Dipartimento di Scienze Economiche e Aziendali

Corso di Laurea magistrale in

ECONOMICS, DEVELOPMENT AND INNOVATION

Artificial Intelligence as a Hybrid Schumpeterian Technological Regime: Entry, Concentration and Industrial Dynamics

Relatore:

Chiar.mo Prof. *Roberto Fontana*

Tesi di Laurea

di

Simone Montagna

Matr. n. **523389**

Anno Accademico 2025-2026

Abstract in Italian

Questa tesi analizza l'intelligenza artificiale come un regime tecnologico schumpeteriano ibrido, concentrandosi sulla relazione tra ingresso delle imprese, concentrazione innovativa e dinamiche industriali. Basandosi sulla letteratura sui regimi tecnologici e sui modelli schumpeteriani dell'innovazione, lo studio interpreta l'IA come un dominio tecnologico caratterizzato dalla coesistenza di dinamiche di tipo Mark I e Mark II.

Da un lato, l'intelligenza artificiale presenta ampie opportunità tecnologiche, applicabilità trasversale ai settori e condizioni di ingresso relativamente aperte, soprattutto nei segmenti a valle e orientati alle applicazioni. Dall'altro lato, l'innovazione nei segmenti centrali e infrastrutturali resta influenzata da una forte cumulatività, da un'elevata appropriabilità e dal controllo di asset complementari quali dati, infrastrutture computazionali, piattaforme e capacità avanzate di ricerca.

L'analisi empirica si basa su un dataset brevettuale relativo alle invenzioni connesse all'IA nel periodo 2006–2024. Il dataset comprende 19.459 brevetti legati all'IA, 95.172 osservazioni brevetto–classe tecnologica e 7.019 richiedenti distinti. L'analisi esamina l'evoluzione temporale dell'attività brevettuale nell'IA, la composizione tecnologica dei brevetti attraverso le classificazioni IPC, la distribuzione dell'attività inventiva tra i richiedenti, i rapporti di concentrazione, l'indice di Herfindahl-Hirschmann e le dinamiche di ingresso.

I risultati mostrano una forte espansione dell'attività brevettuale nell'IA, in particolare a partire dalla metà degli anni 2010, indicando la presenza di ampie opportunità tecnologiche. Allo stesso tempo, la distribuzione dell'attività inventiva risulta fortemente asimmetrica: un ampio numero di richiedenti partecipa all'innovazione nell'IA, ma un gruppo ristretto di imprese leader rappresenta una quota sproporzionata dell'output brevettuale.

Questa evidenza è coerente con l'interpretazione dell'IA come regime tecnologico ibrido, nel quale ampia partecipazione e sperimentazione coesistono con vantaggi cumulativi, leadership persistente e concentrazione selettiva. La tesi sostiene quindi che l'IA non debba essere intesa semplicemente come una tecnologia democratizzante, ma come un sistema tecnologico stratificato, caratterizzato da un'ampia periferia di nuovi entranti e da un nucleo concentrato di innovatori dominanti.

Abstract

This thesis analyses artificial intelligence as a hybrid Schumpeterian technological regime, focusing on the relationship between firm entry, innovative concentration, and industrial dynamics. Drawing on the literature on technological regimes and Schumpeterian models of innovation, the study interprets AI as a technological domain characterised by the coexistence of Mark I and Mark II dynamics.

On the one hand, artificial intelligence displays broad technological opportunities, cross-sectoral applicability, and relatively open entry conditions, especially in downstream and application-oriented segments. On the other hand, innovation in core and infrastructural segments remains shaped by strong cumulativeness, high appropriability, and control over complementary assets such as data, computational infrastructure, platforms, and advanced research capabilities.

The empirical analysis is based on a patent dataset concerning AI-related inventions over the period 2006–2024. The dataset includes 19,459 AI-related patents, 95,172 patent–technology class observations, and 7,019 distinct applicants. The analysis examines the temporal evolution of patenting activity in AI, the technological composition of patents through IPC classifications, the distribution of inventive activity across applicants, concentration ratios, the Herfindahl-Hirschman Index, and entry dynamics.

The results show a strong expansion of AI patenting activity, particularly from the mid-2010s onwards, indicating the presence of broad technological opportunities. At the same time, the distribution of inventive activity appears highly asymmetric: a large number of applicants participate in AI innovation, but a restricted group of leading firms accounts for a disproportionate share of patenting output.

This evidence is consistent with the interpretation of AI as a hybrid technological regime, in which broad participation and experimentation coexist with cumulative advantages, persistent leadership, and selective concentration. The thesis therefore argues that AI should not be understood simply as a democratizing technology, but rather as a stratified technological system, characterized by a broad periphery of new entrants and a concentrated core of dominant innovators.

List of Figures

List of Tables

Introduction

Artificial intelligence has become one of the most important technology of recent decades. Its relevance does not depend only on the rapid diffusion of new applications, but also on its capacity to reshape the organization of innovation, the distribution of technological capabilities and the structure of competition across industries. AI technologies are now used in a wide range of economic activities, including digital services, finance, manufacturing, healthcare, logistics, communication, and scientific research. This broad applicability has led many scholars to interpret artificial intelligence as a general-purpose technology, capable of generating effects across several sectors rather than remaining confined to a single industry.

However, the economic implications of artificial intelligence cannot be fully understood by focusing only on adoption, productivity or labor-market effects. A more structural question concerns the way innovation is organized within the AI technological domain. In particular, it is necessary to understand whether the expansion of AI-related innovations has produced a more open and decentralized technological environment, or whether it has reinforced concentration among a limited number of leading firms. This issue is especially relevant because artificial intelligence combines features that point in different directions. On the one hand, the diffusion of open-source libraries, pre-trained models, cloud services and application programming interfaces has lowered some barriers to entry, allowing many firms and organizations to experiment with AI-based solutions. On the other hand, frontier AI development depends heavily on data, computational infrastructure, advanced research capabilities and complementary assets, which are unevenly distributed across firms.

This tension makes artificial intelligence an interesting case for applying the technological regime framework. In the Schumpeterian and evolutionary tradition, the structure of innovation is not treated as random or purely institutional. Instead, it is shaped by the economic properties of technologies. Technological opportunities, appropriability conditions, cumulativeness, and the nature of the knowledge base influence who can innovate, how firms enter a technological domain, whether innovative leadership persists over time, and whether inventive activity becomes more concentrated or more dispersed. From this perspective, the analysis of artificial intelligence requires more than a description of its growth. It requires an assessment of the regime through which AI-related innovation is produced, accumulated and distributed.

The central argument of this thesis is that artificial intelligence should be interpreted as a hybrid Schumpeterian technological regime. It displays features associated with Schumpeter Mark I dynamics, such as entry, experimentation and the widening of the innovative base. At the same time, it also exhibits features associated with Schumpeter Mark II dynamics, such as cumulativeness, persistence and concentration among leading innovators. The coexistence of these two dimensions suggests that AI is neither a fully open nor a fully closed technological domain. Rather, it appears to combine broad participation with selective and concentrated innovative leadership.

The thesis investigates this argument empirically through a patent-based analysis of AI-related inventive activity over the period 2006–2024. The dataset includes 19,459 AI-related patents, 95,172 patent–IPC observations and 7,019 distinct applicants. Patent data are used as a proxy for inventive activity because they make it possible to observe technological trajectories over time, identify the applicants involved in innovation, and measure the distribution of inventive output across actors and technological classes. Although patents do not capture all forms of innovation, they provide a systematic empirical basis for studying technological regimes and innovation patterns.

The empirical analysis focuses on three main dimensions. First, it examines the evolution and technological composition of AI-related patenting, in order to assess whether the field is characterized by high technological opportunities and a complex knowledge base. Second, it analyses the distribution of patenting activity across applicants, using concentration ratios, applicant distributions and the Herfindahl-Hirschmann Index to evaluate whether inventive output is broadly distributed or concentrated among leading actors. Third, it investigates entry dynamics and technological portfolio strategies, distinguishing between new and continuing applicants and examining the IPC portfolios of the top applicants. This allows the thesis to evaluate not only whether many actors enter the field, but also whether they persist and how leading firms organize their technological capabilities.

The research is guided by six hypotheses derived from the technological regime framework. These hypotheses concern the asymmetric nature of AI opportunities, the role of cumulativeness in core technological segments, the importance of complementary assets for appropriability, the selective nature of entry, the possible emergence of a core–periphery structure, and the strengthening of deepening dynamics over time. The purpose is not to claim that AI inevitably leads to concentration, nor that it automatically democratizes innovation. Rather, the thesis evaluates whether the empirical evidence supports a more nuanced interpretation, in which entry and concentration coexist within the same technological regime.

The structure of the thesis is as follows. Chapter 1 develops the theoretical framework, summarizing the literature on Schumpeterian patterns of innovation, technological regimes, firm entry and industrial concentration. It also discusses the relevance of digital and ICT technologies as cases where high opportunities may coexist with persistent concentration. Chapter 2 applies this framework to artificial intelligence, discussing AI as a general-purpose technology and examining its regime characteristics in terms of opportunities, cumulativeness, appropriability and knowledge base. Chapter 3 presents the dataset, the empirical strategy and the descriptive evidence on AI-related patenting dynamics. Chapter 4 develops the core empirical analysis, focusing on concentration, entry dynamics and the technological portfolios of leading applicants. Finally, Chapter 5 discusses the results in relation to the research questions and hypotheses, summarizes the main findings, and presents the theoretical implications, limitations and possible directions for future research.

Overall, the thesis aims to contribute to the literature on technological regimes and industrial dynamics by providing an empirical interpretation of artificial intelligence as a hybrid regime. The main contribution is to show that the expansion of AI-related patenting does not necessarily imply a simple democratization of innovation. Many actors enter and participate in the technological domain, but innovative output, persistence and technological breadth remain unevenly distributed. In this sense, AI patenting appears to be open in terms of participation, but selective in terms of leadership and cumulative innovative capacity.

Chapter 1 – Technological Regimes, Industrial Structure and Artificial Intelligence: Conceptual Framework and State of the Art

1.1 Motivation and Research Contribution

The dynamics of firm entry and the evolution of industrial concentration lie at the core of market structure analysis.

Entry and concentration shape competition, influence the distribution of market power, and ultimately affect an economy's ability to sustain long-run growth and structural transformation. Within the Schumpeterian and evolutionary tradition, those phenomena are not interpreted as random outcomes or purely institutional results, but rather as endogenous consequences of the innovation and learning processes characterizing specific technologies. From this perspective, the economic properties of technology — how knowledge is generated, accumulated, and appropriated — directly influence the distribution of innovative activity across firms, the probability of entry of new actors, and the level of industry concentration.

Research on AI has expanded rapidly, focusing on its effects on productivity, work organization, skill demand, and long-term growth. Many studies interpret AI as a force reshaping production processes, automating cognitive tasks and affecting labor demand, skill formation and the distribution of economic gains.

Yet, despite the breadth of this debate, AI has rarely been analyzed explicitly as a technological regime. In particular, there is limited analysis linking its economic characteristics to firm entry patterns and the evolution of industrial concentration. Whether and to what extent the economic features of AI influence firm entry and industrial concentration remains an open question. Most attention has focused on adoption effects, while systematic investigation of the underlying technological regime has received less attention.

This gap is especially relevant given AI's specific characteristics. AI shares several traits with digital technologies and General Purpose Technologies: pervasiveness, cross-sector applicability, complementarity with corporate and sectoral innovations, and continuous incremental improvements. At the same time, it has distinctive elements, including the central role of data as a productive input, the importance of computational scale economies, strong cumulativeness from algorithmic learning, and the need for complex technological infrastructures.

These features raise important questions from the perspective of technological regimes. The control of large proprietary datasets and cutting-edge computational capacity may generate dynamic entry barriers, reinforcing incumbent advantages. At the same time, the diffusion of open-source tools, pre-trained models, and cloud services may lower certain technological thresholds, facilitating the entry of new firms in vertical applications. AI therefore combines elements of openness and concentration, making it difficult to classify within a single regime.

In view of these considerations, this chapter pursues two main objectives. First, it reconstructs the conceptual framework developed in the literature on innovation patterns and technological regimes, showing how technological characteristics shape firm entry dynamics as well as industrial structure. Second, it applies this system to artificial intelligence, assessing whether it can be classified as a Mark I, Mark II, or hybrid technological regime, and considering the implications for firm entry and industrial concentration.

The chapter is structured as follows. The next sections trace the origins of the debate on Schumpeterian innovation patterns, distinguishing between Mark I and Mark II configurations. They introduce the concept of technological regime, define its main dimensions, and discuss the link among technological regimes, firm entry, and industrial concentration. The subsequent sections then apply this framework to digital technologies, in particular to artificial intelligence, to identify the remaining unresolved theoretical tensions. The chapter concludes through formulating the research questions that will guide the remainder of the thesis.

1.2 Technological Regimes and Industrial Configurations

The innovation literature investigating patterns of technological change draws on the work of Schumpeter, who offered two intertwined views on the role of innovation inside the capitalist system; these perspectives complement each other. In *The Theory of Economic Development* (1934), innovation stems from the initiative of the entrepreneur and from the entry of new firms: capable of introducing original combinations, they disrupt the prevailing equilibrium and trigger mechanisms of creative destruction. Economic change, from this perspective, is powered by the entry of new actors, previously unseen products, alternative production approaches, and untested organizational arrangements; what follows is constant turbulence, a continuous reshaping of the industrial structure.

In *Capitalism, Socialism and Democracy* (1942), Schumpeter partially revises this view. Observing twentieth-century capitalism, he notes that innovation increasingly concentrates within large firms, endowed with administrative structures and financial resources capable of sustaining systematic research and development. Competitive advantage derives from the ability to accumulate competencies and invest in the long term, thereby generating processes of creative accumulation and a higher concentration of innovative development.

Subsequent studies codified these two configurations into the paradigms known as Schumpeter Mark I and Schumpeter Mark II. The former portrays environments in which innovation is widely dispersed across many firms, distinguished by frequent entry, unstable technological hierarchies, and high industrial turbulence. The latter captures contexts in which a small number of large firms lead innovation, strengthening their position through the accumulation of expertise and assets.

Moreover, these studies consolidate the conceptual framework underlying the Mark I and Mark II classification, heightening its conceptual rigor and interpretative consistency *(Malerba and Orsenigo, 1995; Malerba and Orsenigo, 1996).*

These two views are commonly summarized as Schumpeter Mark I and Mark II.

Accumulated competencies, available financial resources, and organizational routines reinforce incumbent firms' positions and make entry by new competitors more difficult. Later studies, particularly the empirical studies by Malerba and Orsenigo, operationalized this distinction by translating it into two observable dynamics: widening and deepening. Widening is associated with Mark I and describes situations in which innovative efforts spread through the entry of new innovators that continuously reshuffle technological hierarchies. Deepening, linked to Mark II, concentrates inventive activity within a restricted group of firms that consolidate their market advantages over time.

An important finding emerges from these analyses. Such arrangements do not appear randomly across countries or sectors. Rather, they follow the nature of the underlying technology. Patent data analyses show a revealing pattern: holding technology constant innovation styles appear similar across countries; when technology changes, systematically different configurations arise.

It follows that the intrinsic characteristics of technology shape the organization of innovative activity and often prove more decisive than institutional or national structures considered in isolation.

To explain these differences, the evolutionary literature introduces the concept of technological regime: a set of economic circumstances and knowledge characteristics that jointly define how innovation develops within a specific technological domain. The particular combination of technological opportunities, appropriability conditions, cumulativeness of knowledge, and properties of the knowledge base defines technological regimes.

In the framework proposed by Malerba and Orsenigo, the regime is articulated into four dimensions: technological opportunities, appropriability conditions, the cumulativeness of the innovative process, and the characteristics of the knowledge base. These dimensions are not isolated elements but interdependent components: it is their combination—rather than each taken separately—that determines entry probabilities, the distribution of innovation across firms, and the evolution of industrial concentration.

Technological opportunities concern the "richness" of the innovative space: how many innovation pathways exist, how easy it is to generate new ideas, and whether technological progress depends on rare and costly discoveries or, instead, on incremental combinations and widespread experimentation. In high-opportunity technologies, the scope for innovation expands and the relative cost of exploring new routes may decrease, making the role of small firms and new entrants more plausible. However, high opportunities do not automatically imply open and unconcentrated outcomes *(Breschi, Malerba and Orsenigo, 2000)*: Even in the presence of abundant technological opportunities, factors such as high fixed costs, critical complementarities, network effects, or infrastructural requirements may push industries toward more concentrated structures; for this reason, opportunities must be interpreted in conjunction with the other dimensions of the technological regime.

The interaction between technological opportunities and structural constraints is therefore crucial for understanding how innovation patterns translate into specific industrial structures.

Appropriability concerns firms' ability to protect and capture economic returns from innovation. It relies on both formal instruments (patents and intellectual property rights) and informal mechanisms (trade secrecy, technological complexity, lead-time advantages, control of complementary assets, and organizational capabilities that are not easy to copy). In contexts of high appropriability, innovation generates more defensible and durable rents; this reduces market contestability, strengthens incumbents' positions, and makes it more difficult for new entrants to acquire market shares and technological leadership. In contexts of low appropriability, by contrast, innovation is more easily imitated, competitive advantage is shorter-lived, and the competitive space endures more open, encouraging entry and turbulence.

Empirical research further shows that appropriability and patenting do not coincide, patent propensity varies substantially across sectors, and many firms protect innovation through combinations of strategies rather than relying exclusively on patents (*Fontana et al., 2013*). The theoretical point is that the capacity to appropriate the returns from innovation is central in determining the persistence of leadership and the form of competition (*Breschi, Malerba, and Orsenigo, 2000*; *Malerba and Orsenigo, 1997*).

Cumulativeness indicates the extent to which current innovation depends on past innovation and on the accumulation of specific knowledge, competencies, and routines. In highly cumulative regimes, firms that have innovated in the past tend to enjoy an advantage in continuing to innovate: experience reduces the marginal costs of innovation, increases the efficiency of research processes, and raises the likelihood of generating further innovations. This creates mechanisms of path dependence and supports the persistence of technological leadership (*Malerba and Orsenigo, 1996*; *Breschi, Malerba, and Orsenigo, 2000*).

Within regimes with low cumulativeness, by contrast, experience matters less; entry is more feasible, and the hierarchy of innovators tends to be more unstable. Moreover, cumulativeness can take different forms: it may be internal to the firm (learning and capabilities), but it may also depend on infrastructures, standards, and architectures that increase the advantage of those controlling key nodes of the system—an aspect especially visible in digital technologies (*Malerba and Orsenigo, 1997*; *Peneder, 2010*).

The knowledge base concerns the nature of the knowledge used to innovate: its degree of codifiability or tacitness, its complexity, its dependence on science and frontier research, and the requirement for interdisciplinary competencies. A strongly science-based and complex knowledge base is inclined to favor firms endowed with advanced research capabilities, access to highly skilled talent, and strong links with universities and research centers (*Malerba, 2005*). A more modular, codifiable, or user-driven knowledge base may instead facilitate broader innovation diffusion and open spaces for new entrants (*Peneder, 2010*). Even in this case, however, the knowledge base does not determine a single outcome: its effect depends on its interaction with opportunities, appropriability, and cumulativeness (*Breschi, Malerba, and Orsenigo, 2000*).

This conceptual architecture underpins the literature linking technological regimes, firm entry, and transformations in industrial concentration. A crucial step is recognizing that entry is neither a neutral phenomenon nor reducible to a mere count of new firms. Entry is a selective process: many firms enter, few survive, and even fewer grow and manage to influence major technological trajectories (*Malerba and Orsenigo, 1996*).

Breschi, Malerba, and Orsenigo emphasize that the combination of opportunities, appropriability, cumulativeness, and knowledge base determines both the probability of entry and the quality and impact of entry (*Breschi, Malerba, and Orsenigo, 2000*). In regimes characterized by high opportunity and low appropriability, entry may be frequent and contribute to renewal; in regimes with high appropriability and high cumulativeness, the accumulated experience of incumbents becomes a dynamic barrier that limits entrants' ability to emerge and reshape market structure. Analyzing entry, therefore, requires distinguishing between numerous but fragile entry and transformative entry capable of generating new leaders.

The same logic applies to industrial concentration, which is not a fixed outcome or an immutable property. Concentration evolves through the accumulation of capabilities, competitive selection, and the persistence of innovators. In highly cumulative regimes, innovating today increases the probability of innovating tomorrow: this self-reinforcing mechanism tends to produce greater stability and concentration (*Malerba and Orsenigo, 1996*). Appropriability further strengthens incumbents' ability to maintain rents, discourage entry, and consolidate positions (*Breschi, Malerba, and Orsenigo, 2000*).

From this perspective, innovation and concentration are not necessarily in tension. In many technological regimes, innovation may coexist with high concentration or even reinforce it. The relationship between innovation and market structure is therefore shaped by the properties of technology and the strategies derived from them, rather than by a universal trade-off.

A particularly relevant empirical finding is the coexistence of high entry rates and high levels of concentration. This evidence challenges the simple idea that more entry automatically leads to lower concentration. In high-technology sectors, particularly in ICT, entry may be substantial but confined to niches or peripheral segments, while the technological frontier and core segments remain under incumbent control. Corrocher, Malerba, and Montobbio show this outcome across several ICT classes: high opportunities generate experimentation and entry, but do not prevent a strong concentration of innovative activity in core segments (*Corrocher, Malerba, and Montobbio, 2007*). The apparent paradox is resolved by distinguishing between the quantity and quality of entry and by recognizing that sectoral structure may be stratified. Some portions of the value chain are relatively open, while others present stronger cumulative and appropriative barriers.

In recent years, this framework has been further enriched by contributions emphasizing intra-sectoral heterogeneity and selection processes. Peneder highlights that even within the same sector, different innovative behaviors may coexist: science-based firms, production-oriented firms, user-driven firms, highly patent-intensive firms, and firms that innovate without patenting (*Peneder, 2010*). Srholec and Verspagen show that a substantial share of the variance in innovative behavior lies within sectors and that groups of firms with similar strategies may cross traditional sectoral boundaries (*Srholec and Verspagen, 2012*). These results do not weaken the explanatory power of technological regimes. Rather, they highlight their evolutionary nature: regimes do not determine fixed outcomes, but define a framework of constraints and opportunities within which micro-level variety and competitive selection generate the observed industrial structure. Market outcomes emerge from the interaction between technological conditions, firm heterogeneity, and selection mechanisms.

1.3 Technological Regimes in ICT: High Opportunities and Persistent Concentration

This theoretical framework becomes particularly salient when applied to digital technologies, especially Information and Communication Technologies (ICT).

ICTs exemplify a domain in which the four dimensions of the technological regime interact in complex ways, often defying a strict Mark I/Mark II dichotomy. ICTs exhibit high technological opportunities, rapid application development, and a knowledge base characterized by modularity and combinatorial potential. Simultaneously, elements such as infrastructures, standards, core architectures, and platforms introduce strong forms of cumulativeness and appropriability that reinforce industrial concentration.

This configuration challenges the sufficiency of the Mark I/Mark II dichotomy. It suggests the need for a more articulated reading of digital technological regimes: one that combines widening dynamics (entry, experimentation, turbulence) with deepening dynamics (persistence, control over core trajectories, concentration) (*Corrocher, Malerba and Montobbio, 2007*; *Breschi, Malerba and Orsenigo, 2000*).

A central interpretative element concerns the systemic nature of digital innovation. The sectoral systems of innovation approach stress that innovation does not occur in isolation, but through networks of actors, institutions, infrastructures, and learning mechanisms (*Malerba, 2005*).

In ICTs, this dimension is particularly pronounced: hardware and software co-evolve; infrastructures and applications are complementary; standards and protocols coordinate ecosystems. In this context, control over strategic nodes of the ecosystem—data, platforms, standards, computational infrastructures—becomes a crucial factor in shaping entry and concentration dynamics (*Malerba, 2005*; *Peneder, 2010*).

Consequently, appropriability does not coincide with the protection of a single invention. Still, it often depends on control over ecosystems and "bottlenecks" in the architecture: interfaces, APIs, compatibility layers, marketplaces, multi-sided platforms, data, and distribution channels. In digital settings, capturing innovation rents is closely linked to a firm's position within the overall architecture and its ability to orchestrate complementarities (*Peneder, 2010*).

A clear example of architectural bottlenecks in digital ecosystems is provided by mobile application platforms. The App Store of Apple and the Play Store of Google act as gateways between developers and users. Since access to consumers necessarily passes through these platforms, value appropriation depends not only on innovation itself, but on control over the infrastructural layer that coordinates the ecosystem. Firms that govern these strategic nodes are therefore able to capture a substantial share of innovation rents.

A second element concerns cumulativeness, which in ICTs can take both an architectural and an organizational form. In traditional manufacturing sectors, cumulativeness mainly reflects the internal accumulation of capabilities, such as experience, routines, and efficiency improvements. In digital technologies, cumulativeness is often linked to control over standards and platforms. Each improvement of the system reinforces the position of the central actor, creating lock-in effects and increasing returns. As a result, while entry into the market may appear relatively open, access to the core segments of the ecosystem remains limited *(Corrocher, Malerba, and Montobbio, 2007; Breschi, Malerba, and Orsenigo, 2000).*

This framework explains why ICT sectors often display high entry rates without a corresponding de-concentration. Entry is concentrated in application or complementary segments, while control over the technological frontier and infrastructures remains in the hands of a few actors. The outcome is a stratified structure: a concentrated core and a turbulent periphery. This configuration is not anomalous, but rather the result of the combination of high opportunities with strong cumulative and appropriative mechanisms in specific segments (*Corrocher, Malerba, and Montobbio, 2007*).

It is precisely this evidence that weakens the automatic link between entry and de-concentration and calls for a systematic distinction between numerous entry and transformative entry (*Malerba and Orsenigo, 1996*).

ICTs are also frequently interpreted as General Purpose Technologies (GPTs). A GPT is characterized by pervasiveness, cross-sectoral applicability, complementarities with organizational and sectoral innovations, continuous incremental improvements, and often lags in productivity effects. This reading is important because it implies that industrial dynamics related to ICTs do not end within the ICT sector narrowly defined: the technology diffuses downstream, reshapes production processes, and generates indirect effects on entry and concentration in user industries. Pervasiveness further complicates the analysis, as innovation becomes a transversal force that affects structures and competition across multiple domains.

Moreover, the GPT nature amplifies internal heterogeneity: within the same sector, firms with very different roles and strategies may coexist, as shown by *Peneder (2010)* and *Srholec and Verspagen (2012)*. This reinforces the idea that technological regimes should not be interpreted as "sectoral averages" that erase micro-level variety, but as contexts that organize and structure heterogeneous distributions of behavior (*Srholec and Verspagen, 2012*; *Peneder, 2010*).

These dynamics provide a template for analyzing AI, where similar tensions may be amplified. In light of this theoretical framework, the hypotheses guiding the thesis can be formulated.

H1 – Asymmetric opportunities: AI technological opportunities are high in aggregate terms, but their effective contestability is asymmetric, varying across infrastructural and application segments.

H2 – Reinforced cumulativeness in core segments: In frontier and infrastructural AI segments, the innovation process is highly cumulative, generating path-dependent mechanisms that increase the likelihood of persistent technological leadership (*Malerba and Orsenigo, 1996*).

H3 – Appropriability mediated by complementary assets: The capture of innovation rents in AI depends substantially on control over complementary assets such as data, infrastructures, platforms, and distribution channels, reinforcing the positions of actors already present in core segments (*Fontana et al., 2013*; *Peneder, 2010*).

H4 – Numerous but selective entry: AI generates high entry rates in application and downstream segments, but only a limited share of entrants succeeds in influencing central technological trajectories or significantly altering concentration in core segments (*Malerba and Orsenigo, 1996*).

H5 – Persistent core–periphery structure: The evolution of industrial concentration in AI tends to assume a core–periphery configuration, with high concentration and stability in the infrastructural core and greater competitive turbulence in the application periphery (*Corrocher, Malerba and Montobbio, 2007*).

H6 – Strengthening of deepening dynamics over time: In the medium term, the intensification of scale economies in computing, data accumulation, and organizational complementarities may reinforce creative accumulation dynamics in core segments, increasing relative concentration compared to the initial diffusion phase (*Breschi, Malerba and Orsenigo, 2000*; *Malerba and Orsenigo, 1996*).

Taken together, these hypotheses do not leadto the conclusion that AI should inevitably leads to greater concentration. Nor that it permanently democratizes innovation. Rather, they formulate conditional predictions. Structural outcomes depend on the distribution and evolution of the four regime dimensions along the value chain. Their verification requires a differentiated analysis of AI segments and separate measurement of entry, survival, growth, and concentration.

With the formulation of these hypotheses, the theoretical framework and state of the art of the introductory chapter are concluded. Artificial intelligence is interpreted here not merely as a general-purpose technology or a driver of productivity transformation, but as a potential new technological regime whose configuration directly shapes firm entry dynamics and the evolution of industrial structure. The subsequent chapters will empirically investigate these dynamics, to assess whether AI is durably expanding the innovation base or consolidating—through cumulative and appropriative mechanisms—the leadership of a restricted number of actors in the core segments of the technological system (*Breschi, Malerba and Orsenigo, 2000*; *Malerba and Orsenigo, 1996*).

Chapter 2 – AI as a Technological Regime

In the following section, the technological regimes framework is applied to the analysis of artificial intelligence with a clear objective: to understand what kind of technology AI actually is from an economic perspective. The aim is not simply to describe AI as "important" or "transformative," but to understand how innovation works within it, who is able to innovate, under what conditions, and with what implications for market structure and competition *(Malerba and Orsenigo, 1995; Breschi, Malerba and Orsenigo, 2000)*.

A useful starting point is whether artificial intelligence can be interpreted as a General-Purpose Technology. AI displays several characteristics commonly associated with this type of technology, since it is not confined to a single sector but can be applied across a wide range of industries, from healthcare and finance to manufacturing and digital services. Wherever data can be processed or decisions can be optimized, artificial intelligence may potentially be used, and this cross-sectoral applicability represents one of the defining features of general-purpose technologies *(Bresnahan and Trajtenberg, 1995; Cockburn, Henderson and Stern, 2018)*.

This aspect matters because general-purpose technologies tend to reshape not only specific industries but also the broader process of innovation. However, classifying AI as a GPT is not enough to fully understand its economic implications. A more complete interpretation requires analysing how innovation is organized within this technological domain (*Bresnahan and Trajtenberg, 1995; Cockburn, Henderson and Stern, 2018*).

The technological regimes framework is particularly useful for this purpose because it treats technologies as sets of economic and knowledge conditions that shape the organization of innovation. Some technological domains make entry easier and generate a more dispersed structure, while others favor established actors and lead to more concentrated outcomes. Understanding AI therefore requires identifying how technological opportunities, cumulativeness, appropriability and the knowledge base interact (*Malerba and Orsenigo, 1995; Breschi, Malerba and Orsenigo, 2000*).

The first dimension to consider is technological opportunities, meaning how easy it is to generate new ideas and innovations. Artificial intelligence appears to offer a wide range of opportunities. In recent years, new applications have emerged continuously: image recognition, natural language processing, recommendation systems, predictive analytics, and many others. This creates the impression of a highly dynamic and accessible environment, where innovation is frequent and experimentation is widespread (*Cockburn, Henderson and Stern, 2018; OECD, 2024*).

This interpretation is partly accurate. Today, a large number of tools are available at relatively low cost. Open-source libraries, pre-trained models, and cloud-based platforms allow firms and even individual developers to build AI applications without starting from scratch. This significantly lowers the barriers to entry, especially in application-level innovation (*Bommasani et al., 2021; OECD, 2024*).

However, this is only part of the picture. A closer look reveals that opportunities are not evenly distributed across the technological landscape. There are different layers of innovation. At the application level, entry is relatively easy. At the frontier level — where new models are developed, large-scale systems are trained, and core technologies are advanced — the situation is very different. These activities require massive computational resources, access to large proprietary datasets, and highly specialized expertise. As a result, the effective opportunity to innovate is much more limited in these segments (*Bommasani et al., 2021; OECD, 2024; AI Index Steering Committee, 2025*).

This leads directly to the second dimension: cumulativeness. Cumulativeness concerns the extent to which past innovative activity increases the probability of future innovation. In the case of AI, it matters a great deal. Firms that already possess large datasets, advanced infrastructure, and skilled research teams have a strong advantage. They are able to improve their technologies continuously and at a faster pace than others (*Breschi, Malerba and Orsenigo, 2000; OECD, 2024*).

This creates a clear dynamic: those who are ahead tend to stay ahead. The process is not random. It is driven by accumulation. More data leads to better models, better models attract more users, and more users generate more data. This self-reinforcing cycle makes it increasingly difficult for new entrants to compete at the same level. In this sense, innovation in AI is strongly path-dependent (*Breschi, Malerba and Orsenigo, 2000; OECD, 2024*).

The third dimension concerns appropriability, that is, the ability to capture economic returns from innovation. In artificial intelligence, this does not depend primarily on formal protection mechanisms such as patents. Instead, it is largely determined by control over complementary assets. Data is one of the most important of these assets. Access to large and high-quality datasets allows firms to train more effective models, creating a significant competitive advantage (*Teece, 1986; OECD, 2024*).

In addition, computational infrastructure and digital platforms play a crucial role. Firms that control cloud services, distribution channels, or user interfaces are in a much stronger position to scale their innovations and reach the market. In this context, competitive advantage is not just about inventing something new, but about controlling the environment in which that innovation is deployed. Appropriability, therefore, becomes systemic: it depends on a firm's position within a broader technological ecosystem (*Teece, 1986; OECD, 2024*).

The fourth dimension is the nature of the knowledge base. Artificial intelligence is a complex and highly specialized field, requiring advanced knowledge in areas such as mathematics, statistics, and computer science. At the same time, it is also partially accessible. The spread of open-source tools and standardized frameworks allows many actors to use AI technologies without fully understanding their underlying mechanisms (*Breschi, Malerba and Orsenigo, 2000; Bommasani et al., 2021*).

This combination creates a central tension. AI is relatively open at the level of use and application, but more selective at the level of frontier development. This dual structure is essential to understanding the organization of innovation in AI (*Bommasani et al., 2021; OECD, 2024*).

When these four dimensions are considered together, a more nuanced picture emerges. Artificial intelligence does not behave as a purely open system, nor as a fully closed one. Instead, it combines elements of both (*Malerba and Orsenigo, 1995; Breschi, Malerba and Orsenigo, 2000*).

On one side, the availability of tools and the breadth of applications encourage entry and experimentation. This resembles the Schumpeter Mark I model, where innovation is driven by new entrants and a dynamic competitive environment. On the other side, strong cumulative processes and control over key assets lead to concentration, with a limited number of firms dominating the most important technological developments. This is typical of the Schumpeter Mark II model, based on accumulation and persistence (*Malerba and Orsenigo, 1995; Breschi, Malerba and Orsenigo, 2000*).

These two dynamics therefore coexist rather than exclude each other, which is the basis for interpreting AI as a hybrid technological regime.

A useful way to interpret this structure is through the idea of a core–periphery system. At the core, we find large technology firms and leading research organizations that develop the most advanced models and infrastructures. Innovation in this segment is highly cumulative, resource-intensive, and concentrated. In the periphery, we observe a much larger number of actors developing applications based on existing technologies. Entry is easier, experimentation is more widespread, but the individual impact of each actor is more limited (*Corrocher, Malerba and Montobbio, 2007; OECD, 2024; AI Index Steering Committee, 2025*).

A more precise interpretation of this hybrid configuration can be obtained by distinguishing between different layers of the AI value chain. At the infrastructural level, innovation critically depends on access to computational capacity, specialized hardware, large-scale datasets, and advanced research capabilities. These requirements imply high fixed costs and significant scale economies, which tend to favor large incumbent firms and reinforce cumulative advantages over time (*OECD, 2024; AI Index Steering Committee, 2025*).

At the model level, these dynamics are further amplified by the emergence of foundation models. The development, training, and deployment of such models require substantial financial resources, technical expertise, and access to proprietary data, limiting participation to a relatively small number of actors operating at the technological frontier. As a result, innovation in these segments tends to be highly concentrated and strongly path-dependent (*Bommasani et al., 2021; OECD, 2024; AI Index Steering Committee, 2025*).

By contrast, at the application level, the conditions for innovation are markedly different. The widespread availability of open-source libraries, application programming interfaces (APIs), pre-trained models, and cloud-based services significantly lowers entry barriers, enabling a broader set of firms to develop AI-based products and services. In this segment, innovation is more decentralized, experimentation is more frequent, and entry is more widespread, although often limited to downstream and complementary activities (*Bommasani et al., 2021; OECD, 2024*).

This indicates that innovation and concentration are not opposed outcomes in AI, but two dimensions produced by the same layered technological structure.

While upstream and infrastructural segments are characterized by cumulative dynamics and the persistence of leading actors, downstream segments remain more open and dynamic, giving rise to a core–periphery configuration within the technological system (*Breschi, Malerba and Orsenigo, 2000; OECD, 2024*).

In conclusion, the technological regimes framework makes it possible to interpret AI not only as a fast-growing technology, but as a structured innovation system. Opportunities are broad but uneven, innovation is widespread but also concentrated, and entry is possible but not equally consequential for all actors (*Malerba and Orsenigo, 1995; Breschi, Malerba and Orsenigo, 2000*).

This interpretation is crucial because it shapes how we understand both the data and the evolution of market structures. It is not enough to look at how many firms are innovating; what matters is who is able to shape technological trajectories and maintain a leading position over time. The empirical analysis in the following chapters will test whether these patterns are reflected in the data and will provide a more precise assessment of how AI is transforming the structure of innovation.

Chapter 3 – Data and Empirical Strategy

3.1 Data Sources, Patent Sample Construction and Definition of AI Activities

The empirical analysis developed in this chapter aims to assess whether artificial intelligence exhibits the characteristics of a distinct technological regime, in line with the theoretical framework outlined in the previous chapters. In particular, the analysis focuses on patterns of innovative activity, the distribution of inventive efforts across actors, and the degree of concentration within the AI technological domain, with the objective of testing the hypotheses formulated with respect to technological opportunities, cumulativeness, appropriability, and the nature of the knowledge base.

To this end, the study relies on a patent-based dataset covering artificial intelligence inventions over the period 2006–2024. The use of patent data is consistent with the empirical tradition in the economics of innovation, where patents represent an observable and systematic proxy for inventive activity. They allow researchers to trace the evolution of technological trajectories over time, identify the key actors involved in innovation processes, and analyze the distribution of knowledge within specific technological domains. While acknowledging some limitations—such as the fact that not all innovations are patented and that patent propensity varies across sectors—patent data remain particularly suitable for the comparative analysis of technological dynamics and innovation regimes.

The dataset includes 19,459 AI-related patents, 95,172 patent–IPC observations and 7,019 distinct applicants over the period 2006–2024. The analysis uses patents to measure inventive output, patent–IPC observations to measure technological composition, and applicant mentions to measure the distribution of inventive activity across actors. Patent–IPC observations reflect the multidimensional nature of patents, which can be simultaneously classified into multiple technological categories according to the International Patent Classification system. On average, each patent is associated with approximately 4.89 IPC classes, highlighting the strongly combinatorial and interdisciplinary nature of the AI knowledge base. This aspect is particularly relevant within the technological regime framework, as it points to a high degree of variety in innovation trajectories and significant potential for recombination across different technological domains.

The identification of AI-related patents is based on the selection of specific IPC classes commonly associated with artificial intelligence technologies. These include, in particular, classes related to data processing (e.g. G06F), computational methods based on learning models (G06N), pattern and signal recognition (including G10L for speech and language processing), computer vision (G06V), and image and signal processing (G06T). The use of IPC classifications allows for a transparent and replicable definition of the technological domain under analysis, avoiding ambiguities associated with more qualitative or sector-based definitions of artificial intelligence.

From a geographical perspective, the dataset is highly concentrated in the United States, which accounts for the vast majority of the observed patenting activity (approximately 92.7% of the total). This concentration reflects both the leading role of the United States in the development of artificial intelligence technologies and the characteristics of the data source used. Although this distribution may limit the full generalizability of the results to a global context, it provides a coherent setting for analyzing innovation dynamics at the technological frontier.

At the applicant level, the dataset includes 7,019 distinct actors, encompassing large multinational corporations, technology startups, and public and private research organizations. The distribution of patenting activity across these actors is highly heterogeneous. Alongside a large number of actors with only sporadic participation, a significant share of innovative activity is concentrated in a relatively small group of leading firms, including major players in the digital technology sector. This pattern provides a first empirical indication consistent with a core–periphery structure, in which a restricted set of actors dominates the most advanced segments of the system, while a broader periphery is characterized by higher levels of entry and experimentation.

The construction of the dataset also required a data cleaning and harmonization phase aimed at ensuring the consistency of information related to applicants and technological classifications. In particular, it was necessary to standardize applicant names in order to avoid duplications due to variations in naming conventions, as well as to verify the consistency of IPC assignments. These steps are crucial for the analysis of concentration and the distribution of innovative activity, as classification errors could significantly distort the empirical measures employed.

Overall, the structure of the dataset allows for a systematic analysis of the key dimensions of technological regimes. The temporal evolution of patenting activity provides insights into technological opportunities; the distribution of applicants and the concentration of innovation allow for the assessment of appropriability mechanisms; the persistence of actors over time offers evidence on cumulativeness; and the structure of IPC classifications enables the exploration of the characteristics of the knowledge base. In this sense, the dataset constitutes a coherent empirical foundation for testing the hypotheses developed in the previous chapters.

While patent data provide a systematic proxy for inventive activity, they also present limitations. Not all innovations are patented, patent propensity varies across sectors and firms, and the most recent observations, especially those for 2024, may be affected by publication or reporting lags.

Moreover, patents capture invention rather than commercialization, and may over-represent large firms with structured IP strategies. These limitations should be kept in mind when interpreting the results.

3.2 Variables and Empirical Measures

The empirical analysis requires translating the theoretical dimensions of technological regimes into observable and measurable indicators. As discussed in Chapter 1, technological regimes are defined by four main dimensions: technological opportunities, cumulativeness, appropriability, and the nature of the knowledge base, whose joint configuration shapes the structure of innovation processes and market dynamics.

From this perspective, the objective is not to provide an isolated measurement of each dimension, but rather to construct a coherent set of indicators that allows their interaction to be captured within the domain of artificial intelligence.

Given the patent-based nature of the dataset, the variables employed are derived from observable inventive activity, its distribution across actors, and the underlying technological structure. While patents do not capture the entirety of innovation processes, they provide a systematic empirical basis for analyzing the evolution of technological trajectories and for attributing innovative activity to specific economic actors. The measures adopted should therefore be interpreted as proxies of the theoretical dimensions, rather than as direct observations.

Within this framework, *technological opportunities* are proxied through the dynamics of patenting activity over time. In particular, the number of AI-related patents per year is used as an indicator of the vitality and expansion of the technological space. A sustained increase in patenting activity is consistent with the presence of a context characterized by high innovative intensity. This measure is complemented by an indicator of the breadth of technological domains involved, captured through the diversity of IPC classes associated with patents. A high degree of heterogeneity in classifications signals that innovation extends across multiple domains, suggesting the presence of widespread opportunities and a strong capacity for technological recombination.

Cumulativeness is instead analyzed through the persistence of actors within the innovation process. Specifically, attention is paid to the frequency with which the same applicants appear over time among the main producers of patents. The repeated presence of the same actors is indicative of cumulative dynamics, in which accumulated experience and capabilities increase the likelihood of further innovation. This type of evidence is consistent with path-dependent processes, in which technological advantages tend to reinforce themselves over time. Additional insights can be obtained from the stability of the distribution of patenting activity: a relatively stable structure, with limited turnover among leading innovators, suggests a high degree of cumulativeness and the presence of dynamic barriers to entry.

Appropriability is proxied through the degree of concentration of innovative activity across actors. The analysis therefore focuses on the distribution of patents among applicants, with particular attention to the relative weight of the leading innovators. A strong concentration of patenting activity in a limited number of actors is consistent with high appropriability conditions, where dominant firms are able to capture a substantial share of innovation returns. This measure reflects not only the role of formal intellectual property protection, but also the influence of broader mechanisms, such as control over complementary assets, access to resources, and the ability to occupy strategic positions within the technological system.

Finally, the nature of the *knowledge base* is examined through the structure of IPC classifications associated with patents. In particular, the average number of IPC classes per patent is used as an indicator of technological complexity and combinatorial knowledge. A higher value suggests that innovation relies on the integration of knowledge from multiple domains, reflecting a complex and interdisciplinary technological structure.

In addition, the distribution of IPC classes provides insights into the composition of the technological domain, highlighting the main subfields in which innovative activity is concentrated. This configuration is consistent with the characteristics of artificial intelligence, which is inherently interdisciplinary and based on highly integrated innovation processes.

Overall, the variables constructed allow for a systematic empirical analysis of the main dimensions of the technological regime of artificial intelligence, in line with the theoretical framework. The dynamics of patenting activity provide evidence on technological opportunities; the distribution of applicants allows for the assessment of cumulativeness and appropriability; while the structure of IPC classifications offers insights into the nature of the knowledge base. The joint analysis of these dimensions makes it possible to assess whether artificial intelligence exhibits a configuration consistent with a hybrid regime, characterized by the coexistence of high innovative activity and concentration dynamics, in line with the hypotheses developed in the previous chapters.

3.3 Descriptive Evidence on AI Patenting Dynamics

This section presents the first descriptive evidence on the evolution and structure of AI-related patenting activity. The objective is to assess whether the empirical patterns observed in the patent data are consistent with the interpretation of artificial intelligence as a hybrid technological regime. The analysis focuses on three main dimensions: the temporal evolution of patenting activity, the technological composition of AI patents through IPC classifications, and the distribution of inventive activity across applicants.

The first dimension concerns the evolution of AI-related patents over time. As discussed in the previous section, the number of patents filed each year provides a first proxy for technological opportunities. A growing volume of patenting activity suggests that the technological domain is expanding and that firms perceive increasing opportunities for innovation. In the case of artificial intelligence, the dataset covers the period 2006–2024 and includes 19,459 AI-related patents. This allows the analysis to observe the development of AI patenting activity over a sufficiently long-time span, including both the earlier phase of diffusion and the more recent phase associated with the acceleration of data-driven and machine-learning-based technologies.

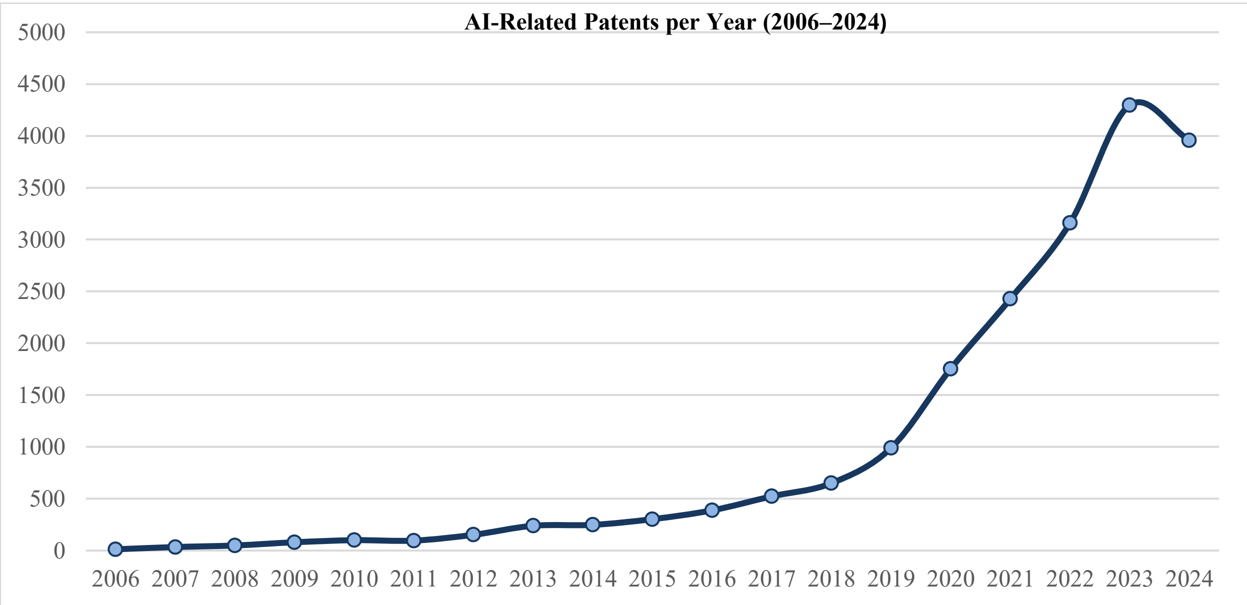

Figure 3.1. Number of AI-related patents by year (2006–2024)

As shown in Figure 3.1, the number of AI-related patents exhibits a strong and sustained increase over the period 2006–2024. While patenting activity remains relatively limited and stable during the early years, a clear acceleration emerges from the mid-2010s, with a particularly sharp increase after 2018. The decrease in 2024 should be interpreted with caution, since recent patent counts may still be affected by publication and reporting lags.

This pattern indicates a clear expansion of the technological space in artificial intelligence, with an increasing number of innovation trajectories being explored over time. The acceleration observed after the mid-2010s is consistent with the transition toward data-driven and machine-learning-based paradigms, which have significantly broadened the scope for innovation.

This acceleration is also consistent with the emergence of artificial intelligence as a general-purpose technology, whose diffusion across sectors tends to generate cumulative innovation waves rather than isolated technological advances.

However, as emphasized in the theoretical framework, high technological opportunities do not automatically imply an unconcentrated or fully open regime.

For this reason, the growth of AI patenting activity must be interpreted together with the distribution of inventive activity across technological classes and applicants.

The second dimension concerns the technological composition of AI-related inventions. The dataset includes 95,172 patent–IPC observations, indicating that AI patents are frequently associated with multiple IPC classes. On average, each patent is linked to approximately 4.89 IPC classes. This is an important empirical result because it suggests that AI innovation is not confined to a narrow technological field, but instead relies on the combination of several knowledge domains. From the perspective of technological regimes, this supports the idea that AI is characterized by a complex and combinatorial knowledge base.

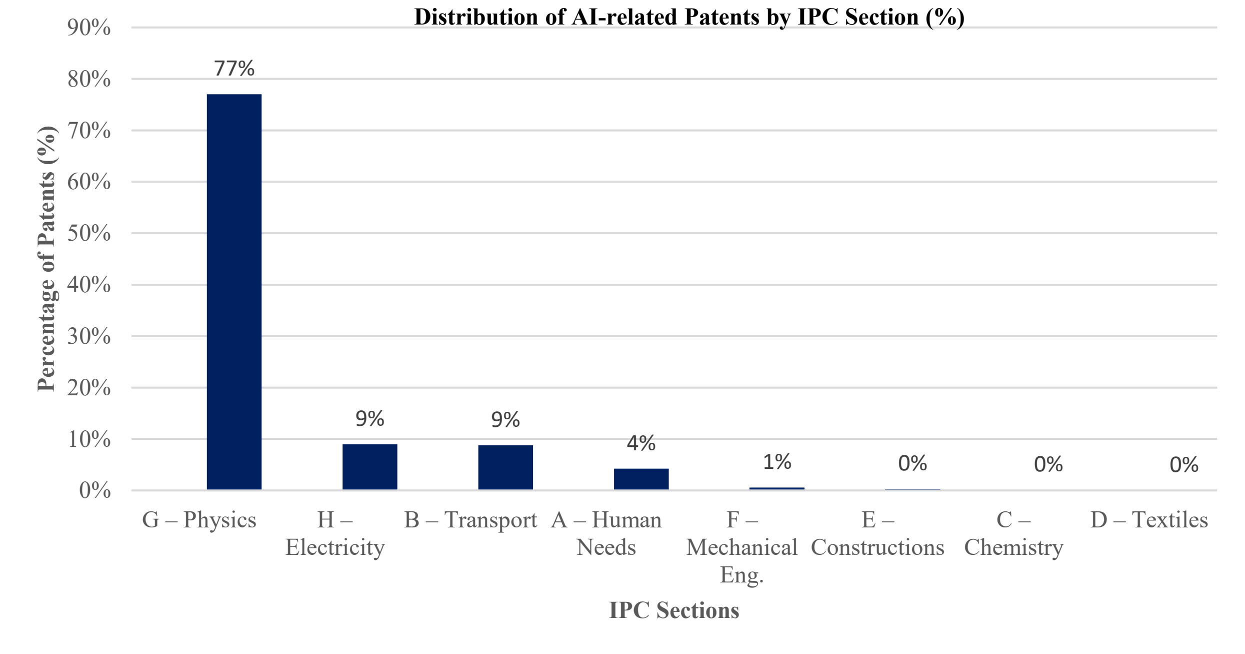

Figure 3.2. Distribution of AI-related patents by IPC section (percentage values)

A particularly relevant empirical result is the dominance of Section G, which alone accounts for approximately 77% of total patent–class observations. This indicates the strong concentration of AI innovation within digital and computational technologies, reinforcing the interpretation of AI as a technology deeply rooted in information processing and algorithmic systems.

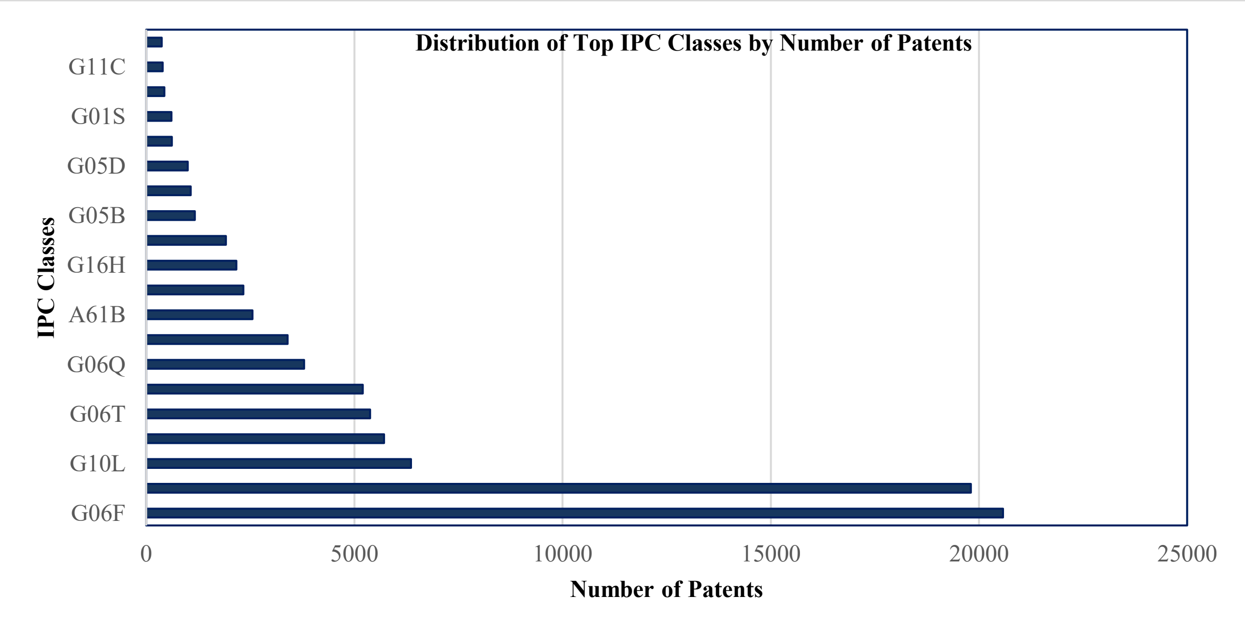

Figure 3.3. Distribution of Top IPC Classes by number of Patents

The distribution of patents across IPC classes shows the central role of digital and computational technologies in the AI domain. The most relevant classes include G06F, G10L, G06T, G06Q, A61B and G16H. These categories capture different technological areas, including data processing, computational models, machine learning, speech and language processing, computer vision, and image processing. The presence of these classes shows that AI does not correspond to a single technological field, but rather to a set of interconnected fields built around computation, data, pattern recognition, and automated decision-making.

This evidence is directly relevant to the knowledge-base dimension of the technological regime. The high number of patent–IPC observations and the presence of multiple dominant technological classes indicate that AI innovation requires the integration of heterogeneous competencies. This supports the interpretation of AI as a technology with a broad and interdisciplinary knowledge base. At the same time, this complexity may generate barriers to entry, especially in frontier segments where advanced technical capabilities, specialized expertise, and complementary infrastructures are required.

The third dimension concerns the distribution of inventive activity across applicants. The dataset includes 7,019 distinct applicants, suggesting that a large number of actors participate in AI-related inventive activity. This is consistent with the idea that artificial intelligence offers broad opportunities for entry and experimentation. However, the mere number of applicants is not sufficient to evaluate the openness of the regime. What matters is whether inventive activity is evenly distributed or concentrated among a limited number of leading actors.

The applicant-level distribution provides a first indication of concentration within the AI technological domain. The presence of a large number of applicants indicates that entry into AI-related patenting activity is possible and relatively widespread. At the same time, the fact that leading applicants account for a significant share of patents suggests that inventive activity is not evenly distributed.

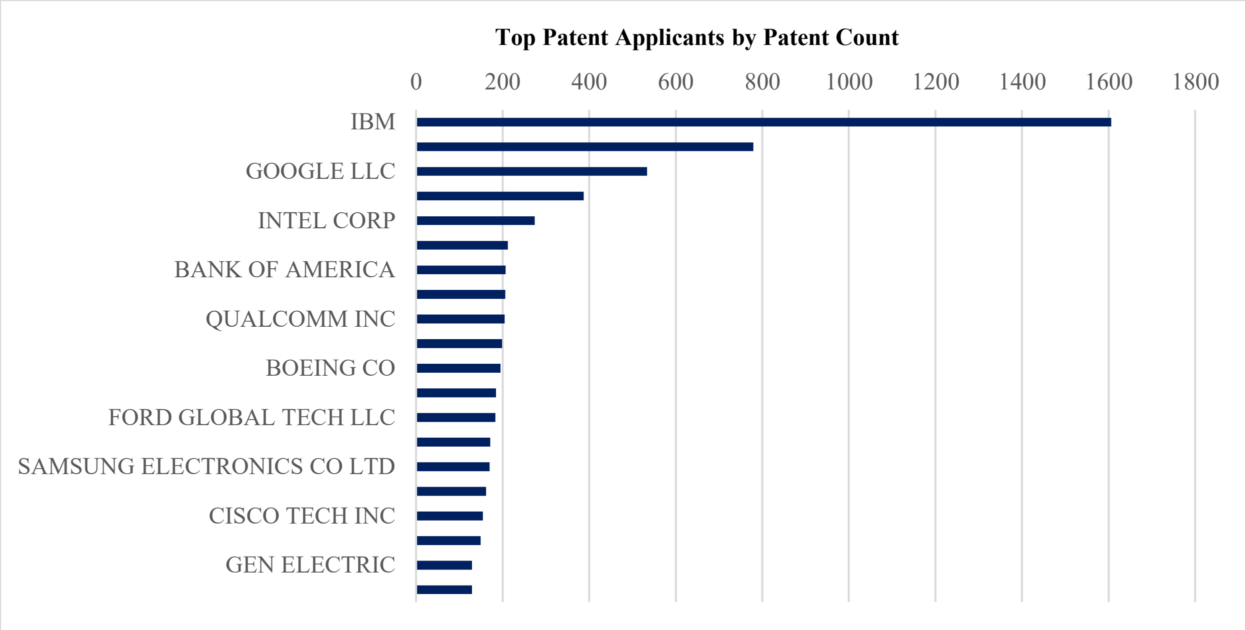

Figure 3.4. Selected leading patent applicants by patent count, highlighting the concentration of inventive activity among major firms.

IBM appears as the leading applicant in the dataset, confirming the role of large technology firms in shaping AI-related technological trajectories.

To better characterize the distribution of inventive activity across applicants, it is useful to move beyond the identification of leading firms and consider the overall structure of patent concentration. In particular, the distribution of patents appears to be highly skewed, with a relatively small number of applicants accounting for a disproportionately large share of total inventive activity, while the majority of actors contribute only marginally.

This type of distribution is consistent with a "long tail" structure, in which a limited core of dominant innovators coexists with a large periphery of occasional or small-scale contributors. From the perspective of technological regimes, such a configuration provides a first indication of the coexistence of entry and concentration: while access to AI-related innovation appears relatively open in terms of participation, the ability to generate a substantial and persistent volume of inventive output is concentrated among a restricted group of actors.

This evidence is particularly relevant for the interpretation of appropriability conditions. A skewed distribution of patents suggests that leading firms are able to capture a significant share of innovation-related returns, not only through formal intellectual property rights, but also through the accumulation of capabilities, access to complementary assets, and organizational advantages that reinforce their position over time.

More generally, the prominence of large technology firms among top patent applicants suggests that innovation in core segments is shaped by actors endowed with substantial technological and organizational capabilities, consistent with Schumpeter Mark II dynamics of accumulation and persistence.

This pattern is consistent with the hybrid nature of the AI regime. On the one hand, the large number of applicants reflects widening dynamics, associated with entry, experimentation, and the diffusion of AI-related capabilities across many actors. On the other hand, the concentration of patenting activity among leading firms points to deepening dynamics, where accumulated capabilities, research infrastructures, and complementary assets reinforce the position of incumbent innovators.

The coexistence of broad participation and concentrated leadership is particularly important for the interpretation of AI as a technological regime.

It suggests that entry exists, but that not all entry has the same relevance. Many actors may participate in AI innovation, especially in application-oriented or downstream segments, while a smaller group of firms appears better positioned to influence core technological trajectories. This distinction is central to the thesis, because it allows the analysis to move beyond a simple count of entrants and to focus instead on the quality and structural impact of entry.

Overall, the descriptive evidence points to three preliminary findings. First, the size and temporal coverage of the dataset suggest that AI represents a highly dynamic technological domain, consistent with high technological opportunities. Second, the structure of IPC classifications indicates a complex and combinatorial knowledge base, in which innovation depends on the integration of multiple technological domains. Third, the distribution of applicants suggests the coexistence of widespread participation and concentrated inventive leadership. Taken together, these patterns are consistent with the interpretation of artificial intelligence as a hybrid technological regime, characterized by both Mark I dynamics of entry and experimentation and Mark II dynamics of accumulation and concentration.

While the descriptive evidence provides a coherent and intuitively consistent picture of AI as a hybrid technological regime, it remains necessary to move beyond qualitative patterns and assess these dynamics through more formal empirical measures. In particular, the observed concentration of inventive activity calls for a systematic quantification, in order to evaluate the extent to which innovation is dominated by a limited number of actors and how this distribution evolves over time.

At the same time, the coexistence of widespread participation and concentrated leadership raises further questions regarding the persistence of innovators and the selective nature of entry. Descriptive evidence alone cannot fully capture whether leading firms maintain their position over time, or whether new entrants are able to emerge and alter the structure of the technological domain.

For these reasons, the following section develops a set of quantitative indicators aimed at measuring the degree of concentration, the persistence of innovative activity, and the dynamic structure of entry within AI patenting. This allows for a more rigorous assessment of the appropriability and cumulativeness dimensions of the technological regime, as well as for a more precise evaluation of the possible emergence of a core–periphery structure.

Overall, the descriptive evidence presented in this chapter provides a first indication of the hybrid nature of artificial intelligence as a technological regime. AI-related patenting expands substantially over the period 2006–2024, suggesting the presence of broad technological opportunities and an increasing number of innovation trajectories. At the same time, the technological composition of patents reveals a complex and combinatorial knowledge base, strongly rooted in digital and computational technologies. Finally, the applicant-level evidence shows that participation in AI inventive activity is widespread, but not evenly distributed across actors.

These preliminary findings are consistent with the idea that AI combines openness and concentration. However, descriptive evidence alone is not sufficient to assess whether this configuration corresponds to a stable industrial dynamic. In particular, it remains necessary to evaluate more systematically the degree of concentration of inventive activity, the evolution of entry over time, and the persistence of leading innovators. These dimensions are crucial for understanding whether AI patenting reflects mainly widening dynamics, associated with entry and experimentation, or deepening dynamics, associated with cumulative advantages and persistent leadership.

The next chapter therefore develops a more focused empirical analysis of concentration, entry and persistence in AI patenting. This allows the thesis to assess whether artificial intelligence can be interpreted as a hybrid technological regime characterized by a core–periphery structure, in which broad participation coexists with concentrated and persistent innovative leadership.

Chapter 4 – Concentration, Entry and Persistence in AI Patenting

4.1 Concentration and the Distribution of Innovative Activity

Building on the descriptive evidence presented in the previous chapter, this section provides a more systematic analysis of the distribution of inventive activity across applicants. The objective is to assess whether the expansion of AI-related patenting has been accompanied by a broad diffusion of innovative activity, or whether patenting remains concentrated among a restricted group of actors.

This distinction is central to the interpretation of artificial intelligence as a technological regime. As discussed in the theoretical framework, high technological opportunities do not necessarily lead to a fragmented or fully open industrial structure. In technologies characterized by strong cumulativeness, appropriability mechanisms and relevant complementary assets, the entry of many actors may coexist with the persistence of concentrated innovative leadership. Therefore, the analysis of concentration is necessary to understand whether AI patenting reflects a genuinely decentralized innovation process or, instead, a structure in which a large number of marginal participants coexist with a smaller group of dominant innovators.

The applicant-level evidence suggests a clearly asymmetric structure.

The dataset includes 7,019 distinct applicants, indicating that participation in AI-related inventive activity is broad. However, the presence of many applicants does not imply that inventive output is evenly distributed. On the contrary, the data point to a skewed distribution, in which a limited number of applicants account for a disproportionately large share of patenting activity, while most actors contribute only occasionally.

This pattern is consistent with a long-tail structure, where broad participation at the periphery coexists with concentration at the core of the technological domain.

To examine this concentration, the analysis first considers the patenting activity of the largest applicants. Table 4.1 reports the twenty applicants with the highest number of AI-related patent mentions. IBM is by far the leading applicant, with 1,606 patent mentions, followed by Microsoft Technology Licensing LLC with 779, Google LLC with 534, and Adobe Inc. with 387. The distance between IBM and the rest of the ranking is substantial, suggesting that leadership in AI patenting is not only concentrated among a small group of firms, but also uneven within that leading group.

Other major applicants include Intel, Amazon, Qualcomm, NVIDIA, Samsung, Apple, Cisco and Baidu, together with firms operating in finance, automotive and aerospace. This composition confirms that AI-related patenting is not confined to a single sector. However, the most intensive patenting activity is mainly associated with large corporations possessing strong technological, financial and organizational capabilities.

Top 20 applicant	
Applicant	**Patent count**
IBM	1,606
MICROSOFT TECHNOLOGY LICENSING LLC	779
GOOGLE LLC	534
ADOBE INC	387
INTEL CORP	274
AMAZON TECH INC	212
BANK OF AMERICA	207
QUALCOMM INC	206
CAPITAL ONE SERVICES LLC	205
NVIDIA CORP	199
BOEING CO	195
ORACLE INT CORP	185
SAMSUNG ELECTRONICS CO LTD	183
APPLE INC	171
FORD GLOBAL TECH LLC	170
MICROSOFT CORP	162
DELL PRODUCTS LP	154
CISCO TECH INC	149
X DEV LLC	129
BAIDU USA LLC	129

Table 4.1. Top 20 applicants by number of AI-related patent mentions.

The concentration ratios reported in Table 4.2 provide a direct measure of the weight of leading innovators within the AI patenting landscape. The top 5 applicants account for 3,580 patent mentions, corresponding to 14.73% of total applicant mentions. The top 10 applicants account for 4,609 mentions, or 18.96% of the total. The top 20 applicants account for 6,236 mentions, equal to 25.66% of total applicant mentions. These figures are calculated using applicant mentions rather than unique patents, because a single patent may be associated with more than one applicant. This approach therefore captures the distribution of inventive activity across applicants more accurately in the presence of co-application.

Group	Applicants included	Patent mentions	Total applicant mentions	Share
Top 5	5	3,580	24,305	14.73%
Top 10	10	4,609	24,305	18.96%
Top 20	20	6,236	24,305	25.66%

Table 4.2. Share of AI-related patent mentions accounted for by the Top 5, Top 10 and Top 20 applicants.

The evidence reported in Table 4.2 should be interpreted in light of the theoretical distinction between participation and leadership. A large number of applicants indicates that AI offers opportunities for entry and experimentation. However, the concentration ratios show that the ability to generate a large and persistent volume of patenting activity is far from evenly distributed. In particular, the fact that only twenty applicants account for more than one quarter of total applicant mentions suggests that broad participation coexists with a selective structure of innovative leadership. AI patenting is therefore not simply a decentralized process involving many equivalent contributors. Rather, it is characterized by a hierarchy in which a small core of applicants plays a much stronger role than the majority of participants.

This pattern is particularly relevant for the appropriability dimension of the technological regime. In artificial intelligence, appropriability does not depend only on formal intellectual property protection. It is also connected to the control of complementary assets such as data, computational infrastructure, research capabilities, digital platforms and distribution channels. Firms that control these assets are better positioned to transform technological opportunities into cumulative inventive output. The presence of companies such as IBM, Microsoft, Google, Amazon, Intel, NVIDIA, Samsung and Apple among the leading applicants is consistent with this interpretation. These actors combine technological capabilities with complementary assets that allow them to accumulate and protect knowledge more effectively than occasional entrants. At the same time, the concentration ratios do not indicate complete closure of the technological domain. The top 20 applicants account for 25.66% of total applicant mentions, which is substantial but still leaves almost three quarters of applicant mentions outside the leading group. This means that AI patenting is neither fully concentrated nor fully dispersed. Instead, the evidence points to a mixed structure: a broad population of applicants participates in AI-related inventive activity, but the intensity of patenting is much higher among a limited set of large and persistent innovators.

A second way to examine concentration is to observe the overall distribution of patents across applicants. Instead of focusing only on the largest firms, this approach considers the entire applicant population and allows the analysis to identify whether most actors contribute substantially or only marginally to AI-related inventive activity. This is important because the number of applicants alone may overstate the openness of the regime. If most applicants hold only one or a small number of patents, then entry is formally broad but innovative leadership remains concentrated.

Patent-count bin	Applicants	Share of applicants
1	5,123	72.95%
2	814	11.59%
3-5	589	8.39%
6-10	240	3.42%
11-20	113	1.61%
21-50	83	1.18%
51-100	27	0.38%
101+	30	0.43%

Table 4.3. Distribution of applicants by number of AI-related patent mentions.

Table 4.3 confirms the strong asymmetry of the applicant distribution. A very large share of applicants is associated with only one or very few patents. In particular, 5,123 applicants, equal to approximately 72.95% of all applicants, appear only once in the dataset. A further 814 applicants, corresponding to 11.59%, are associated with two patents. This means that more than four fifths of all applicants contribute only one or two patent mentions.

By contrast, repeated and intensive patenting activity is limited to a much smaller group. Applicants with between 3 and 5 patents account for 8.39% of the population, while those with 6 to 10 patents represent only 3.42%. The upper tail is even narrower: applicants with more than 100 patent mentions account for only 0.43% of all applicants. This distribution clearly shows that entry into AI patenting is possible, but that entry alone does not necessarily translate into a strong innovative presence.

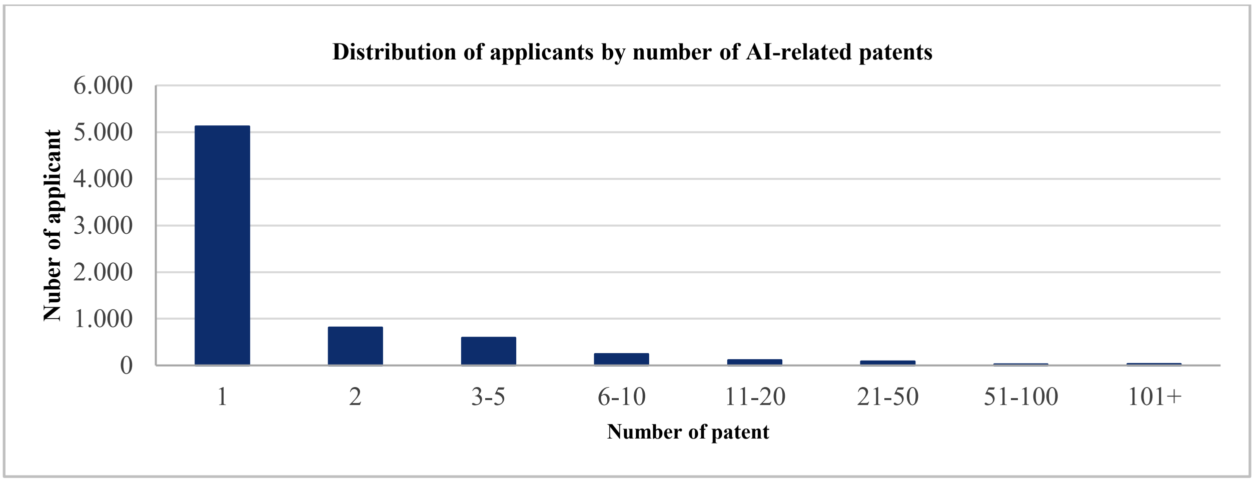

Figure 4.1. Distribution of applicants by number of AI-related patent mentions.

This evidence supports the idea that AI combines openness and selectivity. On the one hand, the existence of thousands of applicants indicates that the technological domain is not closed to new actors. The diffusion of AI tools, software libraries, cloud-based services and pre-trained models may reduce some barriers to participation, especially in downstream and application-oriented activities. On the other hand, the unequal distribution of patents suggests that only a limited group of actors is able to accumulate a substantial stock of inventive activity. This is consistent with a hybrid regime in which Mark I dynamics of entry and experimentation coexist with Mark II dynamics of accumulation and concentration.

The analysis can be further strengthened by using a synthetic concentration indicator, such as the Herfindahl-Hirschmann Index. The HHI captures not only the share of the largest applicants, but the overall degree of concentration in the patent distribution. A higher value indicates that inventive activity is more concentrated, while a lower value suggests a more dispersed structure. In this analysis, the yearly HHI is calculated as the sum of squared applicant shares in each year, multiplied by 10,000.

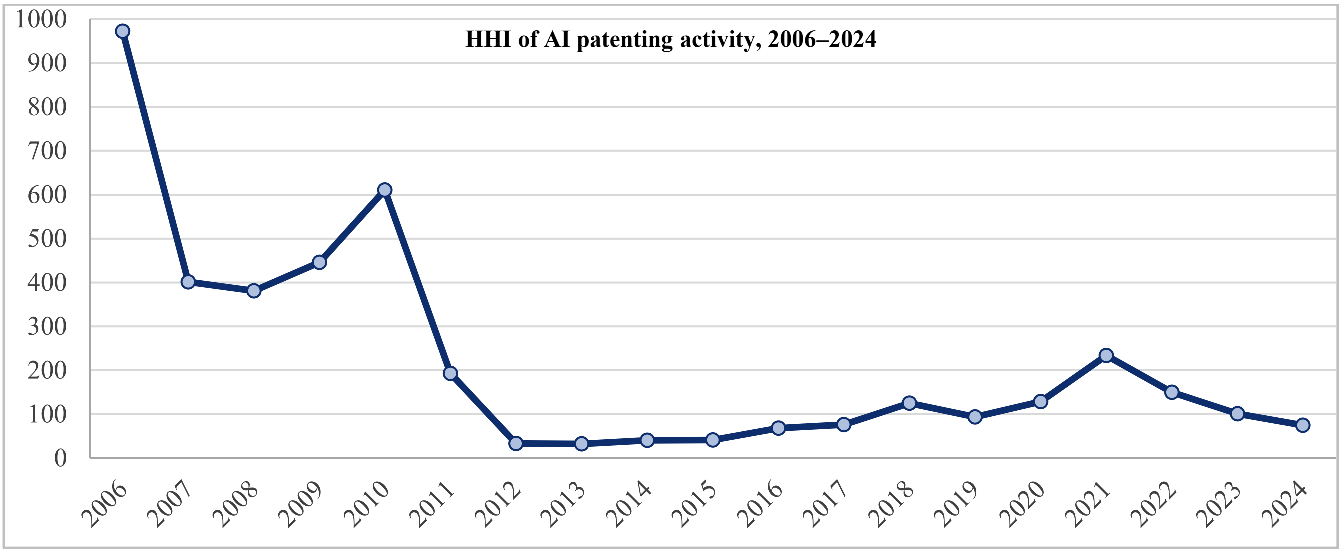

Figure 4.2. Herfindahl-Hirschman Index of AI patenting activity over time.

Note: The vertical axis reports the Herfindahl-Hirschman Index, calculated as $HHI_t = \Sigma_i(s_{i,t}^2) \times 10{,}000$*, where* $s_{i,t}$ *is the share of applicant mentions accounted for by applicant i in year t. Higher values indicate greater concentration of AI-related patenting activity among a smaller number of applicants. The high values observed in the early years should be interpreted with caution, as they partly reflect the smaller number of patents and applicants at the beginning of the observation period.*

The evolution of the HHI over time adds an important dynamic perspective. In the earliest years, the HHI is relatively high, reaching 972.22 in 2006 and 610.43 in 2010. These values should be interpreted with caution, because the number of applicants' mentions in the initial years is still very small. In such cases, even a few active applicants can generate relatively high concentration values.

From 2012 onward, when AI patenting becomes more substantial, the HHI declines sharply and remains at much lower levels for several years. For example, the index is 33.10 in 2012, 32.29 in 2013, and 41.08 in 2015. This suggests that the expansion of AI patenting was accompanied by a broader participation of applicants, reducing the overall degree of concentration when measured across the full applicant population.

However, the later period shows that concentration did not disappear entirely. The HHI increases again in some years, reaching 125.16 in 2018, 128.76 in 2020, and 233.40 in 2021, before declining to 150.04 in 2022, 100.94 in 2023, and 74.95 in 2024. This pattern suggests that AI patenting became more populated over time, but the position of leading applicants remained relevant. In other words, the growth of the technological field widened participation, but did not fully eliminate the structural advantages of more persistent and better-resourced innovators.

Overall, the concentration evidence points toward a clear empirical result: AI-related inventive activity is broad in terms of participation, but asymmetric in terms of output. The large number of applicants confirms that artificial intelligence offers substantial opportunities for entry and experimentation. However, the concentration ratios, the ranking of leading applicants, the long-tail distribution and the yearly HHI all indicate that innovative leadership remains more selective than simple participation counts would suggest.

This result is important for the broader argument of the thesis. It suggests that AI should not be interpreted simply as a democratizing technology because many actors are able to participate in its development. What matters is not only how many firms enter the technological domain, but also whether they are able to accumulate innovative capacity and influence the main technological trajectories. The concentration of patenting activity among leading applicants therefore provides an empirical indication of a core–periphery structure: a broad periphery of occasional or low-intensity innovators coexists with a more restricted core of actors that accounts for a disproportionate share of AI inventive activity.

Thus, the evidence supports the interpretation of AI as a hybrid technological regime. On the one hand, the field displays openness, entry and experimentation. On the other hand, it also shows persistence, accumulation and selective leadership. The expansion of AI patenting has therefore broadened the innovative base, but it has not automatically produced de-concentration.

4.2 Entry Dynamics and the Expansion of the Applicant Base

The evidence presented in the previous section shows that AI-related patenting is characterized by a strongly asymmetric distribution of inventive activity. A large number of applicants participate in the technological domain, but patenting output remains concentrated among a restricted group of leading actors. This section complements the analysis of concentration by focusing on entry dynamics. The objective is to assess whether the expansion of AI patenting over time has been driven by the continuous entry of new applicants, by the persistence of already active applicants, or by the coexistence of both processes.

This distinction is central to the interpretation of artificial intelligence as a hybrid technological regime. In a Schumpeter Mark I configuration, technological opportunities are expected to generate continuous entry, experimentation and turbulence. In a Schumpeter Mark II configuration, by contrast, innovation tends to be dominated by actors that accumulate capabilities over time and maintain a persistent presence within the technological domain. Entry dynamics therefore allow us to evaluate whether AI patenting reflects mainly widening dynamics, associated with the expansion of the population of innovators, or deepening dynamics, associated with cumulative innovative activity by continuing applicants.

For this purpose, applicants are divided into two categories. A new applicant is defined as an applicant whose first observed year in the dataset corresponds to the year under consideration. A continuing applicant is instead an applicant active in a given year who had already appeared in at least one previous year. This distinction makes it possible to separate the yearly expansion of the applicant population from the repeated participation of actors already present in the AI patenting domain.

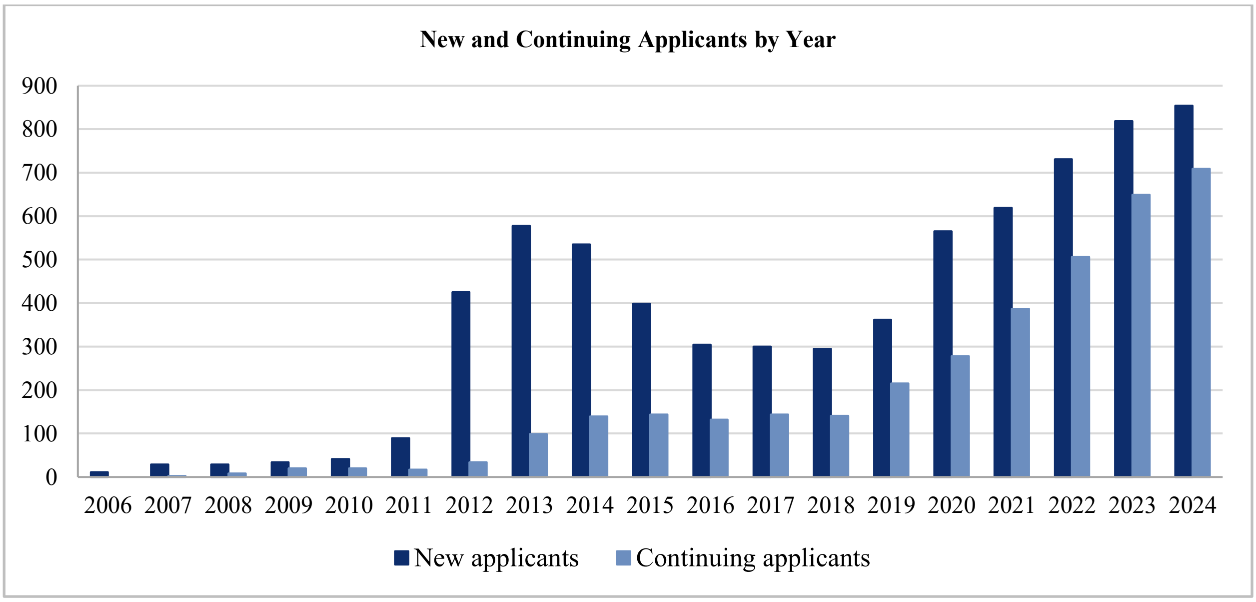

Figure 4.3. Annual number of new and continuing applicants, 2006–2024.

Note: New applicants are applicants whose first observed year in the dataset corresponds to the current year. Continuing applicants are applicants active in the current year but already observed in at least one previous year.

The annual evidence shows that the AI patenting domain experienced a substantial expansion of its applicant base over time.

In the earliest years, the number of active applicants was still limited. In 2006, all 11 active applicants were new entrants, which is mechanically expected at the beginning of the observation period. In the following years, both the number of active applicants and the number of new applicants increased, although the scale of entry remained relatively modest until the early 2010s.

A clear break emerges from 2012 onward. The number of active applicants increased from 106 in 2011 to 459 in 2012, while the number of new applicants rose sharply from 89 to 425. This acceleration indicates that the AI technological domain became increasingly attractive to a broader population of innovators. The expansion continued in subsequent years, with 578 new applicants in 2013 and 535 in 2014. These figures suggest that the early 2010s represented an important phase of widening in AI patenting, during which the technological space opened to a much larger number of actors.

After 2014, the number of new applicants declined for some years, reaching 295 in 2018. However, this decline does not indicate a closure of the technological domain. Rather, it reflects the transition from an initial phase of rapid entry to a more structured phase in which both new and continuing applicants coexist. From 2019 onward, entry increased again substantially. New applicants rose from 362 in 2019 to 565 in 2020, 619 in 2021, 731 in 2022, 819 in 2023 and 854 in 2024. By the end of the period, the number of new entrants was therefore much higher than in the initial phase of the dataset.

At the same time, the number of continuing applicants also increased markedly. Continuing applicants rose from 34 in 2012 to 144 in 2015, 278 in 2020, 506 in 2022, 649 in 2023 and 709 in 2024. This trend is particularly important because it shows that AI patenting is not only expanding through new entry. A growing number of applicants remain active after their first appearance, indicating the emergence of a more persistent population of innovators. In other words, the applicant base is not simply renewed every year; it also accumulates over time.

The coexistence of high entry and increasing continuity is visible in the share of new applicants over total active applicants. In the early years, this share was extremely high, reflecting the still limited size of the applicant population. In 2012, new applicants accounted for 92.59% of active applicants, confirming the strong widening of the technological domain. However, the share of new applicants declined progressively over time, falling to 67.02% in 2020, 61.53% in 2021, 59.09% in 2022, 55.79% in 2023 and 54.64% in 2024.

This decline should not be interpreted as a reduction in entry in absolute terms. On the contrary, the absolute number of new applicants reached its maximum in the final years of the period. Rather, the declining share indicates that the population of continuing applicants grew faster over time. This is consistent with the emergence of cumulative dynamics within the AI patenting domain. As more firms and organizations enter the field, a fraction of them remains active in subsequent years, increasing the weight of continuing innovators in the annual applicant population.

This result is important for the interpretation of AI as a technological regime. The high number of new applicants confirms that artificial intelligence offers broad technological opportunities and allows a large number of actors to enter the patenting domain. This evidence is consistent with Mark I dynamics of widening, experimentation and diffusion. However, the simultaneous growth of continuing applicants suggests that entry is not purely temporary or episodic. Over time, part of the applicant population becomes persistent, indicating that accumulated experience, organizational capabilities and technological positioning matter for continued participation. This second element is consistent with Mark II dynamics of cumulativeness and persistence.

The cumulative expansion of the applicant base further reinforces this interpretation. The number of distinct applicants increased from only 11 in 2006 to 144 in 2010, 658 in 2012, 1,236 in 2013 and 3,069 in 2018. The expansion accelerated again in the final years, reaching 3,996 cumulative applicants in 2020, 5,346 in 2022, 6,165 in 2023 and 7,019 in 2024.

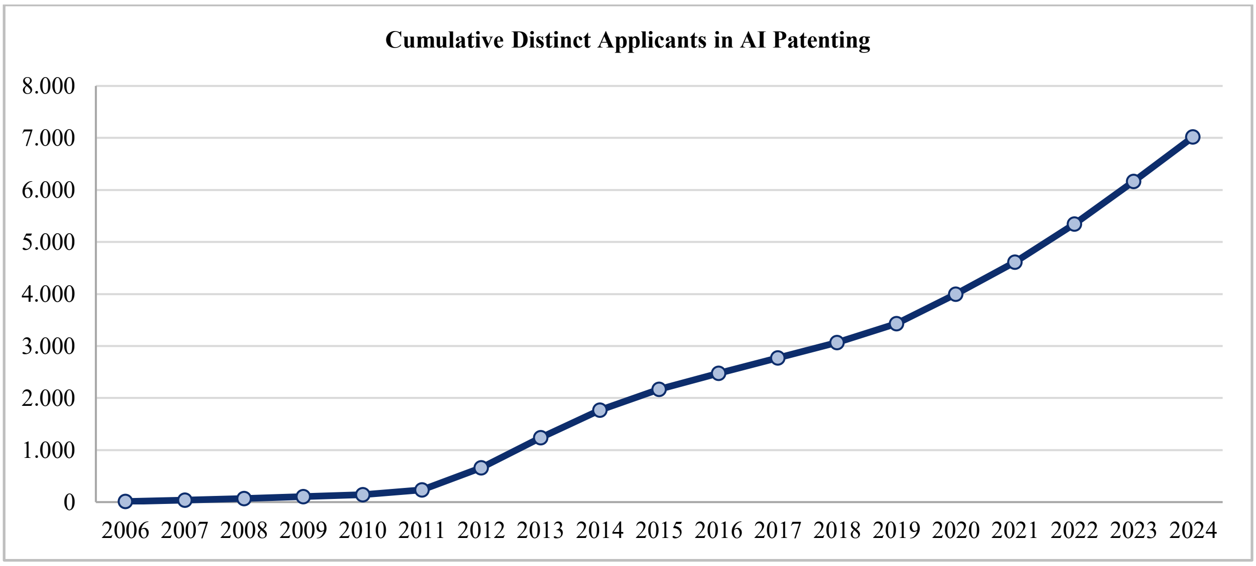

Figure 4.4. Cumulative number of distinct applicants, 2006–2024.

Figure 4.4 reports the cumulative number of distinct applicants entering the AI patenting domain over the full observation period.

The cumulative trajectory shows that AI patenting progressively attracted a very broad population of actors. This pattern confirms that the technological domain did not remain restricted to a small set of early innovators. Instead, the applicant base expanded continuously throughout the period. From the perspective of technological opportunities, this is consistent with the idea that AI provides multiple innovation trajectories and application spaces, allowing different types of firms and organizations to participate in inventive activity.

However, the cumulative expansion of the applicant base must be interpreted together with the concentration evidence presented in Section 4.1. The fact that 7,019 distinct applicants appear in the dataset by 2024 does not mean that innovative activity is evenly distributed across them. As shown previously, most applicants contribute only marginally, while a relatively small group accounts for a disproportionate share of patenting activity. Entry is therefore broad, but not necessarily transformative. Many actors enter the AI patenting domain, but only a limited subset appears capable of accumulating a substantial volume of inventive output.

This distinction between entry and transformative entry is crucial. The empirical evidence suggests that AI lowers some barriers to participation, especially because the technology has a wide range of applications and can be recombined across different technological fields. At the same time, the ability to remain active and accumulate inventive output is more selective. Persistent participation likely requires complementary assets such as research capabilities, data resources, computational infrastructure, specialized human capital and organizational routines. These assets are not equally distributed across applicants, which helps explain why broad entry can coexist with concentration.

The entry evidence therefore strengthens the interpretation of AI as a hybrid technological regime. On the one hand, the continuous arrival of new applicants reflects widening dynamics. The technological domain expands, new actors enter, and the innovation base becomes broader. On the other hand, the growing number of continuing applicants and the concentration of patenting activity among leading firms point to deepening dynamics. Over time, some actors consolidate their presence, while many others remain occasional or peripheral contributors.

Overall, the analysis of entry dynamics shows that the growth of AI patenting has been driven by both new entry and increasing persistence. The applicant base expanded dramatically between 2006 and 2024, reaching 7,019 distinct applicants by the end of the period. This confirms that artificial intelligence is characterized by broad opportunities for participation. At the same time, the declining share of new applicants and the growing number of continuing applicants indicate that cumulative participation becomes increasingly relevant as the technological domain matures.

Taken together with the concentration evidence discussed in Section 4.1, these results suggest that AI patenting is neither fully open nor fully closed. It is open in the sense that many actors are able to enter the technological domain. Yet it is selective in the sense that persistent and intensive innovative activity remains concentrated among a smaller group of actors. This supports the broader argument of the thesis: artificial intelligence operates as a hybrid technological regime, combining Mark I dynamics of entry and experimentation with Mark II dynamics of accumulation, persistence and concentrated leadership.

4.3 Technological Portfolio Strategies of Leading Applicants

The previous sections have shown that AI-related patenting combines broad entry with persistent asymmetries in the distribution of inventive output. Section 4.1 documented that patenting activity is highly skewed across applicants, with a small group of leading actors accounting for a disproportionate share of total applicant mentions. Section 4.2 then showed that the applicant base expanded substantially over time, indicating that artificial intelligence remains open to new entrants and experimentation. Taken together, these results support the interpretation of AI as a hybrid technological regime, in which broad participation coexists with selective and concentrated innovative leadership.

However, concentration and entry indicators do not fully capture the technological strategies of leading applicants. Knowing that a firm patents intensively in AI does not reveal whether its inventive activity is concentrated in a narrow technological area or diversified across multiple technological trajectories. Similarly, observing that a firm remains active over time does not indicate whether it is exploiting an established technological specialization or progressively expanding its knowledge base. For this reason, the analysis is extended to the technological portfolios of leading applicants.

This firm-level perspective is important because the previous evidence mainly described the structure of the technological domain at the aggregate level: the overall distribution of patenting activity, the number of entrants, the concentration of applicant mentions, and the general composition of IPC classes. These indicators are necessary to characterize the AI regime, but they may hide heterogeneity across firms. Within the same technological regime, some applicants may follow broad exploratory strategies, while others may remain concentrated in a smaller set of technological fields. Technological portfolios therefore provide additional evidence on the knowledge-base dimension of the regime and on the heterogeneity of firm-level innovative strategies.

The analysis focuses on the top 15 applicants by number of AI-related patent mentions. For each applicant, the patent portfolio is decomposed by IPC class in order to identify the main technological areas in which inventive activity is concentrated. This firm-level analysis complements the aggregate IPC evidence presented in Chapter 3. While the previous analysis showed that AI patents are concentrated mainly in digital and computational classes such as G06F, G06N, G10L, G06V and G06T, the present section investigates whether the leading applicants display similar or heterogeneous technological profiles.

The evidence shows that the leading applicants differ substantially in the breadth and concentration of their technological portfolios. IBM, the largest applicant in the dataset, records 1,606 patent mentions and 6,976 IPC observations distributed across 88 distinct IPC classes. Its portfolio is mainly concentrated in G06F, which accounts for 36.63% of its IPC observations, followed by G06N with 29.31% and G10L with 5.75%. Together, IBM's top three IPC classes account for 71.69% of its portfolio. This indicates that IBM combines a very broad technological scope with a strong concentration around core computational and AI-related classes.

A similar pattern emerges for Microsoft Technology Licensing LLC, which records 779 patent mentions, 3,572 IPC observations and 47 distinct IPC classes. Its main classes are G06F, G06N and G10L, accounting respectively for 36.31%, 26.34% and 13.38% of its IPC observations. The top three classes represent 76.04% of Microsoft's portfolio, suggesting a relatively concentrated technological profile despite the presence of a broad set of IPC classes.

Google LLC displays a slightly different composition. Its main class is G06N, accounting for 27.62% of IPC observations, followed by G10L with 22.03% and G06F with 20.35%. Its top three classes account for 70.01% of the portfolio. This confirms the centrality of machine learning, language processing and computational methods in the portfolios of the largest digital technology firms.

The portfolios of other leading applicants also show the importance of core computational and AI-related IPC classes, although with different combinations. Adobe is concentrated mainly in G06N, G06F and G06T, reflecting a profile connected to machine learning, computation and image processing. Intel combines G06N and G06F with H04L, indicating the relevance of computational models, data processing and communication technologies. Amazon is characterized by a strong presence of G10L, followed by G06F and G06N, suggesting the importance of speech and language-related technologies alongside computational methods. NVIDIA is mainly associated with G06N, G06T and G06V, which is consistent with its role in machine learning, image processing and computer vision.

The presence of financial applicants such as Bank of America and Capital One Services LLC is also relevant. Their portfolios are dominated by G06F and G06N, but they also include G06Q, which is associated with data processing systems for administrative, commercial and financial purposes. This suggests that AI-related patenting is not confined to firms producing digital technologies directly, but also includes large firms applying AI to specific sectoral activities. Similarly, Ford Global Technologies and Boeing show more sector-specific profiles. Ford combines G06N, G06V and B60W, indicating the relevance of AI, computer vision and vehicle-control technologies. Boeing displays a more sector-specific profile, with a portfolio connected to robotics, automation, computational methods and applied AI-related technologies.

A central element in this analysis is the distinction between technological breadth and technological focus. Technological breadth refers to the number of distinct IPC classes in which an applicant patents. A broader portfolio may indicate an exploratory strategy, in which the applicant applies AI-related knowledge across multiple technological fields or combines different knowledge domains. Technological focus, by contrast, refers to a more concentrated portfolio, where patenting activity is mainly located in a limited number of IPC classes. A focused portfolio may indicate specialization, cumulative learning within specific technological trajectories, or the strategic consolidation of capabilities in areas where the applicant already holds a strong position.

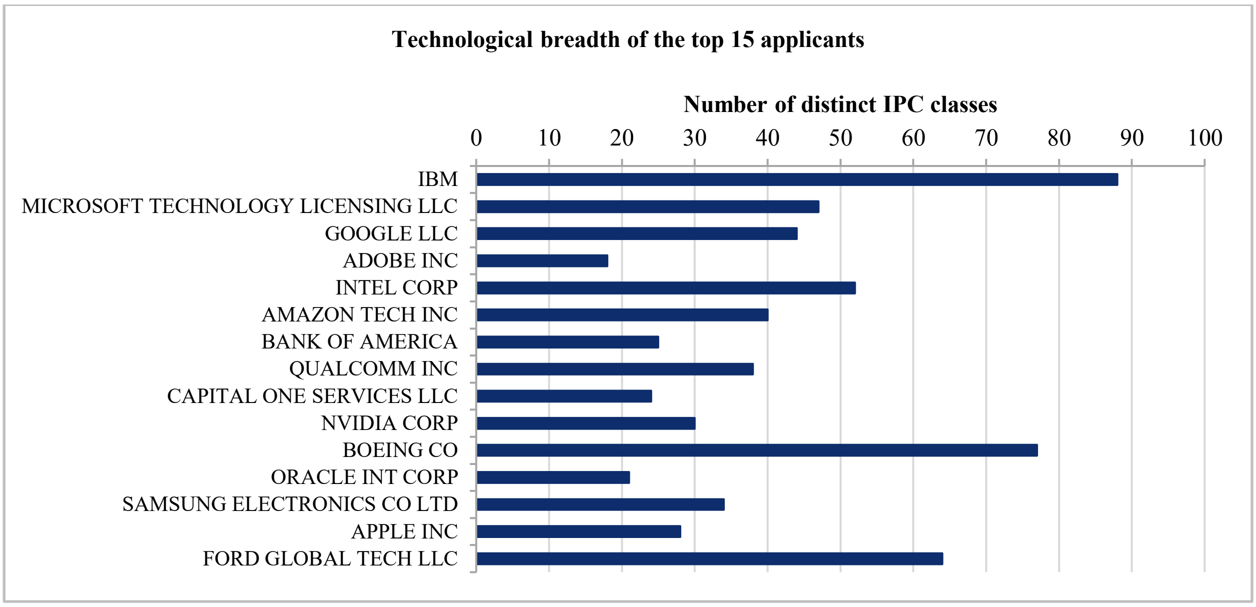

Figure 4.5. Technological breadth of the top 15 applicants.

Note: Technological breadth is measured as the number of distinct IPC classes associated with each applicant's AI-related patent portfolio.

Figure 4.5 highlights clear differences in technological breadth across the top 15 applicants. IBM displays the broadest portfolio, with 88 distinct IPC classes, followed by Boeing with 77 and Ford Global Technologies with 64. Intel also shows a broad portfolio, with 52 distinct IPC classes, while Microsoft Technology Licensing LLC and Google LLC record 47 and 44 classes respectively. These values indicate that some leading applicants operate across a wide technological space and are able to connect AI-related inventive activity to multiple IPC domains.

At the opposite end, some leading applicants appear more technologically focused. Adobe records 18 distinct IPC classes, Oracle 21, Capital One Services LLC 24, Bank of America 25 and Apple 28. These firms are still among the leading applicants in terms of patent mentions, but their portfolios are less diversified in IPC terms. This suggests that innovative leadership in AI does not always take the same form. Some firms achieve leadership through broad technological diversification, while others concentrate their activity in a narrower set of technological trajectories.

The portfolio HHI provides a complementary measure of this distinction. A higher portfolio HHI indicates that an applicant's IPC observations are more concentrated in a limited number of classes, while a lower value suggests a more dispersed technological portfolio. Oracle has the highest portfolio HHI among the top 15 applicants, equal to 2,871.78, followed by IBM with 2,301.09 and Microsoft Technology Licensing LLC with 2,264.05. This indicates that, despite having many IPC observations, these firms concentrate a large share of their activity in a small number of core classes. By contrast, Ford Global Technologies and Boeing show the lowest portfolio HHI values, equal to 765.89 and 871.13 respectively, suggesting more dispersed portfolios. Qualcomm also appears relatively diversified, with a portfolio HHI of 1,233.42 and a top-three IPC share of 52.55%.

This evidence is important because it shows that concentration in patenting activity does not correspond to a single technological strategy. IBM and Microsoft are both large and active patenting organizations, but their portfolios show a high concentration around core computational classes. Boeing and Ford, by contrast, display broader and more dispersed IPC portfolios, reflecting the application of AI-related technologies across more sector-specific and heterogeneous technological domains. Oracle, Adobe, Capital One and Bank of America appear more focused, with a large share of their IPC observations concentrated in a smaller number of classes. Therefore, the core of the AI patenting system is internally differentiated: leading applicants differ not only in how much they patent, but also in how their technological knowledge is organized.

A second dimension concerns the evolution of technological breadth over time. A static view of the portfolio identifies whether an applicant is broad or focused over the whole observation period, but it does not show whether the applicant's strategy changes. For this reason, the analysis also considers the temporal evolution of technological breadth among leading applicants. This allows the section to examine whether firms progressively expand their technological portfolios, remain stable, or converge toward a smaller number of core IPC classes.

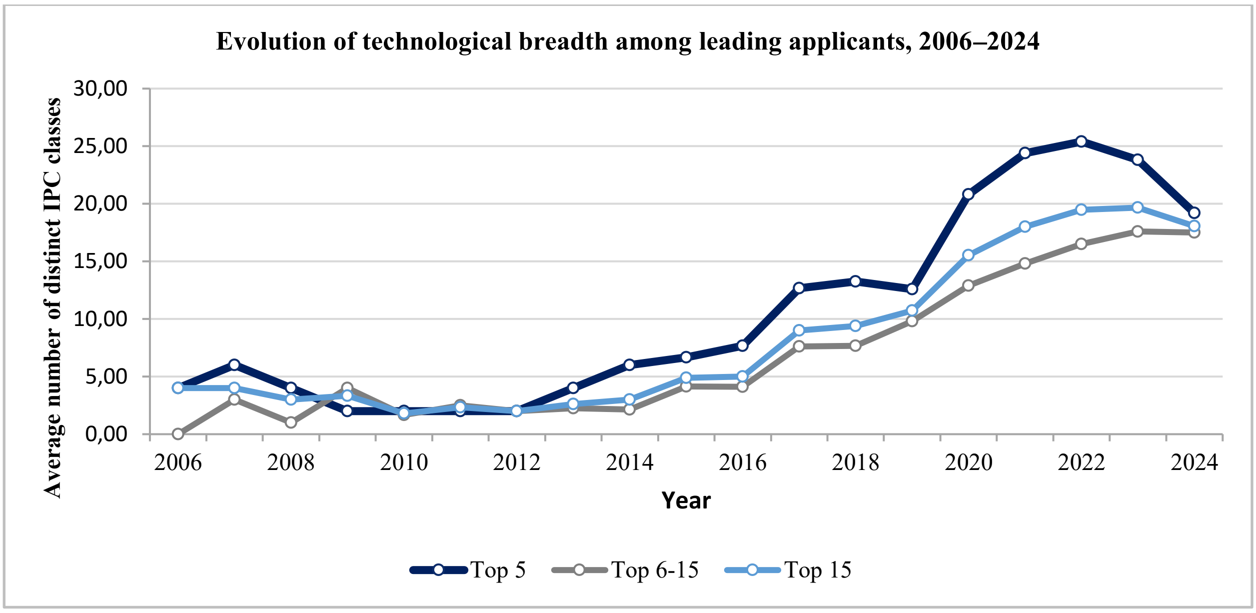

Figure 4.6. Evolution of technological breadth among leading applicants, 2006–2024.

Note: The figure reports the average number of distinct IPC classes associated with the yearly patent portfolios of the top 15 applicants. The comparison distinguishes between the top 5 applicants and applicants ranked from 6 to 15.

Figure 4.6 provides a dynamic view of the technological portfolios of leading applicants. In the early years of the dataset, the average technological breadth of the top applicants is low and volatile. This is expected, because only a small number of top applicants are active in the initial years. For example, in 2006 only one top 15 applicant is active, with an average breadth of 4 IPC classes. In 2010, five top 15 applicants are active, but the average breadth remains low, at 1.8 classes. These early values should therefore be interpreted with caution.

From the mid-2010s onward, however, the technological breadth of leading applicants increases substantially. The average breadth of the top 15 applicants rises from 4.9 IPC classes in 2015 to 9.0 in 2017, 10.73 in 2019, 15.53 in 2020 and 18.0 in 2021. It reaches 19.47 in 2022 and 19.67 in 2023, before declining slightly to 18.07 in 2024. This pattern indicates that, as AI patenting expands, the leading applicants increasingly operate across a broader range of IPC classes. The evidence is therefore consistent with a process of technological portfolio expansion among the most active innovators.

The distinction between the top 5 and the remaining top applicants provides further insight. The top 5 applicants display a particularly strong increase in technological breadth, rising from 6.67 IPC classes in 2015 to 12.67 in 2017, 20.8 in 2020 and 24.4 in 2021. Their average breadth reaches 25.4 in 2022, before declining to 23.8 in 2023 and 19.2 in 2024. Applicants ranked from 6 to 15 also expand their portfolios, but more gradually. Their average breadth increases from 4.14 in 2015 to 9.8 in 2019, 12.9 in 2020, 14.8 in 2021, 16.5 in 2022 and 17.6 in 2023, remaining almost stable at 17.5 in 2024.

This evidence suggests that leading applicants progressively broadened their technological portfolios over time. The strongest expansion is observed among the top 5 applicants, which is consistent with the idea that the largest innovators possess the organizational and technological capabilities required to operate across multiple AI-related domains. At the same time, the growth observed among applicants ranked from 6 to 15 indicates that portfolio expansion is not limited to the very top of the distribution. A broader group of leading applicants also increased the variety of IPC classes associated with their AI-related patenting activity.

From the perspective of technological regimes, this result is particularly relevant. The increase in technological breadth is consistent with the idea that AI is characterized by high technological opportunities and a combinatorial knowledge base. Leading firms do not simply repeat inventive activity within a fixed and narrow domain; rather, they increasingly connect AI-related inventions to a wider set of technological classes. This supports the interpretation of AI as a technology that enables recombination across fields and allows firms to extend their inventive activity into multiple domains.

At the same time, the portfolio evidence also confirms the persistence of technological focus around core classes. Even the broadest applicants remain strongly connected to central AI-related classes such as G06F and G06N. These classes recur across most of the leading portfolios, indicating that technological exploration takes place around a stable computational core. This is important because it qualifies the meaning of diversification. Portfolio expansion does not imply a complete dispersion of innovative activity across unrelated fields. Instead, it suggests a pattern of controlled diversification, in which leading applicants broaden their technological scope while maintaining a strong connection to core computational and machine-learning trajectories.

This finding strengthens the interpretation of AI as a hybrid Schumpeterian technological regime. On the one hand, the growth of technological breadth among leading applicants reflects exploration, recombination and the existence of broad technological opportunities. On the other hand, the recurrence of core IPC classes and the high portfolio concentration observed for several top applicants point to cumulativeness and path-dependent specialization. AI leadership therefore appears to depend not only on the volume of patenting activity, but also on the ability to structure and expand technological portfolios around a persistent knowledge base.

The firm-level evidence also adds an important qualification to the aggregate analysis. Chapter 3 showed that AI-related patenting is mainly concentrated in a limited number of IPC classes at the aggregate level. The present section shows that this aggregate pattern conceals heterogeneous firm strategies. Some applicants, such as IBM, Boeing, Ford and Intel, display relatively broad portfolios. Others, such as Adobe, Oracle, Capital One and Bank of America, appear more focused. This heterogeneity suggests that the core of the AI patenting system is not composed of identical actors. Rather, it includes firms with different technological positions, sectoral backgrounds and portfolio strategies.

Overall, the technological portfolio analysis provides a more nuanced understanding of innovative leadership in AI. Leading applicants are not only important because they account for a large share of patent mentions. They are also relevant because their portfolios reveal how technological capabilities are organized across different domains. The evidence shows that some leading firms pursue broad technological portfolios, while others concentrate their activity in more specific areas. Over time, however, the general tendency among top applicants is toward increasing technological breadth, especially from the mid-2010s onward.

Taken together with the previous sections, the analysis suggests that AI patenting is characterized by three related features. First, the distribution of inventive output is asymmetric, with a limited group of applicants accounting for a large share of patent mentions. Second, the applicant base expands continuously, indicating broad entry and experimentation. Third, among leading applicants, technological portfolios provide evidence of firm-level strategic heterogeneity and increasing technological breadth over time. This combination is consistent with a core–periphery structure in which the core is not homogeneous, but composed of firms with different degrees of diversification, specialization and technological scope.

The next section brings together the evidence on concentration, entry and technological portfolios in order to provide a synthetic interpretation of the empirical results and connect them to the hypotheses developed in the theoretical framework.

4.4 Empirical Synthesis: Concentration, Entry and Technological Portfolios

The empirical evidence presented in this chapter provides a systematic assessment of the structure and evolution of AI-related patenting. The objective was not simply to document the growth of artificial intelligence as a technological field, but to understand whether this growth has been accompanied by a broad diffusion of innovative activity or whether, instead, it has reproduced patterns of concentration, persistence and selective leadership. In this sense, the chapter has examined three complementary dimensions: the distribution of inventive output across applicants, the dynamics of entry into the patenting domain, and the technological portfolios of leading applicants.

Taken together, the results indicate that AI-related patenting is characterized by a dual structure. On the one hand, the technological domain has expanded considerably and has attracted a large number of applicants over time. On the other hand, this expansion has not produced an evenly distributed innovation structure. Participation is broad, but innovative output remains strongly asymmetric. This distinction between participation and leadership is central to the interpretation of AI as a hybrid Schumpeterian technological regime.

Section 4.1 showed that the distribution of inventive activity across applicants is highly skewed. The dataset includes 7,019 distinct applicants, confirming that AI patenting involves a broad population of actors. However, concentration ratios reveal that a restricted group of applicants accounts for a substantial share of total applicant mentions. The top 5 applicants account for 14.73% of total applicant mentions, the top 10 for 18.96%, and the top 20 for 25.66%. This evidence indicates that the presence of many applicants does not imply an equal distribution of inventive output.

The applicant distribution further reinforces this result. A very large share of applicants appears only marginally in the dataset: 72.95% of applicants are associated with only one patent mention, while more than four fifths of the applicant population contributes only one or two patent mentions. This confirms the existence of a long-tail structure, in which a broad periphery of occasional or low-intensity participants coexists with a smaller core of more intensive innovators. Entry into AI patenting is therefore possible, but the ability to accumulate a significant volume of inventive activity is much more selective.

The yearly HHI adds an important dynamic perspective to this evidence. In the early years, concentration is relatively high, partly because the number of active applicants is still limited. As AI patenting expands, the HHI declines, suggesting that the growth of the field is associated with broader participation. However, the persistence of leading applicants and the concentration ratios discussed above show that the expansion of the applicant base does not eliminate asymmetries in innovative output. The evolution of AI patenting therefore does not correspond to a simple process of de-concentration. Rather, it points to a more complex structure in which participation increases while innovative leadership remains unevenly distributed.

Section 4.2 complemented this evidence by examining the dynamics of entry. The analysis showed that the AI patenting domain has been characterized by a continuous inflow of new applicants throughout the period 2006–2024. The cumulative number of distinct applicants increased from 11 in 2006 to 7,019 in 2024, indicating a substantial enlargement of the innovative base. This result supports the view that artificial intelligence offers broad technological opportunities and creates space for entry, experimentation and recombination.

At the same time, entry dynamics reveal that the expansion of the applicant base is not driven only by temporary or one-off participation. The number of continuing applicants also increases markedly over time, reaching 709 in 2024. Moreover, the share of new applicants over total active applicants declines in the later years, falling to 54.64% in 2024. This decline should not be interpreted as evidence of reduced entry in absolute terms, since the number of new applicants reaches very high levels in the final years of the period. Rather, it indicates that the population of continuing applicants becomes increasingly relevant as the technological domain matures.

This finding is important because it shows that AI patenting is shaped by both widening and deepening dynamics. Widening is reflected in the continuous entry of new applicants and the rapid expansion of the cumulative applicant base. Deepening is reflected in the growing presence of applicants that remain active over time. The evidence therefore suggests that AI is neither a purely turbulent Mark I regime nor a purely cumulative Mark II regime. Instead, the field combines broad entry with increasing persistence.

Section 4.3 added a further layer to the analysis by examining the technological portfolios of the top 15 applicants. This firm-level perspective is important because aggregate indicators of concentration and entry do not reveal how leading applicants organize their technological capabilities. The portfolio analysis showed that leading firms are not homogeneous. Some applicants display broad technological portfolios, while others remain more focused around a limited number of IPC classes.

IBM, for example, displays the broadest portfolio, with 88 distinct IPC classes, followed by Boeing with 77, Ford Global Technologies with 64, and Intel with 52. Microsoft Technology Licensing LLC and Google LLC also show relatively broad portfolios, with 47 and 44 distinct IPC classes respectively. These applicants appear able to connect AI-related inventive activity to a wide range of technological domains. By contrast, other leading applicants, such as Adobe, Oracle, Capital One Services LLC, Bank of America and Apple, display more focused portfolios. This suggests that innovative leadership in AI does not follow a single strategy. Some firms lead through broad technological diversification, while others do so through more concentrated and specialized portfolios.

The portfolio HHI reinforces this interpretation. Higher values indicate that an applicant's IPC observations are concentrated in a limited set of technological classes, while lower values suggest a more dispersed portfolio. The evidence shows that some leading applicants, despite their large patenting activity, remain strongly concentrated around core computational classes. Others display more diversified portfolios, often reflecting sector-specific applications of AI in areas such as automotive technologies, robotics, automation or financial data processing. The core of the AI patenting system is therefore internally differentiated: leading applicants differ not only in the intensity of their patenting activity, but also in the structure and breadth of their technological knowledge.

The dynamic evidence on technological breadth further strengthens this result. From the mid-2010s onward, the average number of distinct IPC classes associated with leading applicants increases substantially. The average breadth of the top 15 applicants rises from 4.9 IPC classes in 2015 to 19.67 in 2023, before declining slightly to 18.07 in 2024. The top 5 applicants show an even stronger expansion, reaching 25.4 distinct IPC classes on average in 2022. This pattern suggests that leading applicants progressively broadened their technological portfolios as AI patenting expanded.

However, this broadening should not be interpreted as a complete dispersion of inventive activity. Even among the broadest portfolios, central IPC classes such as G06F and G06N remain highly recurrent. This means that technological exploration occurs around a relatively stable computational and machine-learning core. The evidence therefore points to a pattern of controlled diversification: leading applicants expand their technological scope, but they do so while maintaining a strong connection to central AI-related trajectories.

Overall, the empirical evidence developed in this chapter supports three main findings. First, AI-related patenting is broad in terms of participation but asymmetric in terms of output. Thousands of applicants enter the field, but a limited group accounts for a disproportionate share of inventive activity. Second, entry is continuous and substantial, but the growing role of continuing applicants suggests that cumulative participation becomes more relevant over time. Third, leading applicants differ in their technological portfolio strategies: some display broad and expanding portfolios, while others remain more specialized around selected IPC classes.

These findings support the interpretation of artificial intelligence as a hybrid Schumpeterian technological regime. The Mark I dimension is visible in the expansion of the applicant base, the continuous entry of new actors and the broadening of technological opportunities. The Mark II dimension is visible in the persistence of leading applicants, the concentration of inventive output, the growing importance of continuing applicants and the cumulative organization of technological portfolios around core AI-related classes.

The main conclusion of the chapter is therefore that AI patenting is neither fully open nor fully closed. It is open in the sense that many actors are able to enter and participate in the technological domain. At the same time, it is selective in the sense that only a smaller group of applicants accumulates significant inventive output, persists over time and develops broad or strategically focused technological portfolios. The expansion of the applicant base therefore broadens participation, but it does not automatically democratize innovative leadership.

This result is directly relevant for the broader argument of the thesis. It suggests that the growth of AI patenting should not be interpreted as a simple process of technological democratization. What matters is not only how many actors enter the field, but also whether they are able to remain active, accumulate inventive capacity and influence technological trajectories. From this perspective, the evidence points to a core–periphery structure: a large and expanding periphery of occasional or low-intensity innovators coexists with a more restricted core of persistent, technologically capable and strategically differentiated applicants.

The chapter therefore provides the empirical foundation for the final discussion. The evidence on concentration, entry and technological portfolios shows that artificial intelligence combines high technological opportunities, broad entry, cumulative dynamics and selective leadership. The following chapter builds on these results to assess the research questions and hypotheses formulated in the theoretical framework and to discuss the broader implications of interpreting AI as a hybrid Schumpeterian technological regime.

Chapter 5 – Discussion and Conclusion

5.1 Discussion of Research Questions and Hypotheses

The empirical analysis developed in the previous chapters makes it possible to return to the research questions and hypotheses formulated in the theoretical framework. The central question guiding the thesis was whether artificial intelligence should be interpreted as a Schumpeter Mark I technological regime, characterized by entry, experimentation and turbulence, as a Schumpeter Mark II technological regime, characterized by accumulation, persistence and concentration, or as a hybrid technological regime combining both dynamics.

The evidence presented in Chapters 3 and 4 suggests that artificial intelligence cannot be classified exclusively as either Mark I or Mark II. On the one hand, AI-related patenting displays features that are clearly consistent with widening dynamics. The technological domain expands rapidly over the period 2006–2024, the number of applicants increases substantially, and the distribution of IPC classes indicates a complex and re-combinatory knowledge base. These elements point to high technological opportunities and to the existence of broad spaces for experimentation. On the other hand, the same evidence also shows that innovative output remains highly asymmetric, that leading applicants account for a disproportionate share of patent mentions, that continuing applicants become increasingly important over time, and that technological portfolios remain anchored around core computational and AI-related classes. These elements point to deepening dynamics, cumulativeness and selective leadership.

The discussion of the hypotheses therefore requires a careful interpretation. The empirical results do not simply confirm a fully open and decentralized innovation regime. At the same time, they do not support the opposite interpretation of AI as a completely closed technological domain dominated only by incumbents. Rather, the results point to a stratified and hybrid structure, in which broad participation coexists with persistent asymmetries in innovative output, technological breadth and cumulative presence over time.

This section discusses the six hypotheses formulated in Chapter 1. Each hypothesis is assessed in relation to the empirical evidence presented in the previous chapters. The purpose is not to mechanically accept or reject each hypothesis, but to evaluate whether the results support, qualify or limit the theoretical expectations. This approach is consistent with the nature of the analysis, which is based on patent data and descriptive empirical indicators rather than on a causal econometric model.

H1 – Asymmetric opportunities

The first hypothesis stated that AI technological opportunities are high in aggregate terms, but that their effective contestability is asymmetric, varying across infrastructural and application segments. The empirical evidence is broadly consistent with this hypothesis.

The first element supporting H1 is the rapid expansion of AI-related patenting over time. The dataset includes 19,459 AI-related patents over the period 2006–2024. Patent activity remains relatively limited in the early years, but it accelerates markedly from the mid-2010s onward. This pattern indicates that artificial intelligence developed into a highly dynamic technological domain, characterized by increasing innovative activity and a growing number of technological trajectories. From the perspective of technological regimes, this expansion is consistent with high technological opportunities: firms and organizations appear to perceive AI as a field in which new inventions, applications and combinations can be developed.

The second element supporting H1 concerns the technological composition of AI patents. The dataset includes 95,172 patent–IPC observations, with an average of approximately 4.89 IPC classes per patent. This is an important result because it shows that AI inventions are rarely confined to a single narrow technological category. Instead, they are usually associated with multiple technological classes, suggesting that AI innovation relies on the combination of heterogeneous knowledge domains. The dominance of digital and computational IPC classes, especially G06F and G06N, confirms that AI is rooted in information processing, computation and machine-learning-related technologies. At the same time, the presence of several other IPC classes indicates that AI is not a single technological trajectory, but a broader technological space in which different knowledge components are recombined.

These results support the idea that AI offers broad technological opportunities. However, H1 does not state only that opportunities are high. It states that they are asymmetric in terms of effective contestability. The empirical evidence is consistent with this second part. The dataset includes 7,019 distinct applicants, which indicates broad participation in AI-related patenting. If one looked only at the number of applicants, AI could appear as a highly open and democratized technological domain. However, the distribution of inventive output shows a very different picture. The top 5 applicants account for 14.73% of total applicant mentions, the top 10 for 18.96%, and the top 20 for 25.66%. This means that a very small group of applicants accounts for more than one quarter of total applicant mentions.

The asymmetry becomes even clearer when considering the full applicant distribution. The majority of applicants appear only marginally in the dataset: 72.95% of applicants are associated with only one patent mention, while more than four fifths of the applicant population contributes only one or two patent mentions. This suggests that participation in AI patenting is broad, but the ability to produce a significant and repeated volume of inventive output is much more selective. In other words, many actors are able to enter the technological domain, but only a limited group appears capable of transforming technological opportunities into sustained patenting activity.

This distinction is central to the interpretation of AI as a technological regime. High technological opportunities do not automatically imply that all actors can exploit those opportunities in the same way. The empirical evidence suggests that opportunities are broad at the level of the technological domain, but uneven at the level of effective innovative capacity. Some firms and organizations are able to accumulate patents, operate across several technological classes and maintain a repeated presence over time. Many others enter only occasionally or remain peripheral.

Therefore, H1 is supported. Artificial intelligence displays high technological opportunities in aggregate terms, as shown by the expansion of patenting activity and the breadth of IPC classifications. However, these opportunities are asymmetric in their effective contestability. The ability to convert AI opportunities into persistent and substantial inventive output remains concentrated among a smaller group of actors. This supports the broader argument of the thesis: AI is open in terms of participation, but selective in terms of innovative leadership.

H2 – Reinforced cumulativeness in core segments

The second hypothesis stated that, in frontier and infrastructural AI segments, the innovation process is highly cumulative, generating path-dependent mechanisms that increase the likelihood of persistent technological leadership. The empirical evidence provides meaningful support for this hypothesis, although with some qualifications.

The first piece of evidence comes from the analysis of new and continuing applicants. The entry analysis shows that AI-related patenting is characterized by a continuous inflow of new applicants. However, it also shows that the number of continuing applicants increases substantially over time. Continuing applicants are those that appear in a given year after having already appeared in at least one previous year. Their growth indicates that AI patenting is not driven only by temporary or one-off entry. A growing share of the applicant population remains active after its first appearance.

This result is directly relevant to cumulativeness. In a technological regime characterized by low cumulativeness, one would expect entry and participation to be more unstable, with less evidence of repeated activity by the same actors. In contrast, the increasing number of continuing applicants suggests that previous participation matters. Firms and organizations that enter the AI patenting domain may accumulate knowledge, routines and capabilities that allow them to continue innovating in subsequent years. By 2024, the number of continuing applicants reaches 709, while the share of new applicants over total active applicants falls to 54.64%. This decline does not mean that entry disappears; the absolute number of new applicants remains high. Rather, it means that the stock of previously active applicants becomes increasingly important as the field matures.

The second piece of evidence comes from the concentration of inventive output. The presence of 7,019 distinct applicants suggests broad participation, but the distribution of patent mentions remains highly unequal. A small set of leading applicants accounts for a disproportionate share of patenting activity. This pattern is consistent with cumulative dynamics because repeated innovative activity tends to reinforce the position of already capable actors. Firms that have accumulated technological capabilities, research teams, organizational routines and complementary assets are better positioned to continue producing patentable inventions.

The third piece of evidence comes from technological portfolios. The analysis of the top 15 applicants shows that leading firms do not simply patent randomly across unrelated classes. Their portfolios remain strongly connected to core AI-related IPC classes such as G06F and G06N. Even when leading applicants broaden their technological scope, this expansion takes place around a stable computational and machine-learning core. This pattern can be interpreted as cumulative recombination. Firms expand into new IPC classes, but they do so by building on an established knowledge base.

This is especially visible in the evolution of technological breadth. From the mid-2010s onward, the average number of distinct IPC classes associated with the top 15 applicants increases substantially, rising from 4.9 in 2015 to 19.67 in 2023, before declining slightly to 18.07 in 2024. The top 5 applicants display an even stronger expansion, reaching 25.4 distinct IPC classes on average in 2022. This suggests that leading actors progressively broaden their technological portfolios as AI patenting expands. Such broadening is unlikely to be purely accidental. It reflects the capacity of leading firms to extend their inventive activity across a wider set of technological areas while maintaining a connection to core AI-related capabilities.

However, the support for H2 must be qualified. The evidence does not show a simple or linear increase in overall concentration. The yearly HHI is high in the earliest years, but those values must be interpreted with caution because the number of patents and applicants is still small. As the field expands, the HHI declines sharply and remains at lower levels for several years. This suggests that the expansion of AI patenting is accompanied by broader participation. Therefore, cumulativeness does not produce a complete closure of the technological domain.

The most accurate interpretation is that cumulativeness operates together with entry. AI patenting becomes more populated over time, but within this expanding field some actors become increasingly persistent and technologically capable. This means that the regime does not move from openness to closure in a simple way. Instead, widening and deepening develop simultaneously. New actors continue to enter, but previous participation and accumulated capabilities become increasingly relevant for sustained innovative activity.

For these reasons, H2 is supported with qualification. The evidence supports the existence of cumulative dynamics among continuing and leading applicants, especially in relation to repeated participation and the expansion of technological portfolios around core IPC classes. However, these cumulative dynamics coexist with broad entry and do not eliminate the openness of the technological domain.

H3 – Appropriability mediated by complementary assets

The third hypothesis stated that the capture of innovation rents in AI depends substantially on control over complementary assets such as data, infrastructures, platforms and distribution channels, reinforcing the positions of actors already present in core segments. The empirical analysis provides indirect but coherent support for this hypothesis.

The first point to clarify is methodological. The patent dataset used in this thesis does not directly measure complementary assets. It does not include firm-level information on proprietary datasets, cloud infrastructure, platform ownership, financial resources, distribution channels or access to specialized human capital. Therefore, H3 cannot be tested directly in a causal sense. However, the empirical evidence can still be used to evaluate whether the observed structure of AI patenting is consistent with the theoretical expectation that complementary assets matter for innovative leadership.

The composition of the leading applicants is consistent with this interpretation. The top applicants include large and technologically capable corporations such as IBM, Microsoft Technology Licensing LLC, Google LLC, Adobe, Intel, Amazon, Qualcomm, NVIDIA, Samsung, Apple, Cisco and Baidu, alongside large firms in finance, automotive and aerospace. These actors are not only patent-intensive firms. They are also organizations with significant research capabilities, technological infrastructures, financial resources, data-processing capacity and organizational routines. Their prominence in AI patenting is consistent with the idea that innovative leadership in AI depends not only on formal patents, but also on broader complementary assets.

This is particularly important because appropriability in AI differs from appropriability in more traditional technological contexts. In many fields, patents are a central mechanism through which firms protect inventions and capture returns from innovation. In AI, patents matter, but they are only one part of a broader appropriability system. Data, computation, platforms, cloud services, distribution channels and technical talent can all function as strategic assets that allow firms to transform technological opportunities into repeated inventive output. A firm with access to large datasets, advanced computational infrastructure and strong research teams is better positioned to develop, improve and patent AI-related technologies than a firm without these resources.

The concentration evidence is consistent with this interpretation. Although the applicant base is broad, inventive output remains asymmetric. The top 20 applicants account for 25.66% of total applicant mentions. This suggests that leading firms are able to capture a substantial share of AI-related inventive activity. Such concentration cannot be explained only by the existence of patents as formal intellectual property rights. It is more plausibly linked to the broader organizational and technological capabilities that allow leading firms to repeatedly generate patentable inventions.

The portfolio analysis reinforces this point. Some leading applicants display broad technological portfolios, indicating the capacity to connect AI-related innovation to multiple technological domains. IBM has 88 distinct IPC classes, Boeing 77, Ford Global Technologies 64, Intel 52, Microsoft Technology Licensing LLC 47 and Google LLC 44. A broad technological portfolio requires more than occasional participation. It suggests the ability to coordinate different knowledge areas, manage complex research activities and apply AI-related capabilities across several domains. These are precisely the types of capabilities that are associated with complementary assets.

At the same time, some leading applicants display more focused portfolios. Adobe, Oracle, Capital One Services LLC, Bank of America and Apple appear more concentrated in terms of distinct IPC classes. This does not contradict H3. Rather, it shows that complementary assets may support different strategic configurations. Some firms use their capabilities to diversify broadly across technological domains. Others use them to consolidate specialized positions around selected technological trajectories. In both cases, leading applicants are not simply random entrants; they are firms with the organizational capacity to sustain a recognizable technological strategy.

Therefore, H3 is supported indirectly. The empirical evidence does not directly observe complementary assets, but the identity of the leading applicants, the concentration of patent mentions and the structure of technological portfolios are all consistent with the idea that appropriability in AI is mediated by firm-level capabilities and strategic assets. The hypothesis should therefore be interpreted as theoretically supported and empirically consistent with the observed patterns, while recognizing that a direct test would require additional firm-level data.

H4 – Numerous but selective entry

The fourth hypothesis stated that AI generates high entry rates in application and downstream segments, but that only a limited share of entrants succeeds in influencing central technological trajectories or significantly altering concentration in core segments. The empirical evidence strongly supports this hypothesis.

The first part of H4 concerns the presence of numerous entries. This is clearly confirmed by the data. The cumulative number of distinct applicants increases from 11 in 2006 to 7,019 in 2024. This trajectory shows that AI patenting progressively attracted a very large population of actors. The field did not remain restricted to a small set of early innovators. Instead, year after year, new applicants entered the patenting domain, expanding the overall applicant base.

The annual entry data reinforce this interpretation. New applicants enter the field throughout the observation period. A major acceleration occurs from 2012 onward, when new applicants increase sharply. In the final years of the dataset, entry remains very strong: new applicants rise to 731 in 2022, 819 in 2023 and 854 in 2024. These figures show that AI remains attractive to new actors even in the later phase of the sample. Therefore, the technological domain cannot be described as closed to entry.

However, the second part of H4 is even more important: entry is selective. The presence of many applicants does not mean that all entrants have the same impact on the technological domain. The distribution of applicants by patent mentions shows that most entrants remain marginal. 72.95% of applicants appear only once in the dataset, and more than four fifths of all applicants are associated with only one or two patent mentions. This means that the majority of participants contribute only occasionally to AI-related inventive activity.

This evidence is central because it prevents a misleading interpretation of entry. A high number of entrants could be interpreted as evidence of technological democratization. However, the applicant distribution shows that broad entry does not necessarily translate into broad innovative leadership. Most entrants do not accumulate a significant stock of patents and do not appear as persistent actors within the patenting domain. Entry is therefore numerous, but much of it is limited in intensity.

The concentration ratios further support this conclusion. Despite the large number of applicants, the top 5, top 10 and top 20 applicants account for substantial shares of total applicant mentions. In particular, the top 20 applicants account for more than one quarter of total applicant mentions. This indicates that new entry does not eliminate the disproportionate weight of leading actors. The structure of the field remains hierarchical: a large number of applicants participate, but a smaller group has a much stronger role in shaping the distribution of inventive output.

The distinction between new and continuing applicants also supports H4. While new entry remains strong, the number of continuing applicants increases over time. This suggests that the field is not only characterized by entry, but also by selective persistence. Some applicants remain active after entering, while many others appear only once or sporadically. The ability to continue patenting therefore becomes a key indicator of whether entry is merely formal or whether it reflects a more stable innovative presence.

Therefore, H4 is strongly supported. AI generates broad entry, but this entry is selective in its consequences. Many actors enter the patenting domain, but only a limited subset becomes persistent, patent-intensive or technologically influential. This result is one of the strongest empirical findings of the thesis and directly supports the interpretation of AI as a hybrid technological regime.

H5 – Persistent core–periphery structure

The fifth hypothesis stated that the evolution of industrial concentration in AI tends to assume a core–periphery configuration, with high concentration and stability in the infrastructural core and greater competitive turbulence in the application periphery. The empirical results support this hypothesis, although the patent data capture the core–periphery structure indirectly rather than through a direct mapping of the AI value chain.

The core–periphery structure is visible first in the applicant distribution. The periphery is represented by the large number of applicants that appear only once or a few times. These actors confirm that participation in AI-related patenting is widespread. They represent the broad and turbulent side of the regime: many firms and organizations enter the field, experiment with AI-related inventions and contribute to the expansion of the applicant base. However, their individual weight in total inventive output remains limited.

The core is represented by the smaller group of applicants that accounts for a disproportionate share of patent mentions. IBM is the leading applicant, followed by Microsoft Technology Licensing LLC, Google LLC, Adobe, Intel and other large firms. The top 20 applicants account for 25.66% of total applicant mentions. This is a substantial share considering the size of the applicant population. The empirical evidence therefore points to a structure in which a large periphery of low-intensity applicants coexists with a more restricted core of high-intensity innovators.

The entry evidence reinforces this interpretation. The cumulative number of applicants grows continuously, showing expansion at the periphery. At the same time, the number of continuing applicants increases, showing that part of the applicant population becomes more persistent. This means that the field is not simply renewed every year by completely different actors. Some applicants consolidate their presence over time. This consolidation is consistent with the formation of a core of more stable participants.

The technological portfolio analysis adds an important qualification to H5. The core is not homogeneous. Leading applicants differ substantially in technological breadth and focus. IBM displays the broadest portfolio, with 88 distinct IPC classes, followed by Boeing with 77 and Ford Global Technologies with 64. Other firms, such as Adobe, Oracle, Capital One Services LLC, Bank of America and Apple, display more focused portfolios. This indicates that the core of AI patenting includes different types of leading actors: some pursue broader technological diversification, while others concentrate on narrower technological trajectories.

This internal differentiation matters because it prevents an overly simple view of concentration. The core is not merely a group of identical dominant firms. It is composed of actors with different sectoral backgrounds, capabilities and technological strategies. Some are large digital technology firms, others are semiconductor firms, financial institutions, automotive firms or aerospace firms. Their common feature is not that they innovate in exactly the same way, but that they are able to occupy a more significant position than the majority of applicants in the distribution of AI patenting.

Therefore, H5 is supported with qualification. AI-related patenting displays a core–periphery structure in which broad participation at the periphery coexists with concentrated and more persistent leadership at the core. The evidence also qualifies the hypothesis by showing that the core itself is internally differentiated, with leading applicants adopting different technological portfolio strategies.

H6 – Strengthening of deepening dynamics over time

The sixth hypothesis stated that, in the medium term, the intensification of scale economies in computing, data accumulation and organizational complementarities may reinforce creative accumulation dynamics in core segments, increasing relative concentration compared to the initial diffusion phase. The empirical evidence provides partial and qualified support for this hypothesis.

Several results support the idea that deepening dynamics become more relevant over time. First, the number of continuing applicants increases substantially. This suggests that repeated participation becomes increasingly important as the AI patenting domain matures. The share of new applicants over total active applicants declines in the later years, falling to 54.64% in 2024. This does not mean that entry weakens in absolute terms. On the contrary, the number of new applicants remains very high. However, it does indicate that the stock of continuing applicants becomes more relevant. This is consistent with the idea that cumulative participation strengthens over time.

Second, the technological breadth of leading applicants expands over time. The average number of distinct IPC classes associated with the top 15 applicants rises from 4.9 in 2015 to 19.67 in 2023. The top 5 applicants show an even stronger increase, reaching 25.4 distinct IPC classes on average in 2022. This suggests that leading firms become increasingly capable of extending AI-related innovation across multiple technological domains. Such expansion can be interpreted as a form of deepening because it reflects the accumulation of technological capabilities and the ability to recombine them across a broader portfolio.

Third, despite portfolio expansion, leading applicants remain connected to core IPC classes such as G06F and G06N. This indicates that diversification does not imply a complete dispersion of inventive activity. Instead, leading firms broaden their technological scope while maintaining a stable computational and machine-learning core. This pattern is consistent with cumulative technological development: firms build new inventive trajectories around an existing knowledge base.

However, the hypothesis must be qualified because the concentration evidence does not show a simple monotonic increase in overall concentration. The HHI is relatively high in the earliest years, but this is partly due to the small number of patents and applicants at the beginning of the period. As the field expands, the HHI declines sharply. Later, it fluctuates, with some increases in specific years, but it does not produce a straightforward pattern of continuous concentration. Therefore, it would be inaccurate to claim that AI patenting simply becomes more concentrated over time in aggregate terms.

The better interpretation is that deepening dynamics strengthen within an expanding field. AI patenting becomes broader and more populated, but within this larger population a group of continuing and leading applicants consolidates technological capabilities. Thus, deepening does not replace widening. It develops alongside it. This is exactly what makes AI a hybrid technological regime: the field expands through entry, but accumulated capabilities still matter for persistence and leadership.

For this reason, H6 is partially supported. The evidence supports the strengthening of cumulative and deepening dynamics, especially through the growth of continuing applicants and the expansion of leading firms' technological portfolios. However, the evidence does not support a simple view of increasing aggregate concentration over time. The most accurate conclusion is that AI combines the widening of participation with the deepening of capabilities among leading and continuing applicants.

Overall assessment of the hypotheses

Taken together, the discussion of the hypotheses strengthens the central argument of the thesis. Artificial intelligence is best interpreted as a hybrid Schumpeterian technological regime. The evidence does not support a purely Mark I interpretation, because entry and experimentation coexist with strong asymmetries in inventive output and with the persistent relevance of leading applicants. At the same time, the evidence does not support a purely Mark II interpretation, because the technological domain remains open to a large and growing number of entrants.

The strongest support is found for H1, H4 and H5. H1 is supported because AI displays high technological opportunities, but these opportunities are unevenly translated into significant patenting activity. H4 is strongly supported because entry is numerous but selective: many applicants enter the field, but most remain marginal. H5 is supported because the data reveal a core–periphery structure, with a broad periphery of occasional applicants and a more restricted core of patent-intensive and technologically capable actors.

H2 and H6 are also supported, but with qualifications. The evidence shows cumulative dynamics through the growth of continuing applicants and the technological portfolio expansion of leading firms. However, these dynamics do not eliminate entry and do not lead to a simple linear increase in concentration. Cumulativeness and deepening are present, but they coexist with widening. H3 is supported indirectly: the identity and structure of leading applicants are consistent with the importance of complementary assets, although the dataset does not directly measure data ownership, computational infrastructure or platform control.

The overall conclusion is that AI patenting is open in terms of participation, but selective in terms of innovative output, persistence and technological scope. This distinction is crucial. The growth of the applicant base should not be interpreted as automatic democratization of innovation. What matters is not only how many actors enter the field, but also whether they remain active, accumulate inventive capacity and shape technological trajectories.

The empirical evidence therefore supports a layered interpretation of AI. At the periphery, many actors enter and experiment. At the core, a smaller group of applicants accumulates patenting activity, remains visible over time and organizes technological portfolios around central AI-related classes. This structure captures the coexistence of Mark I and Mark II dynamics and provides the basis for the broader discussion of theoretical and empirical implications developed in the following sections.

5.2 Artificial Intelligence as a Hybrid Schumpeterian Technological Regime

The discussion of the hypotheses leads to the central interpretative result of the thesis: artificial intelligence is best understood as a hybrid Schumpeterian technological regime. This conclusion emerges from the coexistence of two empirical patterns that, if considered separately, could lead to opposite interpretations. On the one hand, AI-related patenting displays clear widening dynamics: the number of patents grows substantially, the applicant base expands continuously, new actors enter the field every year, and the technological knowledge base is broad and combinatorial. On the other hand, the same domain also displays deepening dynamics: inventive output remains strongly asymmetric, leading applicants account for a disproportionate share of patent mentions, continuing applicants become increasingly relevant over time, and the technological portfolios of leading firms remain anchored around core computational classes.

The central point is that these two dynamics are not contradictory. Rather, they are part of the same technological configuration. Artificial intelligence expands the space for participation, but it does not distribute innovative leadership evenly. This is the main reason why AI cannot be interpreted as a purely democratizing technology. The diffusion of AI tools, libraries, pre-trained models, APIs and cloud services allows a growing number of actors to enter the technological domain. However, entry alone does not imply the capacity to persist, accumulate inventive output, or influence the main technological trajectories. In this sense, AI appears open at the margin but selective at the core.

A purely Schumpeter Mark I interpretation would therefore be incomplete. In a Mark I regime, innovation is typically associated with entry, experimentation, turbulence and the continuous reshuffling of technological hierarchies. Some of the empirical evidence clearly points in this direction. The cumulative number of distinct applicants increases from 11 in 2006 to 7,019 in 2024, showing that the AI patenting domain attracted a large and growing population of actors. New applicants continue to enter in large numbers even in the later years of the period, reaching 854 in 2024. This confirms that AI creates broad opportunities for participation and experimentation.

However, this openness should not be overstated. The distribution of inventive output shows that most applicants remain marginal. A very large share of applicants appears only once in the dataset, and more than four fifths of the applicant population contributes only one or two patent mentions. This evidence changes the interpretation of entry. Entry exists, and it is quantitatively important, but much of it is occasional or low-intensity. Therefore, the AI regime does not correspond to a fully decentralized Mark I pattern in which new entrants continuously overturn the hierarchy of innovators. Many actors enter, but only a much smaller group accumulates a relevant innovative presence.

This distinction between entry and transformative entry is one of the most important interpretative results of the thesis. Counting applicants is not enough to understand the structure of AI innovation. A large number of entrants may create the appearance of openness, but the real question is whether these entrants persist, grow and become capable of shaping technological trajectories. The evidence suggests that most entrants remain peripheral. Therefore, AI broadens participation, but this broadening does not automatically translate into a redistribution of innovative leadership.

At the same time, a purely Schumpeter Mark II interpretation would also be incomplete. In a Mark II regime, innovation is usually associated with large firms, accumulated capabilities, persistence and concentration. Several results are consistent with this view. The top 20 applicants account for 25.66% of total applicant mentions, while the top 5 account for 14.73%. Leading firms such as IBM, Microsoft Technology Licensing LLC, Google, Adobe, Intel, Amazon, NVIDIA, Qualcomm, Samsung and Apple occupy a prominent position in the distribution of AI-related patenting. The increasing number of continuing applicants, which reaches 709 in 2024, also indicates that repeated participation becomes more important as the field matures.

Nevertheless, the AI patenting domain does not become closed. The applicant base continues to expand throughout the observation period, and new applicants remain numerous. The HHI does not show a simple and continuous increase in concentration. After the early years, when values are influenced by the small number of patents and applicants, concentration declines substantially as the field expands. This means that AI does not evolve into a fully closed regime dominated only by a fixed set of incumbents. Instead, it remains open in terms of entry while becoming selective in terms of persistence and intensity of innovative output.

The hybrid interpretation captures this combination better than either Mark I or Mark II alone. AI is characterized by widening without full de-concentration and by deepening without full closure. Widening is visible in the expansion of patenting activity, the continuous arrival of new applicants, and the presence of a complex technological knowledge base. Deepening is visible in the concentration of applicant mentions, the growing role of continuing applicants, and the ability of leading firms to expand their technological portfolios around core IPC classes. These dynamics develop together rather than replacing one another.

This interpretation is important because it challenges a common assumption in debates on digital technologies: the idea that greater accessibility necessarily leads to a more democratic innovation structure. In AI, accessibility and concentration coexist. Many firms can access AI tools and experiment with AI applications, but only some actors have the resources needed to accumulate a substantial and persistent inventive position. In my interpretation, this is the key tension of the AI regime: the technology lowers some barriers to participation while simultaneously raising the requirements for leadership.

The core–periphery structure provides a useful way to describe this tension. The periphery consists of the many applicants that appear only once or a few times in the dataset. These actors reflect the openness of the AI domain. They contribute to experimentation, variation and the diffusion of AI-related techniques across sectors. Their presence matters because it shows that AI is not a closed technological field accessible only to a few dominant firms. However, the periphery remains weak in terms of cumulative inventive output. Its contribution is broad but fragmented.

The core consists of the leading applicants that account for a disproportionate share of patent mentions and display stronger technological portfolios. These actors are not only more active in quantitative terms; they also appear more capable of organizing AI-related innovation across multiple technological domains. IBM, for example, displays the broadest technological portfolio among the top applicants, with 88 distinct IPC classes. Boeing, Ford Global Technologies, Intel, Microsoft Technology Licensing LLC and Google also show relatively broad portfolios. This suggests that leading actors are able to connect AI-related inventive activity to a wide range of technological areas.

However, the core itself is not homogeneous. This is an important qualification. Some leading firms display broad and diversified technological portfolios, while others remain more focused around selected IPC classes. Therefore, innovative leadership in AI does not follow a single model. Some firms lead through technological breadth, applying AI-related knowledge across multiple domains. Others lead through specialization, concentrating inventive activity around narrower but strategically important trajectories. This means that the core of the AI patenting system is internally differentiated. It is not simply a group of identical dominant firms, but a set of actors with different technological positions, sectoral backgrounds and portfolio strategies.

This result adds an important firm-level dimension to the technological regimes framework. Technological regimes are often discussed at the level of sectors or technological classes. However, the evidence from AI shows that even within the same regime, leading firms may adopt different strategies. This matters because it prevents an overly mechanical interpretation of concentration. Concentration does not necessarily mean that all leading firms behave in the same way. In the AI regime, some leaders appear more exploratory, while others appear more focused. The structure is concentrated, but not uniform.

The technological portfolio evidence also clarifies the nature of diversification in AI. The increase in technological breadth among leading applicants does not imply random dispersion across unrelated fields. Even when portfolios become broader, they remain strongly connected to core computational and AI-related classes such as G06F and G06N. This suggests a pattern of controlled diversification. Leading firms expand the scope of their inventive activity, but they do so by building on a stable technological foundation. In this sense, AI innovation develops through cumulative recombination: existing capabilities are extended into adjacent and complementary technological domains.

This point is central for understanding the knowledge-base dimension of the AI regime. The dataset includes 95,172 patent–IPC observations, and each patent is associated on average with approximately 4.89 IPC classes. This indicates that AI inventions often combine multiple technological categories. AI is therefore not a narrow technological field, but a complex and interdisciplinary domain built around computation, machine learning, data processing, language processing, computer vision and related areas. However, the ability to exploit this combinatorial knowledge base is not evenly distributed. The broadest and most structured portfolios are observed among leading applicants, suggesting that technological recombination itself may become a source of advantage.

This is a crucial point. A complex knowledge base can create opportunities for many actors, because it generates multiple possible applications and combinations. But complexity can also reinforce the advantage of firms with stronger capabilities, because recombination across domains requires organizational coordination, technical expertise and resources. Therefore, the complexity of AI does not automatically democratize innovation. It creates opportunities, but it also rewards firms that are able to manage complexity at scale.

The hybrid regime interpretation also helps clarify the role of appropriability. In artificial intelligence, appropriability cannot be reduced to patent ownership alone. Patents are important because they provide an observable measure of inventive activity, but the capacity to generate and exploit AI-related inventions depends on broader complementary assets. These include data, computational infrastructure, cloud platforms, specialized human capital, research capabilities and access to distribution channels. The prominence of large technology firms and other major corporations among leading applicants is consistent with this interpretation.

In my view, this is one of the reasons why AI produces a mixed structure. The technology is accessible enough to allow many actors to participate, but demanding enough to make leadership depend on assets that are unevenly distributed. This explains why the same technological domain can generate thousands of applicants and, at the same time, maintain a concentrated core of more capable actors. The contradiction is only apparent. AI is open at the level of use and experimentation, but much more selective at the level of cumulative inventive leadership.

This interpretation also has implications for how the evolution of the AI regime should be understood over time. In the early years, the field is small, and concentration indicators are affected by the limited number of patents and applicants. As AI patenting expands, participation increases and the HHI declines, indicating broader involvement. In the later years, however, continuing applicants become more important and leading firms expand their technological portfolios. This suggests that the AI regime evolves through both expansion and consolidation. The field becomes larger and more populated, but within this larger population some actors accumulate capabilities and become more persistent.

The temporal pattern therefore does not suggest a simple transition from Mark I to Mark II. It suggests the coexistence of Mark I and Mark II dynamics over time. In the expansion phase, AI generates entry and experimentation. As the regime matures, some applicants consolidate their position and deepen their technological capabilities. But this deepening does not eliminate new entry. The result is not replacement of one regime by another, but the simultaneous presence of widening and deepening.

This is why the concept of hybrid regime is not simply a compromise label. It is analytically useful because it captures the actual structure observed in the data. AI patenting is neither fully decentralized nor fully concentrated. It is neither purely turbulent nor fully stable. It is neither open in all dimensions nor closed in all dimensions. Its defining feature is the coexistence of different dynamics across different levels of the innovation system: broad participation at the periphery, selective accumulation at the core, and heterogeneous strategies among leading firms.

This interpretation also modifies the way AI should be understood as a general-purpose technology. The GPT nature of AI helps explain its wide diffusion and the continuous emergence of new applications. However, GPT status alone does not determine the industrial structure of innovation. A technology can be general-purpose and still generate concentration if the capabilities required to lead innovation are unevenly distributed. In the case of AI, general-purpose applicability supports broad entry, while cumulative capabilities and complementary assets support selective leadership. Therefore, AI's GPT character reinforces the hybrid nature of the regime rather than resolving it in favor of openness.

The main interpretative contribution of the thesis is therefore to move beyond two simplified views. The first view is that AI democratizes innovation simply because many actors can access tools and develop applications. The second view is that AI is entirely captured by large incumbents because frontier development requires scale, data and computation. The empirical evidence suggests that both views capture part of the truth, but neither is sufficient. AI democratizes participation more than leadership. It allows many actors to enter, but it does not give all actors the same capacity to persist, accumulate and shape technological trajectories.

This distinction is important also from a broader industrial dynamics perspective. If entry is broad but mostly marginal, then the number of entrants alone is not a reliable indicator of competitive transformation. What matters is whether entry becomes persistent, whether entrants accumulate capabilities, and whether they can move from peripheral participation to more central positions. The evidence in this thesis suggests that such movement is selective. This supports the idea that AI innovation is characterized by a large experimental periphery and a more restricted cumulative core.

Overall, the empirical evidence supports the interpretation of artificial intelligence as a hybrid Schumpeterian technological regime. AI-related patenting expands participation, but it does not distribute innovative leadership evenly. It creates opportunities for many actors to enter, but only a smaller group accumulates substantial patenting activity, persists over time and develops broad or strategically focused technological portfolios. The regime is therefore open in terms of participation, but selective in terms of innovative output, persistence and technological scope.

This conclusion provides the basis for the broader implications discussed in the next section. If AI is a hybrid regime, then both theory and empirical analysis must avoid one-dimensional interpretations. The key issue is not whether AI is open or concentrated in absolute terms, but how openness and concentration coexist, which actors benefit from this coexistence, and under what conditions participation can become persistent and transformative innovation.

5.3 Implications, Limitations and Future Research

The interpretation of artificial intelligence as a hybrid Schumpeterian technological regime has several theoretical and empirical implications. The main implication is that AI should not be understood through a simple opposition between openness and concentration. The evidence developed in this thesis shows that both dimensions coexist. On the one hand, AI-related patenting expands rapidly, attracts a large number of applicants and generates continuous entry. On the other hand, inventive output remains asymmetric, leading applicants account for a disproportionate share of patent mentions, and technological capabilities are unevenly distributed across firms. This means that the growth of AI innovation should not be automatically interpreted as a democratization of innovative leadership.

This point is central. In many discussions on artificial intelligence, the diffusion of tools, models, APIs and open-source libraries is often interpreted as evidence that innovation is becoming more accessible. This interpretation is partly correct, but incomplete. The evidence presented in this thesis suggests that AI does lower some barriers to participation, especially at the level of experimentation and application. However, participation is not the same as leadership. Many actors are able to enter the AI patenting domain, but only a smaller group accumulates substantial inventive output, persists over time and develops broad or strategically focused technological portfolios. In my interpretation, this is one of the most important features of the AI regime: AI democratizes access more than it democratizes innovative power.

From a theoretical perspective, this result confirms the usefulness of the technological regimes framework for studying artificial intelligence. The framework allows the analysis to move beyond general descriptions of AI as a disruptive technology or as a general-purpose technology. Instead, it makes it possible to examine how innovation is organized, who is able to innovate repeatedly, and how inventive activity is distributed across actors. The four dimensions of technological regimes — technological opportunities, appropriability, cumulativeness and knowledge base — help explain why AI can be simultaneously open and selective. High technological opportunities support entry and experimentation, while cumulativeness and appropriability mechanisms reinforce the position of actors that already possess stronger technological, organizational and financial capabilities.

The thesis also suggests that the traditional distinction between Schumpeter Mark I and Schumpeter Mark II should be applied carefully to artificial intelligence. These categories remain analytically useful, but they are not sufficient if used as rigid alternatives. AI does not fit neatly into either model. Mark I dynamics are visible in the continuous entry of new applicants, the expansion of the applicant base and the variety of technological applications. Mark II dynamics are visible in the concentration of inventive output, the growing importance of continuing applicants and the persistence of core technological trajectories around IPC classes such as G06F and G06N. The empirical evidence therefore supports a hybrid interpretation, in which widening and deepening dynamics develop together.

This has an important implication for the literature on technological regimes. The case of AI suggests that technological regimes should not be interpreted as fixed or uniform structures. A regime can contain different dynamics at the same time. It can be open at the periphery and cumulative at the core. It can generate many entrants while still preserving the advantages of leading actors. It can encourage exploration across technological classes while maintaining a stable computational foundation. Therefore, the hybrid regime interpretation is not simply a compromise between Mark I and Mark II. It is a more accurate way to describe a technological domain in which entry, accumulation, experimentation and concentration coexist.

A second theoretical implication concerns the distinction between entry and transformative entry. The evidence shows that many actors enter AI-related patenting, but most applicants appear only marginally. This distinction is crucial because it challenges the idea that the number of entrants alone is enough to assess the openness of a technological regime. A regime may be open in terms of participation while remaining selective in terms of persistence and influence. In the case of AI, the large number of applicants confirms broad participation, but the long-tail distribution shows that most actors do not accumulate a significant volume of patenting activity. Therefore, in my view, the relevant question is not only whether firms enter AI, but whether they are able to remain active and influence technological trajectories.

This result is also relevant for the interpretation of core–periphery structures in digital technologies. The AI patenting domain appears to be characterized by a broad periphery of occasional or low-intensity applicants and a more restricted core of persistent and technologically capable actors. However, the core itself is not homogeneous. The portfolio analysis shows that leading applicants follow different technological strategies. Some display broad portfolios, while others remain focused around a smaller set of IPC classes. This means that concentration should not be interpreted only as the dominance of a few identical large firms. It also reflects heterogeneous firm-level strategies and different ways of organizing technological capabilities.

This point adds a useful qualification to the technological regimes framework. Aggregate indicators are necessary, but they can hide important differences across firms. If the analysis stopped at concentration ratios or HHI, it would show that patenting activity is asymmetric, but it would not explain how leading firms differ from each other. The portfolio analysis shows that some leaders are broad technological explorers, while others are more focused specialists. This is important because innovative leadership is not only about producing many patents. It is also about how firms structure their technological knowledge, how they recombine different domains, and how they position themselves within the broader AI technological system.

From an empirical perspective, the thesis shows the importance of combining multiple indicators rather than relying on a single measure. Patent growth alone would suggest technological expansion. The number of applicants alone would suggest broad participation. Concentration ratios and applicant distributions, however, reveal strong asymmetries in inventive output. Entry indicators show the coexistence of new and continuing applicants. Technological portfolio measures add a firm-level perspective by showing whether leading applicants are broad or focused in terms of IPC classes. Taken together, these indicators provide a much richer interpretation of the AI technological regime than any single measure could offer.

This empirical implication is important because AI is often discussed using very broad indicators, such as the number of patents, investments or firms entering the field. These indicators are useful, but they can easily produce misleading conclusions if interpreted in isolation. For example, a growing number of applicants could suggest democratization, while concentration ratios show that leadership remains selective. Similarly, a decline in HHI may suggest dispersion, while the growth of continuing applicants and the expansion of top firms' portfolios show that cumulative dynamics are still relevant. The contribution of this thesis is therefore to show that the structure of AI innovation can only be understood by looking jointly at growth, concentration, entry, persistence and technological breadth.

The findings also have implications for policy and industrial strategy. If AI is open in terms of participation but selective in terms of leadership, then policies aimed only at increasing the number of entrants may not be sufficient. Supporting entry is important, but the evidence suggests that persistence and capability accumulation matter just as much. Firms and organizations may need access not only to AI tools, but also to data, computational infrastructure, skilled human capital, research capabilities and organizational resources. Without these assets, entry may remain occasional and peripheral.

This interpretation has a practical implication: lowering entry barriers is not enough if the objective is to broaden innovative leadership. A policy that only helps firms adopt AI tools may increase participation, but it may not allow them to become persistent innovators. To move from participation to leadership, firms need deeper capabilities. They need to accumulate technical knowledge, access relevant data, develop internal research capacity and connect AI to their specific technological or sectoral domains. Therefore, the real policy challenge is not simply to make AI accessible, but to make sustained AI innovation possible for a broader set of actors.

At the same time, the evidence suggests that concentration in AI should not be interpreted only negatively. Some degree of concentration may reflect the complexity of the knowledge base and the cumulative nature of technological development. Large and persistent actors may be better able to coordinate complex research activities, invest in infrastructure and operate across multiple technological domains. From this perspective, concentration may partly reflect the economic and technological requirements of AI innovation. The critical issue is therefore not whether concentration exists, but whether concentrated leadership prevents broader diffusion, limits entry into core segments, or blocks the emergence of new technologically capable actors.

This is an important nuance. A simplistic interpretation would treat concentration as automatically harmful and entry as automatically positive. The evidence in this thesis suggests a more careful interpretation. Entry is positive because it increases variety and experimentation, but much entry remains marginal. Concentration can reflect barriers and unequal access to resources, but it can also reflect the need for scale, cumulative learning and coordination in complex technological domains. Therefore, the key issue is the balance between openness and concentration: whether the AI regime allows peripheral actors to move toward more persistent and capable forms of innovation, or whether the core becomes increasingly difficult to contest.

The findings also suggest that AI innovation is layered. The patent data used in this thesis do not perfectly distinguish infrastructural, model-level and application-level innovation, but the coexistence of broad entry and concentrated leadership is consistent with a layered structure. Entry may be easier in downstream or application-oriented activities, while leadership in core technological trajectories may require more cumulative capabilities and complementary assets. This layered interpretation helps explain why AI can simultaneously generate widespread experimentation and persistent concentration among leading innovators.

Despite these contributions, the analysis has several limitations. The first limitation concerns the use of patent data. Patents are a useful and systematic proxy for inventive activity, but they do not capture all forms of innovation. Some AI-related innovation may occur through software development, open-source contributions, trade secrets, model releases, publications, internal organizational changes or informal knowledge accumulation that are not patented. As a result, patent data may over-represent firms with stronger intellectual property strategies and under-represent actors that innovate without relying heavily on patents.

This limitation is particularly relevant in artificial intelligence because AI innovation often occurs through software, data, models and organizational deployment. Not all of these outputs are easily captured by patents. For example, an open-source model, an internal algorithmic system, or a data-processing architecture may be economically important even if it does not appear in patent data. Therefore, the analysis should be interpreted as a study of AI-related inventive activity as captured by patents, not as a complete measurement of all AI innovation.

A second limitation concerns differences in patenting propensity across firms and sectors. Some firms rely heavily on patents to protect inventions, while others rely more on secrecy, speed of innovation, open-source strategies or control over platforms and data. Therefore, patent counts should not be interpreted as a perfect measure of innovative capability. They capture one important dimension of inventive activity, but not the full range of AI innovation strategies. This matters because firms with similar innovative capabilities may appear differently in the data depending on their intellectual property strategy.

A third limitation concerns the geographical structure of the dataset. The empirical analysis is strongly focused on patenting activity associated with the United States. This provides a coherent setting for studying AI innovation at the technological frontier, but it limits the generalizability of the results to the global AI system. Other regions, especially Europe and China, may display different patterns of entry, concentration, applicant composition and technological specialization. A more global analysis could provide a broader comparison of AI technological regimes across institutional and regional contexts.

A fourth limitation concerns applicant-name harmonization. Although the dataset was cleaned and applicant names were standardized, some residual problems may remain. Large firms often patent through subsidiaries, licensing entities or slightly different legal names. This can affect the measurement of concentration, the identification of leading applicants and the construction of technological portfolios. The use of applicant mentions rather than unique corporate groups is appropriate for the available data, but it should still be interpreted with caution.

A fifth limitation concerns the use of IPC classes as proxies for technological portfolios. IPC classifications are useful because they provide a standardized way to identify technological domains, but they do not fully capture the economic layers of AI innovation. A patent classified in a computational IPC class may relate to different stages of the AI value chain, including infrastructure, models, applications or embedded systems. Therefore, IPC-based measures are informative about technological breadth and knowledge composition, but they cannot perfectly distinguish business models or strategic positions within the AI ecosystem.

This limitation is especially important for the interpretation of the core–periphery structure. The thesis identifies a core–periphery structure through patenting intensity, applicant persistence and technological portfolios. However, the data do not directly identify whether a patent belongs to the infrastructure layer, the model layer or the application layer of AI. Therefore, the empirical core–periphery structure observed in the thesis should be interpreted as a patent-based approximation of the broader economic and technological structure of AI, not as a complete mapping of the entire AI value chain.

A final limitation is that the analysis is descriptive and interpretative rather than causal. The thesis identifies patterns of concentration, entry, persistence and technological portfolio strategies, but it does not estimate causal relationships among these variables. For example, the analysis cannot prove that complementary assets cause higher patenting intensity, or that broader IPC portfolios directly lead to persistent leadership. Instead, it shows that the observed empirical patterns are consistent with the interpretation of AI as a hybrid technological regime. A causal analysis would require additional firm-level variables and a different empirical design.

These limitations do not invalidate the results, but they define the boundaries of the contribution. The thesis does not claim to provide a complete causal explanation of AI industrial dynamics. Rather, it provides an empirical and interpretative framework for understanding the structure of AI patenting. In my view, this is still valuable because before testing causal mechanisms it is necessary to clarify what kind of empirical structure exists. The evidence shows that AI patenting is neither purely open nor purely concentrated, and this descriptive result is itself important for guiding future research.

Several directions for future research follow from these limitations. First, future studies could integrate patent data with firm-level information on R&D expenditure, employment, revenues, cloud infrastructure, data access and platform ownership. This would make it possible to test more directly the role of complementary assets in shaping AI innovation and concentration. In particular, it would allow researchers to examine whether firms with stronger data resources or computational infrastructure are more likely to persist in AI patenting or to develop broader technological portfolios.

Second, future research could distinguish more explicitly among different layers of the AI value chain. Separating infrastructure, foundation models, applications and sector-specific uses would make it possible to verify whether entry is indeed broader in downstream activities and more restricted in core technological segments. This would directly address one of the main theoretical claims of the thesis: that AI is open at the periphery but more cumulative and selective at the core.

Third, future research could compare different geographical areas. A comparative analysis of the United States, Europe and China would be particularly useful for understanding whether the hybrid regime identified in this thesis is a general feature of AI or whether it varies across institutional and regional contexts. It is possible that the balance between entry, concentration and state involvement differs significantly across these regions. Such a comparison would help clarify whether the AI regime is globally hybrid or whether different regional regimes coexist.

Fourth, future studies could combine patent data with other sources of innovation data, including scientific publications, open-source repositories, startup funding, merger and acquisition activity and information on foundation models. This would provide a more complete picture of AI innovation beyond patenting. Such an approach would be especially useful because AI innovation often develops outside the patent system, through scientific research, open-source communities and platform-based ecosystems.

Fifth, future research could develop network-based analyses of applicants and technological classes. Applicant–IPC networks could show which firms connect different technological domains and whether some actors occupy more central positions in the technological knowledge structure. Co-applicant networks could also reveal collaboration patterns and the role of universities, research institutes and firms in the development of AI technologies. These approaches would deepen the understanding of AI as a complex and interconnected technological system.

Finally, future work could develop econometric analyses based on richer data. This thesis deliberately focuses on descriptive and interpretative evidence because the objective is to assess the structure of the AI technological regime rather than to estimate a narrow causal effect. However, with additional variables, future research could test whether firm size, R&D investment, prior patenting experience, access to complementary assets or technological breadth predict persistence and leadership in AI patenting. Such an analysis would provide a more formal test of the mechanisms suggested by the present thesis.

Overall, the implications and limitations of the analysis reinforce the main conclusion rather than weakening it. The evidence shows that artificial intelligence expands participation while maintaining selective and cumulative forms of leadership. This result is robust across the different empirical indicators used in the thesis, even if future research can refine the analysis with broader data and more detailed measures. The central contribution remains the identification of AI as a technological regime in which openness and concentration, entry and persistence, exploration and cumulativeness coexist.

5.4 Concluding remarks

This thesis set out to examine artificial intelligence not simply as a fast-growing technology, but as a technological regime with specific implications for entry, concentration and industrial dynamics. The central objective was to understand whether AI-related innovation follows a Schumpeter Mark I pattern of widening, entry and experimentation, a Schumpeter Mark II pattern of accumulation, persistence and concentration, or a hybrid configuration combining both dynamics.

The empirical evidence supports the hybrid interpretation. AI-related patenting expands substantially over the period 2006–2024 and attracts a large number of applicants. This confirms that artificial intelligence offers broad technological opportunities and creates space for participation by many actors. The growth of the applicant base, the continuous entry of new applicants and the broad IPC structure of AI patents all point to a dynamic and expanding technological domain.

However, this expansion does not produce an evenly distributed innovation structure. The evidence shows that inventive output remains highly asymmetric. A small group of leading applicants accounts for a disproportionate share of applicant mentions, while the majority of applicants appears only once or only a few times in the dataset. This means that participation is broad, but innovative leadership is much more selective. In my interpretation, this is the key result of the thesis: AI opens the door to many actors, but it does not give all of them the same capacity to persist, accumulate and shape technological trajectories.

This distinction between participation and leadership is crucial. If AI is judged only by the number of applicants entering the field, it may appear to be a democratizing technology. Yet the deeper empirical evidence suggests a more cautious conclusion. Entry is numerous, but much of it remains marginal. Many actors participate in AI patenting, but only a smaller group develops a persistent and significant inventive presence. Therefore, the expansion of AI-related innovation should not be confused with a full redistribution of innovative power.

The analysis of entry dynamics reinforces this point. New applicants continue to enter the AI patenting domain throughout the period, confirming the openness of the field. At the same time, continuing applicants become increasingly important, indicating that cumulative participation matters as the regime matures. This coexistence of new entry and growing persistence shows that AI is not a purely turbulent regime. It is a regime in which experimentation and accumulation develop together.

The analysis of technological portfolios further strengthens this interpretation. Leading applicants are not only more active in patenting terms; they also differ in the way they organize their technological knowledge. Some firms display broad IPC portfolios, suggesting exploration and recombination across multiple technological domains. Others remain more focused around selected IPC classes, suggesting specialization and cumulative learning. This evidence shows that the core of the AI patenting system is not homogeneous. It is composed of firms with different technological strategies, sectoral backgrounds and degrees of portfolio breadth.

Overall, the thesis shows that artificial intelligence is open in terms of participation, but selective in terms of innovative output, persistence and technological scope. This is why the concept of a hybrid Schumpeterian technological regime is useful. AI displays Mark I dynamics because it generates entry, experimentation and technological variety. At the same time, it displays Mark II dynamics because innovative output remains concentrated, capabilities accumulate over time and leading applicants maintain a stronger position within the technological domain.

The main contribution of the thesis is therefore to move beyond two simplified interpretations. The first is the optimistic view that AI automatically democratizes innovation because tools, models and applications are increasingly accessible. The second is the pessimistic view that AI is entirely closed and controlled by a small number of dominant firms. The evidence suggests that both views are partial. AI democratizes access more than leadership. It expands the number of participants, but the ability to become a persistent and influential innovator remains unevenly distributed.

This result has broader implications for the study of industrial dynamics. It suggests that entry should not be interpreted only as a numerical phenomenon. What matters is not only how many actors enter a technological domain, but also whether they remain active, accumulate capabilities and influence the direction of technological change. In the case of AI, the large number of entrants coexists with a core–periphery structure: a broad periphery of occasional or low-intensity applicants and a more restricted core of persistent, technologically capable and strategically differentiated actors.

The thesis also confirms the importance of looking at multiple empirical dimensions together. Patent growth, applicant counts, concentration ratios, HHI, entry dynamics and technological portfolios each reveal a different part of the AI regime. Taken separately, these indicators could lead to incomplete interpretations. Taken together, they show a more complex structure: AI is expanding but not equalizing; open, but selective; dynamic, but cumulative.

The limitations of the analysis should also be recognized. Patent data capture inventive activity, but not all forms of AI innovation. Some important AI developments occur through open-source projects, publications, trade secrets, software systems or organizational implementation that may not be patented. The dataset also reflects a specific patenting environment and does not fully capture global differences across regions such as Europe, China and the United States. Moreover, the analysis is descriptive and interpretative rather than causal. For this reason, the results should be read as evidence of structural patterns in AI patenting, not as a complete causal explanation of AI industrial dynamics.

Despite these limitations, the evidence provides a coherent answer to the research question. Artificial intelligence should be interpreted as a hybrid technological regime in which widening and deepening coexist. The growth of AI patenting broadens participation, but it does not eliminate asymmetries in innovative leadership. The regime creates opportunities for many actors, yet it also rewards those with cumulative capabilities, complementary assets and the ability to organize technological knowledge across core domains.

The final conclusion is therefore clear: artificial intelligence does not simply democratize innovation, nor does it fully close it. It creates a layered and stratified innovation structure. Many actors can enter and experiment, but fewer can persist, accumulate inventive output and shape the central technological trajectories. This hybrid configuration is likely to remain one of the defining features of AI-related industrial dynamics, especially as technology continues to diffuse across sectors while its core capabilities remain demanding in terms of data, computation, research intensity and organizational capacity.

Bibliography

AI Index Steering Committee. (2025). *The AI Index 2025 Annual Report*. Stanford Institute for Human-Centered Artificial Intelligence.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., Donahue, C., et al. (2021). *On the opportunities and risks of foundation models. arXiv.*

Breschi, S., Malerba, F., & Orsenigo, L. (2000). Technological regimes and Schumpeterian patterns of innovation*. The Economic Journal*, 110(463), 388–410.

Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies: *"Engines of growth"? Journal of Econometrics,* 65(1), 83–108.

Castellacci, F. (2007). Technological regimes and sectoral differences in productivity growth. *Industrial and Corporate Change, 16*(6), 1105–1145.

Cockburn, I. M., Henderson, R., & Stern, S. (2018). *The impact of artificial intelligence on innovation.* NBER Working Paper No. 24449.

Corrocher, N., Malerba, F., & Montobbio, F. (2007). Schumpeterian patterns of innovative activity in the ICT field. *Research Policy, 36*(3), 418–432.

Fontana, R., Martinelli, A., & Nuvolari, A. (2021). Regimes reloaded! A reappraisal of Schumpeterian patterns of innovation, 1977–2011. *Journal of Evolutionary Economics, 31*, 1495–1519.

Fontana, R., Nuvolari, A., Shimizu, H., & Vezzulli, A. (2013). Reassessing patent propensity: Evidence from a dataset of R&D awards, 1977–2004. *Research Policy, 42*(10), 1780–1792.

Malerba, F. (2005). Sectoral systems of innovation: A framework for linking innovation to the knowledge base, structure and dynamics of sectors. *Economics of Innovation and New Technology, 14*(1–2), 63–82.

Malerba, F., & Orsenigo, L. (1995). Schumpeterian patterns of innovation. *Cambridge Journal of Economics, 19*(1), 47–65.

Malerba, F., & Orsenigo, L. (1996). Schumpeterian patterns of innovation are technology-specific. *Research Policy, 25*(3), 451–478.

Malerba, F., & Orsenigo, L. (1997). Technological regimes and sectoral patterns of innovative activities. *Industrial and Corporate Change, 6*(1), 83–118.

OECD. (2024). *Artificial intelligence, data and competition*. OECD Artificial Intelligence Papers, No. 18. OECD Publishing**.**

Orsenigo, L. (1995). Technological regimes, patterns of innovative activities and industrial dynamics: A survey of empirical evidence and some theoretical models. *Cahiers d'économie et sociologie rurales, 37*, 26–67.

Peneder, M. (2010). Technological regimes and the variety of innovation behaviour: Creating integrated taxonomies of firms and sectors. *Research Policy, 39*(2), 323–334.

Schumpeter, J. A. (1934). *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle*. Harvard University Press.

Schumpeter, J. A. (1942). *Capitalism, socialism and democracy*. Harper & Brothers.

Srholec, M., & Verspagen, B. (2012). The voyage of the Beagle into innovation: Explorations on heterogeneity, selection, and sectors. *Industrial and Corporate Change, 21*(5), 1221–1253.

Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6), 285–305.

Van Dijk, M. (2000). Technological regimes and industrial dynamics: The evidence from Dutch manufacturing. *Industrial and Corporate Change, 9*(2), 173–194.