriate level of AI autonomy. Different applications require different governance intensities. Administrative assistance, evidence retrieval, terminology clarification, workflow coordination, or other comparatively low-risk functions may tolerate greater levels of automation. Communication involving diagnosis, treatment, significant emotional vulnerability, ambiguous scientific interpretation, or consequential decision-making requires stronger human involvement.

Human-in-the-Loop governance should therefore be designed proportionately rather than symbolically. Human oversight is unlikely to generate value merely because a person formally appears at the end of an automated workflow. The relevant question is whether qualified professionals possess sufficient information, authority, time, and organizational responsibility to intervene meaningfully when contextual interpretation is required.

A third managerial implication concerns scientific provenance. The survey findings suggest that stakeholders place meaningful value on verifiable sources, while the technical analysis demonstrates the importance of evidence-grounded generation within pharmaceutical environments. Organizations should consequently treat traceability not as a secondary compliance feature but as a central component of communication design.

Where AI-assisted systems retrieve, summarize, adapt, or generate medical information, organizations should be able to establish which evidence supports the output, how that evidence was selected, and which human or technological processes verified its appropriateness. Retrieval architectures, structured knowledge repositories, metadata, and content-governance systems may support this objective, but they do not remove the need for professional review.

Transparency produces a comparable managerial implication. AI disclosure and data-use explanations should not be treated exclusively as legal obligations to be minimized within consent mechanisms. They can instead be understood as components of stakeholder relationship management. The field research indicates substantial interest in knowing whether health-related content has been generated or modified through AI and substantial discomfort surrounding opaque behavioral personalization.

This does not mean that greater disclosure will automatically create acceptance. The survey did not experimentally establish such a causal effect. Transparency is better understood as one condition of legitimate technological deployment: it enables stakeholders to interpret the interaction more accurately, understand the role of automation, and make more informed decisions concerning participation.

The findings also indicate that organizations should reconsider how AI-driven communication performance is evaluated. Traditional digital marketing metrics such as engagement, click-through rates, conversion, reach, and production efficiency remain commercially relevant, but they may be insufficient in regulated healthcare contexts. A highly engaging communication asset can simultaneously be intrusive, scientifically misleading, emotionally inappropriate, or inconsistent with stakeholder autonomy.

Performance evaluation may therefore require complementary measures concerning information accessibility, evidence traceability, user understanding, perceived intrusiveness, error frequency,

escalation effectiveness, stakeholder control, and the appropriateness of human oversight. Responsible AI communication should be evaluated according to the quality of the relationship it creates between technological capability and stakeholder need, not exclusively according to behavioral response.

Finally, the comparison of the three cases examined in Chapter 5 suggests that responsible innovation also includes the organizational capacity to decide **not** to automate. Technological restraint should not necessarily be interpreted as resistance to innovation. Where synthetic representation, autonomous communication, behavioral targeting, or emotional simulation conflicts with stakeholder vulnerability or organizational values, deliberate non-deployment may itself constitute a legitimate governance decision.

The managerial objective is therefore not the construction of the most automated pharmaceutical communication ecosystem possible. It is the development of an organizational capability capable of determining which forms of automation create meaningful value, which require stronger safeguards, and which should remain outside the boundaries of technological delegation.

6.4 Research Limitations

The findings and conceptual propositions developed in this thesis should be interpreted in light of several methodological and analytical limitations.

The first limitation concerns the exploratory field research. Participants were recruited through digital personal and online networks using a non-probability convenience sampling strategy with elements of snowball dissemination. The final analytical sample consisted of 147 respondents aged between 18 and 29 years and was strongly concentrated between the ages of 21 and 26. Consequently, the sample cannot be considered statistically representative of Generation Z in Italy or internationally.

Self-selection may have further influenced the results. Individuals who voluntarily completed a questionnaire concerning Artificial Intelligence may possess greater technological interest, familiarity, or engagement than those who chose not to participate. This consideration is especially relevant given the high reported frequency of AI use within the sample.

The survey was also cross-sectional and predominantly based on self-reported attitudes. Responses therefore represent stated perceptions under hypothetical scenarios rather than observed behavior in real healthcare environments. For example, respondents who indicated that clearer explanations of data use would increase their comfort with personalization were reporting an anticipated reaction rather than responding to an experimental manipulation. The study consequently cannot establish causal relationships between transparency, personalization, trust, or technological acceptance.

Similarly, the moderate association identified between frequency of AI use and comfort with receiving AI-mediated health information should not be interpreted as demonstrating that repeated AI use causes greater acceptance. Other variables not measured within the questionnaire may contribute to the relationship.

The instrument itself was intentionally concise and exploratory. This supported participation and reduced respondent burden, but necessarily limited the depth with which individual psychological and communication constructs could be measured. The qualitative component consisted of one

optional open-ended question and therefore provides contextual depth rather than a comprehensive qualitative investigation.

The case-study component presents additional limitations. Eurofarma's Scrolling Therapy, ASCO's The Most Beautiful Sound, and Dove's The Code were purposively selected because they represent different configurations of technological utility, scientific translation, and ethical boundary-setting. Their selection was analytical rather than statistically representative. They should therefore not be interpreted as reflecting the overall state of AI adoption across the pharmaceutical industry.

The Dove case also originates from an adjacent personal-care and wellness sector rather than pharmaceutical communication directly. Its inclusion provides an external governance benchmark concerning synthetic representation, but limits the extent to which sector-specific conclusions can be transferred directly to pharmaceutical organizations.

The case analysis also relies substantially on publicly available campaign and organizational documentation. Internal performance data, implementation processes, stakeholder interviews, longitudinal outcomes, and proprietary decision-making criteria were not available. Consequently, the cases illustrate possible configurations of responsible technological deployment rather than demonstrating measurable causal effects on patient outcomes, stakeholder trust, or brand performance.

A further limitation concerns the Data + Empathy framework itself. The framework represents a conceptual synthesis developed through theoretical analysis, case interpretation, operational modeling, and exploratory stakeholder research. It has not been experimentally validated, implemented across multiple pharmaceutical organizations, or directly compared with alternative AI-governance frameworks. Claims concerning its effectiveness should therefore remain prospective.

This limitation is particularly relevant to RQ4. While the framework identifies mechanisms that may plausibly support credibility—including transparency, evidence grounding, stakeholder utility, appropriate human oversight, and accountable governance—the present research does not demonstrate longitudinal improvements in brand credibility or stakeholder relationships.

The regulatory analysis is also necessarily selective. Pharmaceutical communication operates across heterogeneous national legal systems, therapeutic areas, product classifications, and platform environments. The thesis focuses on selected European and United States regulatory principles and uses additional cases to examine specific areas of regulatory friction. It does not constitute an exhaustive comparative analysis of global pharmaceutical law.

Finally, the technological environment examined throughout this thesis is evolving rapidly. Large Language Models, generative media systems, recommendation architectures, privacy-preserving technologies, and regulatory standards are likely to continue changing. Some operational assumptions discussed within the thesis may therefore require revision as technological capabilities, organizational practices, and regulatory expectations evolve.

These limitations do not negate the analytical contribution of the research. Rather, they establish the boundaries within which its findings should be interpreted and provide a foundation for future empirical investigation.

6.5 Future Research and Concluding Remarks

Future research can develop the findings of this thesis across several complementary directions.

A first priority concerns the empirical examination of stakeholder attitudes across larger and more representative populations. Probability-based or stratified samples could investigate whether the patterns observed among the present respondents are replicated across different national contexts, demographic groups, healthcare experiences, and levels of AI familiarity. Comparative research could also examine whether patients, caregivers, healthcare professionals, and general consumers establish different boundaries concerning technological autonomy.

Experimental designs would provide an important second direction. Future studies could manipulate factors such as AI disclosure, evidence-source visibility, degree of personalization, human supervision, conversational tone, and error occurrence to determine how these variables causally influence trust, perceived credibility, willingness to engage, and information comprehension.

Such research would be particularly valuable for distinguishing between the effects of technological visibility and technological authority. The present study suggests that respondents may accept direct AI interaction in lower-risk circumstances while preferring stronger human involvement as sensitivity increases. Experimental research could therefore examine whether stakeholder responses depend primarily on whether AI is visible, what function it performs, what degree of decisional authority it receives, or the consequences associated with potential failure.

The Data + Empathy framework itself represents another important avenue for investigation. Future studies could operationalize its components into measurable dimensions and test the proposed principle of risk-proportionate functional asymmetry across different communication scenarios. This could eventually enable the development of practical evaluation tools through which organizations classify AI use cases according to stakeholder vulnerability, evidentiary requirements, privacy exposure, reversibility, regulatory sensitivity, and required human oversight.

Organizational research within pharmaceutical companies would also provide substantial value. Interviews or comparative case studies involving Medical, Legal and Regulatory teams, marketers, data scientists, healthcare professionals, compliance officers, and external agencies could examine how AI-governance decisions are made in practice. Such research could determine whether the functional divisions proposed in this thesis correspond to actual enterprise workflows and identify organizational barriers to implementing risk-proportionate governance.

Longitudinal studies are particularly necessary to address the relationship between AI-mediated communication and brand credibility. Trust is unlikely to be a static response to a single technological interaction. It may accumulate, deteriorate, or recover across repeated experiences, disclosures, successful interactions, and observable errors. Studying stakeholder relationships over time would therefore allow future research to examine the longitudinal relationship between AI-mediated communication and credibility, which the present study did not measure.

Further investigation is also required regarding synthetic empathy. The present field research examined a single hypothetical empathetic statement and identified substantial ambivalence toward this form of communication. More sophisticated research could compare different conversational styles, degrees of anthropomorphism, disclosure conditions, clinical contexts, and levels of human

supervision. This would help distinguish resistance to artificial emotional simulation from resistance to automated healthcare communication more broadly.

From a regulatory perspective, future research should continue examining how AI-mediated communication modifies conventional distinctions between advertising, information, recommendation, personalization, and automated decision support. As generative technologies increasingly participate in content creation and stakeholder interaction, the boundaries between technological provider, communication intermediary, and accountable organizational actor are likely to require further conceptual and regulatory clarification.

Ultimately, the future of Artificial Intelligence in pharmaceutical communication is unlikely to be defined by a simple choice between automation and human involvement. The more significant challenge concerns how responsibilities are distributed between them.

Artificial Intelligence offers substantial opportunities to process scientific complexity, support regulatory workflows, identify informational needs, adapt communication, and extend organizational capabilities. Yet technological possibility does not independently determine communicative legitimacy. The greater the sensitivity of the stakeholder context, the more important questions of transparency, evidence, autonomy, accountability, and human responsibility become.

The findings of this thesis therefore support a shift in strategic focus. The central objective should not be to make pharmaceutical communication maximally automated, nor to preserve human intervention merely as a symbolic response to technological anxiety. The objective should be to establish the **appropriate relationship between computational capability and human responsibility**.

Within this relationship, Data provides the capacity to understand, retrieve, predict, and adapt. Empathy ensures that technological capability remains connected to actual human circumstances rather than becoming an end in itself. Governance transforms both dimensions into accountable organizational decisions.

The future of responsible pharmaceutical digital communication may therefore depend less on how closely Artificial Intelligence can imitate human communication and more on how effectively organizations can determine **where technology creates genuine value, where human responsibility remains indispensable, and how the boundary between the two should change as context and risk evolve**.

In this sense, the central contribution of the Data + Empathy framework is not a prescription for maximum automation or technological restraint. It is a principle of **appropriate technological deployment**: using computational intelligence where it can strengthen scientific, informational, and operational capabilities, while preserving meaningful human judgment wherever communication becomes consequential to individual autonomy, trust, and health.

References

- Advertising Standards Council of India. (2021). *Guidelines for influencer advertising in digital media*.
- Amini Khiabani, G., & Rezaeian, Z. (2026). Modeling ethics in AI-driven marketing strategies to enhance market demand: Evidence from the Iranian pharmaceutical industry. *Journal of Ethics in Entrepreneurship, Technology and Decision Making*. <https://doi.org/10.1108/JEETDM-01-2026-0001>
- Amugongo, L. M., Mascheroni, P., Brooks, S., Doering, S., & Seidel, J. (2025). Retrieval augmented generation for large language models in healthcare: A systematic review. *PLOS Digital Health*, 4(6), e0000877. <https://doi.org/10.1371/journal.pdig.0000877>
- Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Chen, X., Olsson, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., Kerr, J., ... Kaplan, J. (2022). Constitutional AI: Harmlessness from AI feedback. *arXiv preprint arXiv:2212.08073*. <https://doi.org/10.48550/arXiv.2212.08073>
- Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.
- Catalani, V., Negri, A., Townshend, H., Simonato, P., & Al-Imam, A. (2021). The market of sport supplement in the digital era: A netnographic analysis of perceived risks, side-effects and other safety issues. *Emerging Trends in Drugs, Addictions and Health*, 1, Article 100015. <https://doi.org/10.1016/j.etdah.2021.100015>
- Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology*, 39(5), 752–766. <https://doi.org/10.1037/0022-3514.39.5.752>
- Chantre, T., Moreira, I. A., Oliveira, M. C., & Sousa, H. (2025). Does your generation influence your perception of the application of Artificial Intelligence in medicine? *European Archives of Oto-Rhino-Laryngology*, 282, 5375–5380. <https://doi.org/10.1007/s00405-025-09578-4>
- Chou, W. Y. S., Oh, A., & Klein, W. M. P. (2018). Addressing health-related misinformation on social media. *JAMA*, 320(23), 2417–2418. <https://doi.org/10.1001/jama.2018.16865>
- Consumer Protection Act, No. 35 of 2019 (India).
- Dentsu Creative. (2023). *Eurofarma – Scrolling Therapy* [Campaign case study]. Cannes Lions International Festival of Creativity. <https://www.dentsucreative.com/work/scrolling-therapy>
- Deohans, A. D., Nandeshwar, M. D., Kharate, S. A., & Deshmukh, S. P. (2026). Social media influencers and pharmaceutical products: A review of risks, misuse, and regulatory challenges. *International Journal of Current Research and Techniques*, 16(2), 104–118. <https://doi.org/10.5281/zenodo.20068600>
- Dietary Supplement Health and Education Act of 1994, Pub. L. No. 103-417, 108 Stat. 4325 (1994).

- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
- Directive 2001/83/EC of the European Parliament and of the Council of 6 November 2001 on the Community code relating to medicinal products for human use. (2001). *Official Journal of the European Communities*, L 311, 67–128.
- Dove. (2024). *The Code: 20 years of the Campaign for Real Beauty and the commitment to real women in the age of AI* [Press release and campaign documentation]. Unilever. <https://www.dove.com/us/en/stories/campaigns/the-code.html>
- Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407. <https://doi.org/10.1561/04000000042>
- European Medicines Agency. (2017). *Guideline on good pharmacovigilance practices (GVP): Module VI – Collection, management and submission of reports of suspected adverse reactions to medicinal products (Rev 2)*. EMA/873138/2011.
- Federal Trade Commission. (2023). *Health products compliance guidance*. U.S. Federal Trade Commission. <https://www.ftc.gov/business-guidance/resources/health-products-compliance-guidance>
- Fugh-Berman, A. (2008). Prescription tracking and public health. *Journal of General Internal Medicine*, 23(8), 1277–1280. <https://doi.org/10.1007/s11606-008-0630-0>
- Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*. <https://doi.org/10.48550/arXiv.2312.10997>
- Gell, S., Dave, S. P., Willis, E., Vitale, E. J., Woloshin, S., & Heiss, R. (2026). Prescription drug promotion by social media influencers: A systematic scoping review. *JAMA Network Open*, 9(3), Article e262738. <https://doi.org/10.1001/jamanetworkopen.2026.2738>
- Gomes, A. B., & Sultan, A. (2024). Problematizing content moderation by social media platforms and its impact on digital harm reduction. *Harm Reduction Journal*, 21(1), 45–60. <https://doi.org/10.1186/s12954-024-00987-x>
- Grey Health & Wellness. (2023). *ASCO: The Most Beautiful Sound* [Campaign case study]. Cannes Lions International Festival of Creativity. <https://www.grey.com/work/the-most-beautiful-sound>
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Yan, Y., Ishii, E., Yeung, Y. J., Del-Agua, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1–38. <https://doi.org/10.1145/3571730>
- Jones, C. H., Madhavan, S., Natarajan, K., & Corbo, M. (2024). Rewriting the textbook for pharma: How to adapt and thrive in a digital, personalized and collaborative world. *Drug Discovery Today*, 29(3), 103890. <https://doi.org/10.1016/j.drudis.2024.103890>

- Karpukhin, V., Oguz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., & Yih, W. (2020). Dense passage retrieval for open-domain question answering. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 6769–6781. <https://doi.org/10.18653/v1/2020.emnlp-main.550>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 9459–9474. <https://doi.org/10.48550/arXiv.2005.11401>
- Liang, B. A., & Mackey, T. K. (2011). Prevalence and global health implications of social media in direct-to-consumer drug advertising. *Journal of Medical Internet Research*, 13(3), e64. <https://doi.org/10.2196/jmir.1775>
- Lim, M. S. C., Molenaar, A., Brennan, L., Reid, M., & McCaffrey, T. (2022). Young adults' use of different social media platforms for health information: Insights from web-based conversations. *Journal of Medical Internet Research*, 24(1), Article e23656. <https://doi.org/10.2196/23656>
- Mansfield, K. L., Ghai, S., Hakman, T., & Ballou, N. (2025). From social media to artificial intelligence: Improving research on digital harms in youth. *The Lancet Child & Adolescent Health*. [https://doi.org/10.1016/S2352-4642\(25\)00012-3](https://doi.org/10.1016/S2352-4642(25)00012-3)
- Marchant, G. E., Allenby, B. R., & Herkert, J. R. (Eds.). (2011). *The growing gap between emerging technologies and legal-ethical oversight: The pacing problem*. Springer.
- McKinsey & Company. (2024a). *Meet Gen Z: Social media and digital tools are key in healthcare*. McKinsey Healthcare Systems & Services. <https://www.mckinsey.com/industries/healthcare/our-insights/meet-gen-z-social-media-and-digital-tools-are-key-in-healthcare>
- McKinsey & Company. (2024b). *Harnessing AI to reshape consumer experiences in healthcare*. McKinsey Global Institute.
- Mohamed, J. M. M., Ahmad, F., & Venkatesan, K. (2025). The role of artificial intelligence becoming a business success in pharma marketing and pharma universities. In *International Conference on Sustainability and Digitalization*. Springer.
- Nogueira, R., & Cho, K. (2019). Passage re-ranking with BERT. *arXiv preprint arXiv:1901.04085*. <https://doi.org/10.48550/arXiv.1901.04085>
- Noy, N., Gao, Y., Jain, A., Narayanan, A., Patterson, A., & Taylor, J. (2019). Industry-scale knowledge graphs: Lessons and challenges. *Communications of the ACM*, 62(8), 54–62. <https://doi.org/10.1145/3331166>
- OECD. (2023). *Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449)*. Organisation for Economic Co-operation and Development. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- Pervaiz, K., & Bawa, S. S. (2024). The impact of artificial intelligence-driven digital marketing strategies on pharmaceutical consumer behavior. In *International Conference on Optimization and Data Science*. Springer.

Redondo, I., & Aznar, G. (2018). To use or not to use ad blockers? The role of consumers' love of gadgets, privacy and avoidance of advertising. *Telematics and Informatics*, 35(6), 1607–1617. <https://doi.org/10.1016/j.tele.2018.04.006>

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). (2016). *Official Journal of the European Union*, L 119, 1–88.

Rodríguez-Hernández, M. D., Gil-Izquierdo, Á., & García, C. J. (2024). Health claims for sports drinks—Analytical assessment according to European Food Safety Authority's scientific opinion. *Nutrients*, 16(4), 512. <https://doi.org/10.3390/nu16040512>

Sabapathy, G. (2025). Artificial intelligence in pharmaceutical commercial operations: Transforming sales performance and patient engagement in the digital era. *Journal of Computer Science and Technology Studies*, 7(8), 636–651. <https://doi.org/10.32996/jcsts.2025.7.8.74>

Sorrell v. IMS Health Inc., 564 U.S. 552 (2011).

Stein, J.-P., & Ohler, P. (2017). Venturing into the uncanny valley of mind: The influence of mind attribution on the acceptance of human-like characters in a virtual reality setting. *Cognition*, 160, 43–50. <https://doi.org/10.1016/j.cognition.2016.12.010>

Tickle-Degnen, L., Zebrowitz, L. A., & Ma, H.-I. (2014). Culture, gender and health care stigma: Practitioners' response to facial masking in Parkinson's disease. *Social Science & Medicine*, 73(1), 95–102. <https://doi.org/10.1016/j.socscimed.2011.05.009>

Unilever. (2024). *Dove reaffirms commitment to 'real beauty' with pledge never to use AI to represent women* [Corporate press release]. Unilever PLC. <https://www.unilever.com/news/press-releases/2024/dove-reaffirms-commitment-to-real-beauty/>

U.S. Code of Federal Regulations. (2023). Postmarketing reporting of adverse drug experiences, 21 C.F.R. § 314.80.

U.S. Food and Drug Administration. (2022). *Dietary supplements: Structure/function claims*. U.S. Department of Health and Human Services. <https://www.fda.gov/food/dietary-supplements>

U.S. Food and Drug Administration. (2023). Prescription drug advertising: General regulations (21 CFR § 202.1). U.S. Code of Federal Regulations.

Verma, S., Tiwari, R. K., & Singh, L. (2024). Integrating technology and trust: Trailblazing role of AI in reframing pharmaceutical digital outreach. *Intelligent Pharmacy*. <https://doi.org/10.1016/j.ipha.2024.01.005>

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12. <https://doi.org/10.1145/3298981>

Zheng, L., Chiang, W. L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., & Gonzalez, J. E. (2023). Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. *Advances in Neural Information Processing Systems (NeurIPS)*, 36, 46595–46623.
<https://doi.org/10.48550/arXiv.2306.05685>

Zillul, M., & Shoeb, M. (2025). Digital marketing and artificial intelligence in healthcare: Revolutionizing patient engagement and service delivery. *Journal of Primeasia*.