



University
of Bremen

Faculty of Business Studies and Economics

University of Bremen

Doctoral Thesis

Technological change and absorptive capacity:

An evolutionary perspective of Artificial Intelligence

A thesis submitted to the Doctoral Commission Dr. rer. pol. of the University of Bremen in fulfilment of the requirements for the degree of Dr. rer. Pol.

Matheus Eduardo Leusin

Date of submission: March 25, 2022.

Date of colloquium: June 17, 2022.

First examiner:

Prof. Dr. Jutta Günther

Second examiner:

Prof. Dr. Martin G. Möhrle

Bremen, 2022

ABSTRACT

This thesis takes an evolutionary theoretical perspective to analyse knowledge dynamics linked to the development of Artificial Intelligence (AI) innovations. An introductory chapter presents the thesis' overarching theoretical framework, research questions, main findings, contributions, policy implications, and limitations. This chapter is followed by three empirical papers that focus on distinct knowledge dynamics related to the creation of AI innovations. The papers provide a comprehensive overview of the technological development of AI, how this development was reflected in countries' exploration of this technology, and how it affected the innovative activities of firms that introduced AI innovations. The combined findings from these papers, it is argued, allow a better understanding of how specific technologies and global technological development can impact firms' ability to learn and innovate.

Keywords: *Artificial Intelligence; Evolutionary economics; Absorptive capacity; Geography of innovation; Economic Complexity; Technological relatedness; Technological Change;*

Contents

1	CHAPTER 1: TECHNOLOGICAL CHANGE AND LEARNING – THE CASE OF ARTIFICIAL INTELLIGENCE	
	1.1. Introduction.....	2
	1.2. Theoretical background.....	4
	1.2.1. Evolutionary Economics	4
	1.2.2. Learning and technological paradigms, trajectories, and paths	8
	1.2.3. Absorptive capacity and relatedness	11
	1.2.4. The geography of innovation and Economic Complexity.....	12
	1.3. What is Artificial Intelligence.....	15
	1.3.1. AI’s history and technological development	15
	1.3.2. AI’s potential and international policy responses	18
	1.4. Research questions.....	20
	1.4.1. The geographical learning of AI.....	22
	1.4.2. AI learning from a firm perspective.....	26
	1.5. Overview of chapters	28
	1.5.1. Patenting Patterns in Artificial Intelligence: Identifying National and International Breeding Grounds (Chapter 2)	30
	1.5.2. Break-through or Break-in? How AI Becomes a Part of National Technological Trajectories (Chapter 3).....	32
	1.5.3. The Development of AI in Multinational Enterprises – Effects upon Technological Trajectories and Innovation Performance (Chapter 4)	34
	1.6. Contributions, implications and future research	36
	1.6.1. Main findings	36
	1.6.2. Theoretical contributions	38
	1.6.3. Policy implications.....	44
	1.6.4. Limitations and future research	46
	References	50
2	CHAPTER 2: PATENTING PATTERNS IN ARTIFICIAL INTELLIGENCE: IDENTIFYING NATIONAL AND INTERNATIONAL BREEDING GROUNDS	
	2.1. Introduction.....	56
	2.2. Literature review: Breeding Ground indicators and related literature.....	57
	2.3. Data collection.....	59
	2.4. Comparison of strategies for identifying AI patents	66
	2.5. Method and results	68
	2.5.1. AI specialists by Revealed Technology Advantage	69

2.5.2.	National AI Breeding Grounds.....	71
2.5.3.	International AI Breeding Grounds	74
2.6.	AI Techniques: Evolution, specialisations and Breeding Grounds.....	77
2.7.	Discussion	79
2.8.	Conclusions.....	81
	References	83
	Appendix A: Query for identifying AI application IDs.	85
	Appendix B: IPC-search strategy adopted in papers used for comparison.	86
	Appendix C: Variables considered and their definitions, according to the authors.....	87

3

CHAPTER 3: BREAK-THROUGH OR BREAK-IN? HOW AI BECOMES A PART OF NATIONAL TECHNOLOGICAL TRAJECTORIES

3.1.	Introduction.....	89
3.2.	Theoretical background.....	92
3.2.1.	The concept of relatedness	92
3.2.2.	Relatedness and STI policies.....	93
3.2.3.	Research gap	94
3.2.4.	Research questions.....	95
3.3.	Data and methods	97
3.3.1.	Data collection and identification	98
3.3.2.	Method	99
3.4.	Empirical analysis	102
3.4.1.	Identifying technologies linked to AI in the ATS.....	102
3.4.2.	General and AI-specific specialisations in the GTS.....	104
3.4.3.	General and AI-specific specialisations in technological classes at the 4-digits IPC level	110
3.5.	Discussion and concluding remarks	113
3.5.1.	Summary of the main findings	113
3.5.2.	Contributions.....	114
3.5.3.	Policy implications.....	116
3.5.4.	Limitations and future research directions	118
	References	119
	Appendix A: Search strategy applied in PATSTAT 2019 spring version for identifying AI patents. ...	123
	Appendix B: Log 10 of number of AI patents by Japan, the United States (USA), South Korea, and China.....	124
	Appendix C: Technological fields connectivity considered in the Global Technological Space (GTS) network.	125

Appendix D: Tabular values of countries specialisations over the considered intervals (D1), followed by their respective plots (D2). 126

Appendix E: AI relatedness and specialisations over the considered intervals for IPC subclasses. . 128

4 CHAPTER 4: BREAK-THROUGH OR BREAK-IN? HOW AI BECOMES A PART OF NATIONAL TECHNOLOGICAL TRAJECTORIES

4.1. Introduction..... 131

4.2. Theoretical Background and Research questions 133

4.2.1. Digital Technologies and AI 133

4.2.2. Relatedness and Innovative Performance..... 134

4.2.3. Hypotheses development 136

4.3. Data & Method 139

4.3.1. Constructing the Dataset..... 139

4.3.2. Matching AI Adopters to Non-adopters 140

4.3.3. Estimating the Effects of AI Adoption 141

4.3.4. Dependent Variables Considered..... 142

4.3.4.1. Path dependency..... 142

4.3.4.2. Innovative Performance 143

4.4. Empirical Findings..... 144

4.4.1. Overview of Dataset 144

4.4.2. Estimated effects..... 145

4.4.2.1. Effects on Relatedness 145

4.4.2.2. Relatedness effects across sectors and the self-selection bias..... 147

4.4.2.3. Effects on Innovative Performance 150

4.5. Discussions and Conclusions 155

4.5.1. Discussion 155

4.5.2. Contributions..... 158

4.5.3. Limitations and future research 159

References 160

Appendix A: Constructing a dataset combining Patents and Multinational Enterprises data (2011-2019)..... 163

Appendix B: General descriptives of the complete dataset. 174

Appendix C: Data about 30 larger AI adopters in 2016 (Table C1) and 2019 (Table C2). 176

Appendix D: Data about AI adopters before and after matching..... 178

Appendix E: Average distance of every Nace sector to the AI technological cluster. 180

Appendix F: Dynamic effects for the self-selection analysis for the first (Figure F1) and second (Figure F2) classification methods used. 184

Appendix G: Distribution of distinct considered samples regarding the size of the companies. 186



PERSONAL CONTRIBUTIONS TO THE PAPERS OF THE CUMULATIVE DISSERTATION	187
ACKNOWLEDGEMENTS.....	188

List of Figures

Figure 1: Author keywords co-occurrence over the years.	18
Figure 2: Comparison of results between the selected queries.....	67
Figure 3: Description of the steps followed.	69
Figure 4: Top 20 Patent Offices according to (and sorted by) the sum of their RTAs over three periods.	70
Figures 5 (top) and 6 (bottom): Top 15 Patent offices which are considered National AI Breeding Grounds, according to the Nat Breeding Ground_Countryp and the Nat Breeding Ground_Weighted_Countryp indicators, respectively.....	73
Figures 7 (top) and 8 (bottom): Top 15 Patent Offices which are considered International Breeding Grounds, according to the Int Breeding Ground_Countryp and the Int Breeding Ground_Weighted_Countryp indicators, respectively.....	76
Figure 9: Evolution of the patents of each of the AI-related techniques considered, sorted by the total number of patents for each technique.	77
Figure 10: ATS and AI-core technologies over the considered intervals and IPC fields.....	103
Figure 11: GTS for IPC fields.	105
Figure 12: Four types of specialisations considered.	106
Figure 13: Specialisations in the considered technological fields for Japan (13a), United States (13b), South Korea (13c), and China (13d).....	108
Figure 14: Share between coinciding and general specialisations for the four considered countries at the technological field level.	109
Figure 15: General (15a) and AI-specific specialisations (15b) across patent subclasses for Japan, the United States, South Korea, and China.	111
Figure 16: Share between coinciding and general specialisations for the four considered countries at the subclass level for ten codes (16a) and all codes (16b).....	113
Figure 17: Estimated Dynamic Treatment Effects of AI adoption on the relatedness of all MNEs (17a) and when holdings and head offices are excluded (17b).....	147
Figure 18: Estimated effects for the number of patents owned by the entire MNE for the entire (18a) and for the smaller samples with larger companies (18b and 18c).	153
Figure 19: Estimated effects for Number of patents per unit of turnover (19a and 19b) and R&D expenses per patent (19c and 19d).....	154

List of Tables

Table 1: Overview of the dissertation’s Chapters by their main characteristics..... 29

Table 2: AI techniques considered and definitions 62

Table 3: Keywords used for the patent search..... 64

Table 4: National and International Breeding Ground leaders for each AI technique considered..... 78

Table 5: Types of technological space considered. 102

Table 6: Estimates for the aggregate treatment effects of AI adoption on firms’ relatedness. 146

Table 7: Results estimated for the self-selection bias across the considered two distinct knowledge distances to AI. 150

Table 8: Estimated effects for the three distinct dependent variables related to innovative performance..... 151

Author: Matheus Eduardo Leusin

Abstract: In this introductory chapter, I present the overarching theoretical framework, research questions, findings, and contributions of this thesis. I start by motivating the focus of the thesis on AI technologies and learning (Section 1), before I introduce an evolutionary theoretical framework on learning, both from a geographic and from a firm-level perspective (Section 2). Then, I present various aspects of AI's historical development (Section 3), which are followed by an outline of the research questions addressed in this thesis (Section 4). Next, I provide an overview of the three subsequent chapters that compose the core of the thesis (Section 5) and summarize the main findings, complementarities between the chapters, and policy implications (Section 6). This chapter ends with a discussion of the theoretical contributions of the thesis, as well as its limitations and suggestions for future research (Section 7).

Keywords: Artificial Intelligence; Technological development; Technological learning; Evolutionary Economic Theory; Relatedness;

JEL Classification: D22; D83; O3; O14; O25; O33

Publication

This is the introductory paper of this cumulative thesis submitted to the Doctoral Commission of Bremen in fulfilment of the requirements for a Dr. rer. pol. degree.

1.1. Introduction

The term “Artificial Intelligence”, or AI for short, was proposed in the 1950s. It refers to the creation of a machine with abilities of perception, cognition, and action close to human abilities (Li & Jiang, 2017; Russell & Norvig, 2016). Much time would pass before AI started to shift from theory to commercial applications, which only has been happening in a significant way in the last decade (WIPO, 2019). Today’s AI applications are much narrower than what its original proposition would suggest, but expectations were never higher. These expectations are not only due to AI’s significant improvements in performance but also due to its technological particularities. These particularities include AI’s potential to be incorporated into other technologies (Nilsson, 2009) from a very broad range of technological fields (Righi et al., 2020). This broad applicability is reflected in AI innovations, which include applications that improve social well-being in areas as diverse as climate change, health, transportation, and public safety, among others (Fujii & Managi, 2018; Harhoff, Heumann, Jentzsch, & Lorenz, 2018). Regarding its economic potential, it is expected that companies and countries adopting AI will gain significant advantages over global markets and industries, whereas late adopters may become locked out of these markets (Cockburn, Henderson, & Stern, 2018; European Commission, 2017; Klinger, Mateos-Garcia, & Stathoulopoulos, 2018).

Accordingly, a “global race” has been deployed in the last decades, with an increasing number of countries investing in their local development of AI (Harhoff et al., 2018; Klinger et al., 2018). In this policy context, national targets are set for AI development, and large flows of investment follow. However, AI’s technological development is a moving target, making the design and implementation of such policies particularly hard. It is also currently unknown how AI’s broad technological range may affect local knowledge development. In what areas are AI’s commercial applications being developed? Should countries invest in incorporating AI into the sectors in which they already have some specialisation advantage, or should they aim at developing new competencies in the most relevant areas of AI? Furthermore, considering the increasing role of businesses in AI development, what are the effects of AI adoption at the firm level? Should innovation policies target companies from specific sectors and avoid others, or should AI be incentivised to diffuse across many sectors?

This thesis approaches these questions. Taking an evolutionary view on learning (Dosi, 1982; Nelson & Winter, 1982), it analyses how distinct knowledge bases influence the emergence of

AI innovations. The evolutionary view postulates, among other things, that geographical spaces agglomerate particular technologies, capital, institutions, and skills that shape the local development of knowledge (Christopher Freeman, 1987; Lundvall, 1992). Firms, likewise, accumulate knowledge in a path dependent process and change their routines accordingly (Nelson & Winter, 1982). This particular evolutionary perspective on learning makes AI's broad applicability an especially relevant case. By analysing how distinct knowledge bases affect the emergence of AI knowledge, one can better understand which areas of AI should be targeted by national innovation policies. A focus on AI is also relevant to advance the evolutionary view on how new knowledge affects existing knowledge bases. By considering how the same technology is absorbed by actors with very distinct knowledge bases, one can better understand the commonalities of learning and its effects on technological trajectories (Dosi, 1982).

The thesis is based on three scientific articles (Chapters 2-4), which analyse AI's development from a country-level and a firm-level perspective. Countries leading AI development are identified (Chapter 2) and used as a reference to understand how successful technological trajectories with this technology emerge locally. The technological competencies of these leading countries are examined and compared to AI innovations deployed locally (Chapter 3). Finally, it is analysed how AI changes the technological trajectories of firms with distinct knowledge bases. The focus is both on how AI adoption changes the innovations created by firms in comparison to their previous innovations, and on the possible effects of AI adoption on firms' innovative performance (Chapter 4).

Together, these three scientific articles provide a comprehensive overview of the technological development of AI, how this development was reflected on countries' exploration of this technology, and how the most recent wave of AI development affects the innovative activities of firms that introduce AI innovations. The papers also make theoretical contributions to the evolutionary view on learning and technological trajectories (Dosi, 1982; Nelson & Winter, 1982). Two novel mechanisms are identified as affecting this view. One arguably emerges due to particular learning affordances provided by specific technologies. This mechanism is particularly helpful to explain differences in firms' technological performance after learning about specific technologies. The second mechanism emerges due to dynamics generated by technological change. This mechanism is useful for further understanding the emergence of lock-in effects.

The subsequent Section 2 presents the main principles of the evolutionary economic theory to understand learning, both from a firm and a geographical perspective. Section 3 discusses the technological characteristics of AI, its historical development, and the current expectations about this technology. Section 4 explains the research questions addressed, and Section 5 links these research questions to the findings from the three scientific papers. This section also clarifies how the articles complement each other. Finally, Section 6 ends this introductory chapter by highlighting the main findings, theoretical contributions, policy implications, and limitations of this thesis, besides providing an outlook for future related research.

1.2. Theoretical background

1.2.1. Evolutionary Economics

The bedrock of the current evolutionary view on economics can be traced to the book “An evolutionary theory of economic change”, by Nelson and Winter (1982)¹. At the core of the book, the authors address how technologies and economic production are affected by changes in economic knowledge. Subsequent research further developed the seminal ideas proposed in the book – e.g., Dosi (1988), Winter (1987), Bresson (1987) –, consolidating one important line of current evolutionary theory in economics.

The evolutionary theory proposes a focus on innovation as the main driver of economic development. However, this proposition had been made also in some earlier works, most notably by Joseph Alois Schumpeter (1939). In very broad terms, the evolutionary theory can be understood as the theory about how society or the economy learn (Dosi & Nelson, 1994). It stresses the role of random events, suboptimal decisions, systematic errors, trials, discoveries, opportunities, and how these are affected by distinct timely-related events (ibid). In this sense, the evolutionary theory is an alternative to the neoclassical view of economics, which proposes that observed economic variables can be explained as the result of rational actors making choices that maximize their utilities. Conversely, the evolutionary perspective stresses how agents present rule-guided behaviours that are context and time-specific. This

¹ Although an evolutionary perspective has been used before this, e.g., in the works of Malthus and Marx, as highlighted in Dosi & Nelson (1994).

time-dependent dynamic allows agents to experiment and learn new rules, which further affects their behaviour.

The theoretical basis from evolutionary thinking branched into several other existing areas – e.g., Evolutionary economic geography (Boschma & Frenken, 2018), Evolutionary Behavioural Economics (Burnham et al., 2016), Evolutionary Theory of the Firm (D. Teece & Pisano, 1994; D. J. Teece, Rumelt, Dosi, & Winter, 1994) –, besides influencing the creation of specific perspectives with distinct foci. Examples of the latter are National Innovation Systems (Christopher Freeman, 1987; Chris Freeman, 1995), which focus on explaining national patterns of innovation, and Technological Innovation Systems (Bergek, Jacobsson, Carlsson, Lindmark, & Rickne, 2008; Hekkert, Suurs, Negro, Kuhlmann, & Smits, 2007), which focus on explaining the emergence and development of specific technologies.

All these branches and theoretical frameworks share the same core characteristics, which allow classifying them together under the evolutionary theory. They all try to explain some phenomenon over time, or how time allowed this phenomenon to develop in the way it did. They also account for the existence of mechanisms that are random and systematically generate or renew some of the variables considered (Dosi & Nelson, 1994). This randomness is linked to a combination of the mistake-driven processes of learning and discovery. A selection mechanism acts upon the units involved in these processes, preserving the determinants that improve their “fitness” to the environment in which they are.

As the economic evolutionary view is largely based on principles first introduced under a biological perspective, it is worth comparing both perspectives for a better understanding of the evolutionary economic theory. In this sense, Dosi and Nelson (1994) present four building blocks of evolutionary thinking in economics that are comparable to evolutionary biology: i) a fundamental unit of selection (in biology seen through the genes); ii) a mechanism that links the selection to a level of entities (under the biological perspective, specific phenotypes emerging as the result of the genotypes that were selected by the environment); iii) some interaction process, which allows the emergence of the selection dynamics; and iv) mechanisms that generate variations in the fundamental units, which are then reflected on how these units respond to their environments (under the biological view: A mechanism that creates random variations in genes, which are then reflected on phenotypes).

Still according to Dosi and Nelson (1994), the unit of selection (i) in economics typically refers to individuals, organisations, countries, etc. Like the DNA in biology, these units are also composed of more fundamental units that influence their behaviour. These fundamental units may be technologies, policies, behavioural patterns, cultural traits, mental categories, representations, rules, etc. They are modified and improved over generations following some rules of transmission, and the modifications influence some observable aspects (the phenotype in biology) of the unit of analysis. At the example of technological evolution under a firm perspective, the authors highlight how organisations are a composition of several distinct technologies. The selection of these technologies is up to the firm (typically related to choices on R&D, portfolio diversification, decisions rules related to investments, etc.), but the firm's success resulting from this selection depends on external selection forces from the environment, i.e., the market. The market selects relatively complex products or technological systems instead of individual elements of technological knowledge, thus penalizing or rewarding whole organisations instead of, for example, their specific technological choices. "Organisational routines" and "competencies" are the typical fundamental units used to explain how organisations evolve and are selected over time.

Dosi and Nelson (1994) further explain the mechanisms and criteria of selection (ii), which relate to how the "fitness" of units is ultimately considered. As already mentioned, fitness refers to some aspect relevant for the unit to be selected by its environment. The environment, represented by the market, is typically divided into two perspectives: The product-market and the financial market. Each of these markets has distinct factors affecting the behaviour of a firm, with some of them often conflicting. The financial market assesses firms based on their cash flows, accounting profits, expectations about future profits, historical successes, technological portfolio, etc. In the product market, the criteria include the quality of products, prices, after-sale servicing, marketing networks, delays, etc. One can see here an important difference between the economic and biological evolutionary perspectives: Whereas the latter deals with selection mechanisms that are invariantly stable – e.g., rates of reproduction, the efficiency of accessing food, etc. –, the former deals with factors that are often changing due to economic and social circumstances. Thereby, the question of "when" plays a large role in explaining the success or failure of a given unit in the economic evolutionary perspective.

The interaction process (iii) and the generation of variations (iv) are interconnected (ibid). It is through these two mechanisms that the departure from the neoclassical theory becomes clearer. They both highlight the lack of possible ex-ante optimal solutions by proposing that discovery and experimentation are processes that occur to a large extent at random. These processes, which lead to the emergence of novelty, are specific to the context in which the unit of analysis is. This contradicts basic assumptions from the neoclassical theory, which does not assume the existence of systematic mistakes associated with ignorance or lack of understanding (Dosi & Nelson, 1994). It instead assumes that actors have a correct understanding of their choices and possible consequences. The evolutionary theory, in turn, assumes that actors take actions in environments characterised by novelty so that appropriate actions do not yet exist but rather they are to be learned. The standard unit of efficiency is thus defined by the most efficient existing firms and not by the theoretically possible maximum efficiency. This leads to a perspective of learning actors that strongly contrasts the rational actors considered in the neoclassical view. By neglecting the aspects of learning, the neoclassical theory does not account for the particularities of distinct contexts that affect actors' decisions. These particularities are a major motivation for the evolutionary understanding of learning as a path dependent dynamic process (ibid).

Finally, the variations (iv) are often seen through changes in the rule-guided behaviours that actors have. In the seminal work of Nelson and Winter (1982) they are labelled "routines". These routines are shaped by agents' history, pre-existing knowledge, value systems, and prejudices (Dosi & Nelson, 1994). As there is no guarantee that optimal routines exist at all, there is always the possibility of improving them. As much as these routines are changed by interactions with other actors (iii), they are also shaped by the imperfect and mistake-prone process that characterises learning.

From the broad range of research foci that economic evolutionary studies have, one can argue that particular attention has been given to the nature of the learning process, mechanisms of adaptation, processes of discovery and selection underlying economic growth, theory of the firms and their routine creation, as well as dynamics of industrial organisation. The units of selection (i) and the criteria to define "fitness" (ii) change in these contexts depending on the research aim, whereas a focus on actors' interactions (iii) and how variation is created (iv) are often implicit in the underlying assumptions of these studies.

1.2.2. Learning and technological paradigms, trajectories, and paths

Since learning is at the core of the evolutionary theory, plenty of research efforts concern understanding the distinct mechanisms linked to learning. The process of learning is understood to be cumulative, path dependent, and interactive (Dosi, 1982; Nelson & Winter, 1982). This means that learning is influenced by very distinct factors. It is influenced by i) firm-specific factors i.e., the knowledge that an organisation accumulated, ii) the environment i.e., how much knowledge is currently known about existing problems, and iii) interactions between a variety of actors such as firms, suppliers, clients, universities, etc.

For firm-specific factors (i), a particular and critical element is the knowledge they accumulated, which is to be formative of their traits. Much of the evolutionary thinking on firms' evolution and market dynamics focuses on the organisational and behavioural traits of firms (Dosi & Nelson, 1994). Following the seminal work of Nelson and Winter (1982), these traits are often approximated through the notion of routines. The authors distinguish between three main types of routines: i) standard operating procedures, which determine how the firm acts in its daily procedures of production; ii) investment routines, which comprehend behaviours that affect the growth or decline of firms as a function of their profits; and iii) deliberation routines, which involve firms' search for better ways of doing things.

All of these routines occur under a specific environment (ii). From a learning perspective, this environment represents the collection of knowledge about a given topic and its possible solutions. An innovation can be understood in this environment as a possible solution to a specific problem. One branch of the evolutionary view on learning highlights that innovations are created by recombining distinct pieces of knowledge in novel ways (Frenken, Van Oort, & Verburg, 2007; Van den Bergh, 2008).

It is useful to approximate the environment in which firms operate through the notion of technological paradigms (Dosi, 1982). The concept of technological paradigm tries to capture the nature of the available technological knowledge and organisational procedures in a given moment. More broadly, a technological paradigm comprehends the general outlook of the productive problems faced by firms (ibid). It refers to the collection of understandings about particular technologies that are shared, including views about their potential and limitations.

This shared understanding influences the prevailing views of how to make things better, which influences firms' innovative activities (Dosi & Nelson, 1994).

Thus, the idea of technological paradigm links firms' routines to an environment of existing productive problems and shared understandings about how to solve them. Each technological paradigm has its own technological trajectory, which refers to the improvement path taken with the technology to which the paradigm refers. Technological development in a technological trajectory can lead to continuous or discontinuous changes. Continuous changes refer to technological progress along the existing trajectory, whereas discontinuous changes are associated with the emergence of new paradigms (Dosi, 1982).

Therefore, firms' learning and routines are influenced by their perceptions of opportunities available in the current technological trajectory. These perceptions affect what firms believe they can improve, and hence how they innovate (deliberation routines). Another fundamental dimension of the environment (ii) is its particular "fitness" criteria, which influences which kind of innovations are selected. The fitness criteria are to be different across technological trajectories, and even more distinct across technological paradigms.

These dynamics help to understand how learning can be influenced by random events occurring in the environment of firms. Random events influence the selection criteria. At the innovation level, the technological performance in solving a problem is just one factor out of several that contribute to its selection over alternatives. For example, better-performing innovations fail if they are too expensive, if they do not easily complement other existing technologies, if there is a lack of suppliers for specific components, if buyers have specific demands, etc. All these possibilities contribute to the idea of firms being continuously learning, both about the technological paradigms in which they are and about the consequences of their decisions. Hence, not only a great part of learning is generated by random events – e.g., one actor made a discovery that allowed improving the selected technology, making it significantly superior in performance in comparison to a potential substitute – but also the selection process of successful routines is random – e.g., one technology may become more successful for being related to a larger network, influencing current profits that will be generated according to the firm's past decisions.

Another factor affecting learning (and the selection of innovations) is the presence of interactions (iii). As already mentioned, interactions complement firms' learning. Profitable

innovations created by any given actor are, with a lag, ultimately imitated by other firms in the industry. Thus, the profitability of a firm is determined not only by what it is doing but also by what its competitors do. These interactions explain some stylized facts, such as the differences seen in firms' performance. Successful technological innovation generates profits for a firm, leading it to grow over its competitors. This growth is limited by how fast competitors learn about the new technological innovation. Over time, firms using more profitable technologies grow, and the more profitable technologies are imitated and adopted by competitors, replacing old technologies. At the same time, successful technologies receive more investments, being even further developed than alternatives². As a result, the productivity as a whole of an economy or industry is based on its aggregated "technical advance". This advance, in turn, reflects both the improvement of individual technologies and the replacement of less productive technologies for more productive ones (Dosi, 1982; Nelson & Winter, 1982).

Finally, there are some aspects related to potential drawbacks of learning that are worth mentioning. In particular, the accumulation of knowledge may create a bias towards established routines. The conflict between new and old routines is especially highlighted in the concepts of technological lock-in and competence destroying innovations. The basic idea of technological lock-in is that firms and markets may get "stuck" in an inferior technological alternative due to network effects (Arthur, 1989; Cowan, 1990). This means that the established routines of a collection of firms from a network do not adjust for absorbing a more promising technological alternative. This can cause this network of firms to be selected out of their environments. Competence destroying innovations, in turn, refer to what happens to existing industries or firms when a new potentially disruptive technology emerges³. Tushman and Anderson (1986) coined the term "competence destroying technical advance" to characterise how such technologies may render existing skills and experience irrelevant if they cannot be transferred from the replaced technology to the new one. Similar to lock-in, these technologies may also lead markets and incumbent firms to be deselected from their environments.

² Accordingly, the selected alternatives advance in their technological trajectory.

³ Here again is seen a parallel with the early work of Schumpeter. The term coined the term "creative destruction" to describe how new technologies may replace existing ones over time, potentially leading to the disruption of existing firms (Schumpeter, 1942).

1.2.3. Absorptive capacity and relatedness

The concept of absorptive capacity was proposed in Cohen and Levinthal (1990) to describe a firm's ability to recognise the value of new knowledge, assimilate it, and apply it to commercial ends. According to the authors, this ability is largely affected by how the firm's prior knowledge relates to the new knowledge. The authors examine the cognitive structures that underlie learning, linking several stylized facts about learning to the concept of absorptive capacity, e.g., evidence from research on memory suggesting that accumulated prior knowledge increases the ability to learn, remember, and use the existing knowledge (Hilgard & Bower, 1966); evidence showing that experience in one learning task may influence and improve the performance of some subsequent similar learning tasks (Ellis, 1965); and the fact that the preconditions for successful learning are the same as the conditions for problem-solving, and therefore, for the creative process (Simon, 1985).

Following the evolutionary view, Cohen and Levinthal (1990) stress the role of learning as a cumulative, path dependent process. The authors also highlight the role of knowledge diversity in facilitating the innovative process by enabling an individual to make novel associations between distinct types of knowledge. They then transfer these insights from the individual to a firm perspective. From a firm perspective, learning takes place through the interactions formed between several individuals related to the organisation (i.e., not only the individuals within the organisation, but also other actors like buyers and suppliers). These links form the knowledge structure that creates absorptive capacity, i.e., the ability to learn as a function of related knowledge. This ability is also to influence firms' decisions regarding the allocation of resources for innovative activity. The authors still hypothesize that absorptive capacity is a by-product of firms' routine activities and that as such, it is not likely that absorptive capacity can be created directly through investments.

Breschi, Lissoni, and Malerba (2003) propose a way to measure absorptive capacity through the indicator of knowledge relatedness. By linking firms' innovations to their knowledge bases, the indicator is able to capture several stylized facts related to firms' innovative activities. Between the stylized facts captured by relatedness, one can cite: Large firms are multi-technology corporations that combine a usually greater number of technologies to develop and produce a lower number of products and services (Granstrand, 1998); corporate technological diversification changes only slowly over time (Cantwell & Andersen, 1996); and

profiles of technological diversification differ across firms due to history, distinct initial conditions, market incentives, distinct institutional settings and other factors (Antonelli, Krafft, & Quatraro, 2010; Ivarsson, Alvstam, & Vahlne, 2015; Le Bas & Sierra, 2002), despite being very similar among large firms producing similar products (Breschi et al., 2003; D. J. Teece et al., 1994).

The indicator of relatedness encompasses three dimensions of knowledge: i) knowledge proximity, related to learning processes from unintended learning spillovers or intended local learning; ii) knowledge commonalities, related to the possibility of firms' innovative activities spanning over more than one technology as a result of the same type of knowledge being used in more than one technology; and iii) knowledge complementarities, related to firms' need to master more than one technology for being able to develop new products and/or processes. In essence, relatedness captures the fact that firms extend their innovative activities across knowledge-related technological fields as a consequence of both unintended and intended learning processes and of knowledge features of relatedness and its links (ibid).

Based on the proposed relatedness indicator, Breschi et al. (2003) present evidence that firms diversify their innovative activities mainly by exploring related technologies, with larger innovators being typically more coherent in their technological trajectories than smaller innovators. Subsequent studies provide evidence that firms innovating in related areas have higher survival rates (Colombelli, Krafft, & Quatraro, 2013), lower coordination costs (Nesta, 2008), and perform better in knowledge transfer and creation (Weber & Weber, 2010). Relatedness has also been shown to moderate positively an inverted U-shaped relationship between technological diversification and technological performance (Kim, Lee, & Cho, 2016; Leten, Belderbos, & Van Looy, 2007). It eases the burden of having a widely diversified portfolio by reducing the costs of learning, but just to a limited extent. Too much diversification – despite offering further opportunities for cross-fertilisation and technology fusion – may hurt innovative performance due to higher coordination and integration costs (Leten et al., 2007). Harmful effects of excessive diversification can be attenuated by firms that develop sufficient competencies around their core technologies (Kim et al., 2016).

1.2.4. The geography of innovation and Economic Complexity

Geographies have characteristics that favour the emergence of specific innovation patterns shaped by relatedness. This is because geographical spaces agglomerate particular

technologies, capital, institutions, and skills that shape how innovations are created (Christopher Freeman, 1987; Lundvall, 1992).

Based on the premise that a country is composed of a collection of firms, Hidalgo, Klinger, Barabási, and Hausmann (2007) incorporate the concept of relatedness into a network perspective to show that countries are also strongly influenced by it. Using international trade data, the authors show that the relatedness between products affects how countries explore new goods. Over time, countries move through a “product space” by developing goods close to those for which they already have established some comparative advantage, whereas more distant products are hardly reached (Hidalgo et al., 2007)⁴.

The product space framework was later generalised as the knowledge space (Rigby, 2015), which extends the initial framework to other domains. Other than products, Balland (2016) stresses that the knowledge space framework has been used, for example, for representing the relatedness between technologies found in a given patent dataset, forming then a technological space (Rigby, 2015), or for representing scientific domains found in a given publication dataset, through the scientific space (Boschma, Heimeriks, & Balland, 2014). The idea behind these distinct knowledge spaces is based on the same premise: Relatedness is a reliable indicator of the costs of moving from one type of knowledge to another.

The knowledge space framework is commonly depicted as a network in which nodes represent knowledge categories, such as technological or scientific fields, and the position and links between nodes depict their relatedness (Balland, 2016). This framework became a pillar of Economic Complexity (Hidalgo, 2021) and was rapidly incorporated into the location-based analyses of the Evolutionary Economic Geography (EEG) literature. Its consistency to measure

⁴ Another core idea linked to the approach presented in Hidalgo et al. (2007) refers to complexity. The authors highlight that some products are more sophisticated than others, demanding a larger number of different inputs, whereas others are more easily produced. Accordingly, these more-sophisticated products are located in densely connected core areas of the product space, whereas the simplest ones occupy the less-connected peripheral areas. This sophistication is later labelled “complexity” in Hidalgo and Hausmann (2009). The authors show empirically that countries’ complexity, measured by their diversity of capabilities and the interactions between these capabilities, is correlated with countries’ level of income, whilst deviations from this relationship are predictive of future growth. Both relatedness and complexity would thus explain why poor countries hardly develop more competitive exports and fail to converge to the income levels of rich countries: As they have a narrow set of competitive advantages and these are mostly in products with low complexity, they fail to traverse their product spaces to become more diverse and/or develop the more profitable complex products.

the relation between the emergence and development of distinct forms of knowledge gave new tools to the EEG to explain historical patterns of spatial concentration of knowledge.

The empirical evidence produced in this literature shows that knowledge concentration patterns are strongly linked to existing local specialisations, including technological trajectories seen in cities (Boschma, Balland, & Kogler, 2014; Rigby, 2015), regions (Buarque, Davies, Hynes, & Kogler, 2020; Colombelli, Krafft, & Quatraro, 2014; Ejdemo & Örtqvist, 2020; Van Den Berge & Weterings, 2014), and countries (Hidalgo & Hausmann, 2009; Petralia, Balland, & Morrison, 2017). These strong links are due to the fact that relatedness affects the entrance and exit of firms in specific geographies. Boschma, Balland, et al. (2014), for example, show that the entry probability of a new technology in US cities increases by 30% if the level of relatedness with existing technologies in the city increases by 10%, while the exit probability of an existing technology decreases by 8%. Rigby (2015), also considering US cities, shows how relatedness changes over time and the key role that it plays in maintaining acquired competencies. Similarly, Kuusk and Martynovich (2021) find that changes in inter-industry relatedness influence regional employment growth.

The consideration for how distinct technologies relate to each other has also been crucial to explain variations in local technological development (Balland, 2016; Boschma, Balland, et al., 2014; Petralia et al., 2017; Rigby, 2015; Santoalha, Consoli, & Castellacci, 2021) and the emergence of new industries (Colombelli et al., 2014; Feldman, Kogler, & Rigby, 2015; Neffke, Henning, & Boschma, 2011; Tanner, 2016). It also explains, for example, the emergence of knowledge related to radical technologies. In particular, Tanner (2016) shows that existing knowledge in areas relevant to fuel cell industries explains the local emergence of this disruptive technology. The higher the variety of specialisations of regions in related fields, the more likely it is that they develop this type of industry (ibid). The literature also highlights that relatedness itself can change over time due to local-specific factors. In particular, Juhász, Broekel, and Boschma (2021) find that the co-location of technologies and relatedness not only change over time but also affect each other: The more two technologies overlap within spatial distributions, the greater is the change in their relatedness. As relatedness between two technologies increases, so does the probability of them being co-located in the same geographical space (ibid).

Finally, the role of relatedness has also been linked to the innovation performance of geographies. The literature includes evidence that the exploration of related technologies is linked to higher production of knowledge (Kogler, Rigby, & Tucker, 2013) and higher innovative output (Aarstad, Kvitastein, & Jakobsen, 2016; Castaldi, Frenken, & Los, 2015; Delgado, Porter, & Stern, 2014; Solheim, Boschma, & Herstad, 2018; Solheim, Boschma, & Herstad, 2020). Conversely, unrelated knowledge is linked to technological breakthroughs (Castaldi et al., 2015) and radical innovations (Solheim et al., 2018; Solheim et al., 2020). The diminishing returns of focusing on related innovations are also recognised (Antonelli et al., 2010; Ejdemo & Örtqvist, 2020), with intermediate relatedness levels being pointed out as the optimal diversification strategy for geographies (Alshamsi, Pinheiro, & Hidalgo, 2018).

1.3. What is Artificial Intelligence

1.3.1. AI's history and technological development

The term “Artificial Intelligence” was proposed in the 1950s to refer to the creation of a machine with abilities of perception, cognition, and action close to human abilities (Li & Jiang, 2017; Russell & Norvig, 2016). Since this initial proposition, AI went through a sequence of periods with high as well as low attention and activity, which create some scepticism when discussing the potential of this technology today.

The first “golden age” of AI was marked by worldwide investment into AI research. It occurred in the two decades following the initial proposition of the term AI. A first “winter” followed between 1974 and 1980. During this period, many governments, including the ones that were leading AI research such as the British and Japanese governments, withdrew their funding for AI development (Li & Jiang, 2017; WIPO, 2019). AI had another boom in the 1980's associated with the development of knowledge-based expert systems, which was followed by yet another downturn. This second AI “winter”, which took place from 1987 to 1993, was the result of a clash of expectations: The once-promising expert systems had severe computational limitations that could not be resolved due to the necessity of pre-defining sets of rules in these systems in each application. This requirement makes expert systems very inflexible when dealing with new or unexpected situations.

During the 1990's, AI entered a new era, characterised primarily by the arrival of data-learning techniques, such as machine learning. These techniques “learn” their own rules from the data

received and do not require an expert to define strict sets of rules a priori. This makes data-learning techniques the antithesis of rule-based techniques, such as expert systems. Machine learning techniques developed rapidly, in particular, with the proposition of deep learning models by Hinton and Salakhutdinov (2006). The progress of advanced algorithms related to deep learning resulted in outstanding performance gains, while at the same time dramatically reducing computational costs in training computer models (Brynjolfsson, Mitchell, & Rock, 2018; Cockburn et al., 2018; Li & Jiang, 2017). In addition to this, the development of high-performance parallel computing chips and the availability of large datasets increased AI's performance and applicability (Klinger et al., 2018; Li & Jiang, 2017). This consolidated the decline of logic-based models⁵, and AI became data-driven (WIPO, 2019).

Therefore, the current boom of AI is strongly related to the use of data-driven methods and to the increasing availability of data. There is a ubiquity of information and communication technologies (ICTs) in business and society (Alcácer, Cantwell, & Piscitello, 2016). AI can become similarly ubiquitous since it can be deployed in every place in which there are ICTs. This potential ubiquity seems achievable when one considers the recent flourishing of AI-related inventions, which are shifting AI from theory to commercial application (WIPO, 2019). This shift has already enabled AI to expand from specific applications in computer science to include a wide range of distinct functional applications involving the most diverse areas and industries. Li and Jiang (2017) underline that AI is already being a powerful tool to make people's lives better by reducing the human workload while improving work experience and conditions. Among the areas already using AI technologies extensively, the authors highlight engineering, business, medicine, weather forecasting, manufacturing, and service systems. Similarly, WIPO (2019) points out that the greatest opportunities for AI existing today lie outside the software industry, in sectors such as agriculture, healthcare, and manufacturing.

This diffusion of AI into a wide range of technological sectors makes it more difficult to define the technological domains of this technology. AI has become an umbrella term to describe technologies that incorporate some kind of human-like intelligence in very narrow tasks. AI

⁵ It is important to understand clearly the fundamental difference between knowledge-based systems and machine learning techniques. The former has as its main focus the development of a priori rules to analyse an amount of data. The latter, on the other hand, works the other way around, using the a priori data to infer "rules" or predictions from the patterns found. Thus, to the extent that the complexity of problems considered by AI has increased, it is visualized that machine learning techniques started to stand out as the main solution because they managed to minimize the extensive programming required a priori to deal with more complex problems.

technologies also share some other important characteristics: They are modular (Nilsson, 2009), transversal (Righi et al., 2020), and digital (D. J. Teece, 2018). Modularity means that AI can be combined with other technologies, while transversality refers to the possible use of AI across a variety of technological sectors. The digital aspect, in turn, means that AI has particular properties linked to convergence and generativity⁶ that allow, among other things, its embeddedness into non-physical products and the emergence of combinatorial innovations (Yoo, Boland Jr, Lyytinen, & Majchrzak, 2012). AI can be used to transform existing nondigital products into digital ones, and, through recombination with other digital products, lead to the creation of further innovations with new functionalities.

One can identify some aspects of AI's technological evolution in Figure 1. It highlights distinct keywords that have been used in AI-related publications⁷ over the past decades. It shows the decline of expert-systems since the 1990's, which is followed by the emergence of keywords related to data-driven techniques (e.g., neural networks, machine learning, deep learning⁸). One can also observe the increasing use of terms related to recent technologies, like "big data" and "internet of things" in the most recent period.

⁶ Convergence refers to the action potential of embedding digital technologies in nondigital artefacts. This enables coupling several distinct products or tools into one. Generativity, in turn, refers to the action potential of digital technologies for producing innovations characterised by unprompted change driven by large, varied, and uncoordinated actors (Yoo et al., 2012).

⁷ The bibliographic data used in this work comes from a search conducted on 11/13/2018 in the Web of Science database. Using the term "Artificial Intelligence" and the "Topic" search option, a total of 33,094 publications were found and had their information collected. The software VOSviewer was applied to this data to create this figure. For better visualisation, very similar terms (e.g. "multi-agent system", "multi-agent systems" and "multiagent system") were excluded, as well as the keyword "Artificial intelligence", used for collecting the documents. A description of VOSviewer, as well as its adopted algorithms, is highlighted in van Eck and Waltman (2009).

⁸ For more information about these and other relevant AI techniques and their definitions, please refer to Table 2 from Chapter 2.

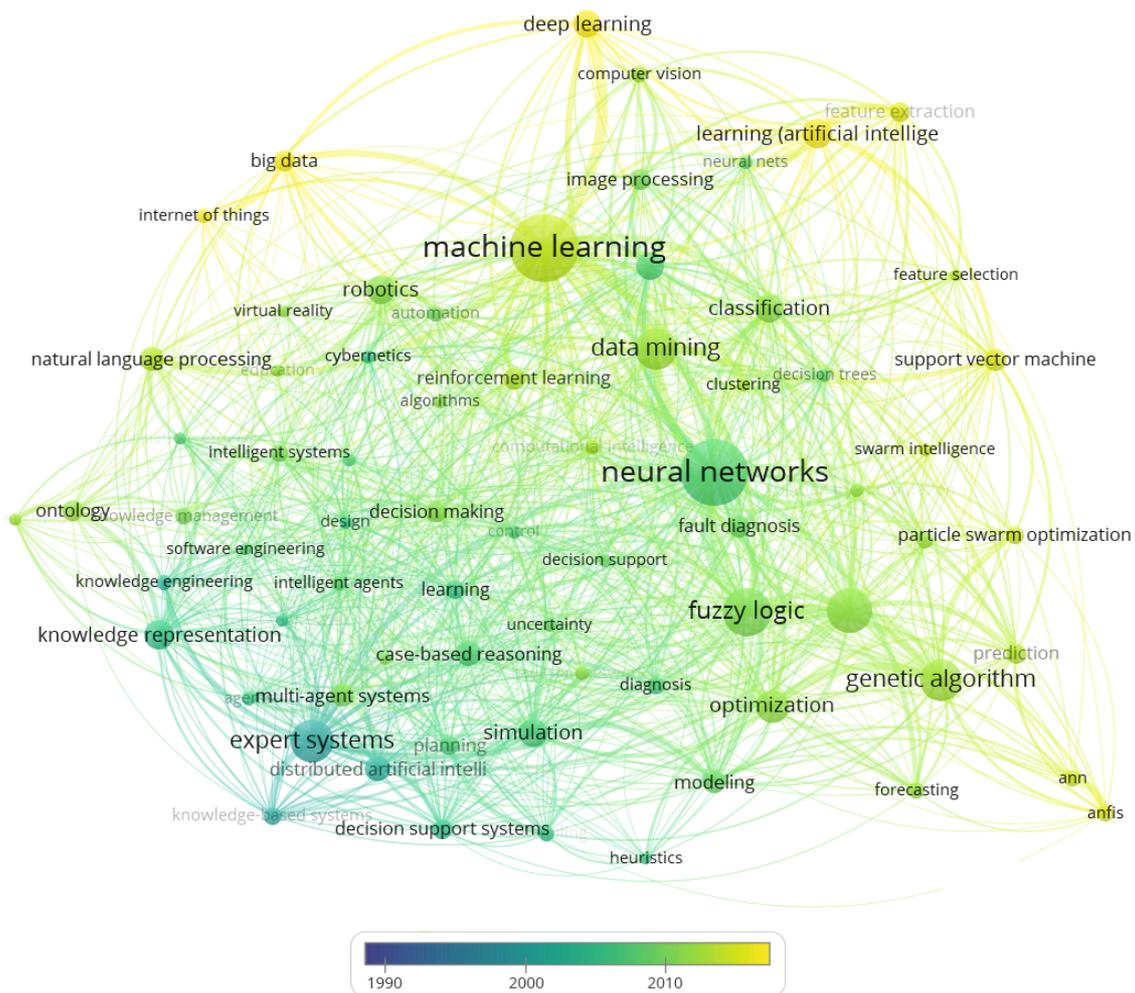


Figure 1: Author keywords co-occurrence over the years.

1.3.2. AI's potential and international policy responses

The high performance of AI technologies in narrow tasks and the widespread development of AI-related applications influence the current high expectations about this technology. For instance in the policy arena, one can find the assumption that AI is starting to disrupt entire business value chains by automating existing business processes, uncovering new value from data, and augmenting human decisions and actions (European Commission, 2017). Similarly, it is proposed that AI's automation of complex tasks has the potential to enhance or turn obsolete the existing knowledge and capabilities of firms (Paschen, Pitt, & Kietzmann, 2020).

At the same time, it is recognised that much of AI's potential is yet to be deployed. Fujii and Managi (2018) stress that by enhancing and creating efficiencies, AI technologies can play a critical role in the future in improving social well-being in areas as diverse as economic development, public welfare, environmental protection, service robotics, healthcare, education, low-resource communities, public safety, employment, and entertainment.

Harhoff et al. (2018), in turn, highlight AI's potential for solving relevant societal issues in areas such as climate change (energy efficiency), medicine (new diagnostic possibilities), and transportation (more efficient logistics chains and intelligent mobility solutions).

Cockburn et al. (2018) point out that AI also holds the potential to be an "Invention of a Method of Invention" (IMI). In this way, AI can be used to "automate discovery" across several different domains of application. Cockburn et al. (2018) highlight that there has been only a handful of IMIs⁹ as general as AI's learning techniques, having all of these a huge impact on humankind, both directly in the form of inventions and indirectly through the reformulation of technological paradigms. The authors also anticipate that AI will likely change the nature of scientific and technical development, having large and unanticipated impacts across the economy and society in general.

In economic terms, the European Commission (2017) points out that AI can contribute to an increase of 13.33 trillion euros in the global economy between 2017 and 2030, of which 55% should originate from productivity gains and 45% from consumption-side effects. In response to this growing relevance and its related possible effects on the world economy, countries have increasingly invested in policies to support their national development of AI. Harhoff et al. (2018) cite the United States, Japan, South Korea, France, and Finland as examples that have dramatically increased their investments in this technology in the last decades. In China, AI has become a high political priority with clear targets and strategies defined for overtaking the United States and making China the new world leader of AI by 2030. Similarly, the European Commission is trying to define a joint strategy so that European countries do not fall behind, calling on EU Members to develop national and collective action plans for the diffusion and progress of AI at the EU level (Harhoff et al., 2018).

There seems to be a sense of urgency given in this AI "global race", which stems from the understanding that countries mastering this technology early will gain advantages over global markets and industries (Klinger et al., 2018). The European Commission (2017) highlights that the impact of AI on productivity is so extensive that companies that fail to adopt it can rapidly lose a significant market share or even become obsolete. Cockburn et al. (2018) point out that if there are increasing returns to the development of primary AI inputs such as data, early or

⁹ Among these IMIs that are arguably as general as AI, Cockburn et al. (2018) cite optical lenses (which allowed the invention of glasses, telescopes, and microscopes) and digital computing.

aggressive entrants are likely to create substantial and long-lasting competitive advantages merely through the control of these inputs. Klinger et al. (2018) present a similar view, noting that the window of opportunity for new entrants may simply close as new machine learning research centres become dominant. The authors show evidence that China is already a world leader in AI in some areas, while early prominent European countries are losing their relevance more and more rapidly. The analysis presented in WIPO (2019) seems to confirm this trend.

The loss of European prominence, in turn, is alarming not only for the affected countries but also for the AI progress in general. Harhoff et al. (2018) point out that whenever AI is discussed in Europe, policymakers highlight the importance of creating a “third path” for its development, which should differ from both the “Chinese state-led” and the “Silicon Valley market-driven” approaches. This third path should follow specific European principles including, among other values, clear criteria of ethics and regulation and the vision of joint development at the EU level (ibid).

Such criteria often restrict one of the fuels of AI, namely data. Data is essential for developing commercial AI applications (Furman & Seamans, 2018; Taddy, 2018) and a strategic input for local industrial development (Goldfarb & Trefler, 2018), which makes this type of input potentially relevant to the development and diffusion of AI. It is recognised that a larger volume of data enables better business performance (Goldfarb & Trefler, 2018) and the emergence of specific AI innovations (Beraja, Yang, & Yuchtman, 2020). The European guidelines, as exemplified by the General Data Protection Regulation¹⁰, heavily restrict the collection and storing of data in the EU, which may potentially hinder the creation of some AI applications. At the same time, the European path is a good solution for the deployment of more ethically-based AI. Concerns over data collection often stress potential violations of consumer privacy (Tucker, 2018) and the need for policies that could prevent security problems arising from the use of this data (Jin, 2018).

1.4. Research questions

The research questions that guide this thesis focus on knowledge dynamics related to AI. An important relevant characteristic of AI in this sense is its potential to be a General-purpose Technology (GPT). GPTs technologies are pervasive, have an inherent potential for technical

¹⁰ Available at <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02016R0679-20160504>.

improvements over the years, and induce the creation of related “innovational complementarities” (T. F. Bresnahan & Trajtenberg, 1995). From a macroeconomic perspective, these technologies solve the problem of diminishing returns: The cluster of a GPT and its related innovations bring technical and market opportunities, and consequently economic growth¹¹. Albeit many authors discuss AI as being a GPT (e.g., Brynjolfsson et al. (2018); Cockburn et al. (2018); Klinger et al. (2018); Lipsey, Carlaw, and Bekar (2005); Trajtenberg (2018)), this thesis departs from this perspective for two main reasons.

The first is that the GPT theory is strongly connected to a neoclassical macro view on the role of technologies. This makes the insights from this theory less useful to an evolutionary perspective that focuses on knowledge dynamics. The second reason relates to suggestions on how an evolutionary perspective could be incorporated into the GPT theory, presented in Cantner and Vannuccini (2012, 2017). Particularly, this literature proposes adopting an evolutionary view on competition between technologies at the micro level to explain the macro effects seen for a GPT. This would explain macro effects as emerging by competition between distinct technologies to become dominant (Cantner & Vannuccini, 2012) or through upstream technological competition for downstream markets (Cantner & Vannuccini, 2017). I argue that this view, for the AI case, would unnecessarily direct the analysis towards explaining successful technological trajectories with AI as a result of technological choices. However, the main focus of this thesis is to understand how knowledge dynamics (i.e., instead of technological dynamics) are linked to the emergence of successful technological trajectories with AI. This particular focus allows, for example, to explore the possibility that local technological trajectories with AI may become successful because they are well linked to existing knowledge bases, even if they are not the best existing technological alternatives nor follow AI’s dominant design.

In this thesis, knowledge dynamics related to AI are analysed under two complementary perspectives: Local geographically bounded dynamics and firm-specific dynamics. In the first part of the thesis, the research questions focus on local geographically bounded learning by identifying countries leading AI’s development (Chapter 2), and further analysing their

¹¹ The GPT cluster and the remaining “application sectors” are impacted in different ways. While the former leads to the development of a vertical externality and improvements in the GPT’s performance, the latter produces a horizontal externality and its impact depends on the technological distance of each sector to the GPT core. More distant technological domains would need a greater number of ICs to be developed before adopting the GPT, generating laggard sectors during the development of such technologies (T. Bresnahan, 2010).

technological trajectories (Chapter 3). The research identifies the specialisations of selected countries as a reference to understand how they influenced the emergence of AI-specific specialisations. In the second part of the thesis, the research questions focus on how AI learning influences the innovative activities of firms. More precisely, the research analyses how the introduction of AI affects firms' subsequent innovations in terms of relatedness and performance (Chapter 4).

All the mentioned paper-based chapters use patents as an indicator for innovation. The first paper, in particular, proposes a strategy for identifying AI patents that is adopted in all three papers. It also compares this strategy to alternative approaches. Other specificities of these three papers are highlighted in Section 1.5.

1.4.1. The geographical learning of AI

The first part of this thesis investigates which countries lead the development of AI globally, and how distinct the development of AI is across these countries. This is arguably the first step for understanding how technological learning is affected by geographic-specific factors. This kind of understanding contributes not only to advancing the evolutionary view on technological change, but also offers insights about suitable innovation policies aimed at the local development of AI. Given AI's potential to increase the difference in economic performance across countries, it is highly desirable that successful policies are implemented worldwide so that less developed countries do not fall further behind economically by missing AI-enabled technological opportunities.

Two perspectives are considered to identify countries that lead AI development globally. The first perspective considers which countries are leading in the creation of AI innovations, whereas the second identifies countries leading in terms of commercializing AI innovations.

Accordingly, the first research question addresses:

1. Which countries are leading the creation of AI innovations? Which countries are leading in the commercialization of AI innovations?

It is assumed that differences in the identified leaders across the two perspectives are a first indication of specific local knowledge dynamics with AI. In this case, it is assumed that these factors are primarily linked to institutional frameworks. Moreover, the thesis aims at

understanding whether differences in local learning influence the AI innovations created in a country. Considering this, the following research question is proposed:

2. What are the technological specialisations in AI that are distinct for the countries that globally lead the development of AI innovations?

A technological specialisation refers to a revealed comparative advantage that is developed in a specific technology. Considering knowledge at the country level, it means that the country has a production level of knowledge that is higher than the global average in a given technology. That is to say that the country is innovating in this particular technology at a higher rate than the global average. The technologies examined here are the typical AI techniques used to create AI innovations, e.g., expert systems, neural networks, deep learning, support vector machine, etc. The techniques considered are based on the typical AI techniques proposed in WIPO (2019, pg. 24).

Specialisations in distinct AI technologies across geographies are considered a second indication of the influence of local-specific factors on knowledge dynamics related to AI. In this case, it is assumed that these differences are primarily due to distinct local knowledge bases.

If such differences exist, it is important to analyse how they influence local technological trajectories with AI. Local knowledge bases influence how a new technology is learned. Knowledge bases can indicate how similar a new piece of knowledge like AI is in comparison to the existing local knowledge. They also reflect local factors that affect learning. This is because geographies agglomerate particular technologies, capital, institutions, and skills that shape knowledge development (Christopher Freeman, 1987; Lundvall, 1992).

When focusing on countries leading AI development, one can better understand how AI's successful technological trajectories are created locally. Are AI-leading countries learning specifically about the main knowledge linked to this technology, or are they combining AI with their existing knowledge bases? The geographical perspective on relatedness is particularly useful to answer this kind of question. It allows identifying what are the technologies most related to AI and comparing them to the specialisations held by selected countries (in this case, the ones leading AI development). This perspective also allows comparing how countries' existing specialisations relate to their AI-specific specialisations.

Thus, the thesis considers next AI's innovation patterns to understand the development of this technology. Here, the research employs a "recombinant innovation" process perspective, which considers innovations as being created by combining distinct pieces of knowledge in novel ways (Frenken et al., 2007; Van den Bergh, 2008). Thus, AI innovations are assumed as being created through combining distinct pieces of knowledge.

Looking specifically at AI innovations, the thesis disentangles what is the "state-of-the-art" AI knowledge in distinct intervals of time. This disentanglement is based on two mechanisms. The first, proposed in Dosi and Grazzi (2006), highlights that innovations reflect the technical routes taken to solve the problems typical of a technology. The second, proposed in Arthur (2009), highlights that successful innovations are repeated in following innovations. Thus, it is expected that the use of distinct technologies in AI innovations will create specific innovation patterns over time. In these innovation patterns, the technologies that better solve the problems typical of AI will be repeated. This repetition shows that these technologies are more important than others to the deployment of AI innovations. As they are arguably linked to the best technical solutions¹² of the typical AI problems, these technologies are considered the "state-of-the-art" of AI knowledge in a given interval of time.

These successful technologies are termed "AI-core technologies". It is argued that they reflect AI's technological progress, and the best technical routes to solve AI's typical problems in a given interval of time. Accordingly, the following two sets of research questions are proposed:

3. Which technologies are used to create AI innovations? Are there any "AI-core technologies" essential to the development of AI innovations?
4. Do innovation patterns and "AI-core technologies" change over time?

By answering these research questions one can better understand what were the best technical solutions deployed in AI innovations over time, and whether they changed as AI evolved. This allows the comparison of AI's technological progress – here understood as the accumulation of knowledge about a technology – to countries' innovation patterns in their local learning of this technology.

¹² These best technical solutions can be broadly linked to the idea of dominant design discussed previously. However, whereas dominant design takes a more definitive assumption in the sense that it is supposed to last, the "best technical solutions" highlighted here refer to more provisory solutions that can still be further improved and changed through technological development.

To make this comparison, the specific AI innovations that leading countries developed are considered. More specifically, the thesis analyses how these countries deployed “AI-specific specialisations” over time through the development of AI innovations. Conversely to Research Question 2, all technologies are considered now (instead of just the ones related to AI). The following research question is proposed:

5. How do changes in AI’s global innovation patterns affect the development of “AI-specific specialisations” by countries leading AI development?

It is expected that countries leading AI’s development are to be particularly prone to adopting its most successful technical routes. The absence of such relation would arguably indicate that countries leading AI development are not necessarily leading in the progress of AI-specific knowledge. That is to say that these countries are not using the state-of-the-art knowledge accumulated about AI globally. As a result, the AI innovations created locally in these countries do not follow the most successful technological trajectories taken with this technology.

Finally, another important factor to understand the differences in the geographical development of AI concerns the identification of how this technology is incorporated into existing technological trajectories. In this sense, one needs to separate countries’ “general” technological trajectories from their “AI-specific” trajectories. This is done by isolating the specialisations that countries have when all of their innovations are considered (i.e., “General specialisations”) from the specialisations they have when only their AI innovations are considered (i.e., “AI-specific specialisations”)¹³. Assuming that local innovations are a proxy for the knowledge available locally, “General specialisations” reflect the broad knowledge base that a country has, whereas “AI-specific specialisations” reflect the knowledge that a country has in AI. Hence, the last research question proposed under the geographic perspective is:

6. How do the “AI-specific specialisations” and “General specialisations” of a country interact to generate new local technological trajectories in AI?

Based on the findings from Tanner (2016) for fuel cell industries, it is expected that countries leading AI development have “General specialisations” in a variety of fields highly related to

¹³ This measure is similar to the one adopted in Hidalgo et al. (2007), which looks at the comparative advantages that countries have in the exportation of products. But here, instead of considering the products that countries export, the focus is on the innovations that countries produce.

AI when they start learning about this technology. This would explain the technological leadership of these countries in AI. But, as AI knowledge is accumulated, it is expected that “AI-specific specialisations” are created locally in fields that are technologically relevant to AI learning, regardless of the direction of the “General” technological trajectory. This differs from the alternative in which “AI-specific specialisations” would emerge in fields following countries’ existing specialisations.

1.4.2. AI learning from a firm perspective

To complement the geographical perspective on AI learning, it is necessary to understand the effects of AI learning at the micro-level. Accordingly, the second part of this thesis aims at comprehending how AI affects firms’ absorptive capacity and technological trajectories. It is also considered whether learning is affected in the same way across distinct sectors and whether learning can be improved through AI adoption. For the latter, the effects on firms’ innovative performance are analysed.

Innovations are used as a proxy for understanding how firms’ learning is affected: If innovations following AI adoption are different from the firm’s previous innovations, it is assumed that new knowledge domains are being learned; if they are increasingly similar, it is inferred that the firm is filling its existing knowledge gaps. In the context of this research, “AI adoption” is considered as the introduction of an AI innovation – measured by patent applications – into a firm’s technological portfolio. This includes both the possibility of a firm developing an AI patent by itself or acquiring such a patent from another company.

The first issue concerns identifying what are the effects that AI has on the assimilation of new knowledge. AI may allow firms to explore unrelated knowledge domains – e.g., through its “Invention of a Method of Invention” potential, or by offering firms the possibility to enter into a new technological paradigm – or to learn more about related domains – e.g., by uncovering new patterns from datasets related to the companies’ operations, or by making traditional learning more efficient through automation of some activities. Accordingly, it is examined how the adoption of AI affects firms’ technological trajectories, which are proxied by knowledge relatedness. The focus is on understanding whether innovations created after AI adoption are similar or not to firms’ previous knowledge bases. Thus, the following research question is proposed:

7. Does the introduction of an AI innovation into a firm technological portfolio affect the relatedness of its subsequent innovations?

Given the expectations about AI's potential, it is assumed that its adoption is to be linked to a significant change in firms' technological trajectories. Thus, relatedness significantly increases or decreases due to AI adoption.

Yet another relevant aspect concerns the possibility of changes in relatedness being different across firms from distinct technological sectors. AI may affect learning differently according to the distance of firms' knowledge bases to AI knowledge. The three dimensions of knowledge that encompass the concept of relatedness are explanatory of how these differences can play out. For firms from technological sectors that are related to AI, knowledge proximity, knowledge commonalities, and complementarity of knowledge can contribute to an increase in relatedness. But for firms from sectors technologically distant from AI, AI knowledge will be a new piece of unrelated knowledge. In this case, only knowledge complementarities can arise. This dimension, which arises from the need to use distinct technologies together, would emerge when firms combine AI with another piece of their existing knowledge. This dimension alone may not be enough to affect relatedness, in case firms are very distant from AI and, therefore, have difficulties innovating with unrelated knowledge. There is thus the possibility that relatedness remains unchanged or even decreases (i.e., due to AI being an unrelated piece of new knowledge).

Understanding these dynamics is particularly important when one considers the broad range of technological sectors that may be targeted by AI-innovation policies. Hence, it is important to understand if knowledge dynamics created by AI adoption affect learning differently across distinct technological sectors. Considering this, the following research question is proposed:

8. Do firms from knowledge domains close and distant from AI's knowledge change the relatedness of their innovations in the same way following the introduction of an AI innovation?

Knowledge domains are proxied by technological sectors. It is assumed that different technological sectors have distinct knowledge bases as a result of the technologies and routines that they typically adopt. By separating firms' innovations according to their sectors, one can measure how similar the innovations from a specific sector are to AI innovations.

Finally, another way in which AI may impact firms' knowledge dynamics is by increasing their capacity to learn, which would be reflected in an increase in firms' innovative performance. AI can affect firms' absorptive capacity in two main ways: Due to its "Invention of a Method of Invention" potential, or by affecting the relatedness of firms. Changes in relatedness are recognised as influencing the creation of knowledge and innovations (Kim et al., 2016; Leten et al., 2007; Weber & Weber, 2010). Accordingly, the last research question proposed is:

9. Does the introduction of an AI innovation into firms' technological portfolios lead to a significant change in their innovative performance?

It is expected that such change occurs due to the specific affordances linked to AI (e.g., "Invention of a Method of Invention" potential, digital technology that can be recombined with others, transversal technology, etc.).

1.5. Overview of chapters

This section highlights the main individual contributions of each of the three scientific papers that compose this thesis¹⁴. An overview of the three papers is presented in Table 1.

¹⁴ Although the scientific papers are unified under the broad focus of this dissertation on learning, the terminology presented in this section is specific to the arguments presented in the papers. This results in a higher focus on terms linked to innovation (which is overall understood in this dissertation as the result of learning and recombining knowledge in novel ways), technological capabilities (a.k.a. specializations or competitive/comparative advantages, used as a proxy for the accumulation of knowledge in specific fields or technologies), and technological portfolio (which is used as a proxy for the knowledge accumulated by companies).

Chapter	Subject of analysis	Main focus	Research questions addressed	Data sources	Method	Temporal scope
2	Countries (all)	Identify which countries are leading AI development, and which AI technologies they apply in their innovations.	<p>1. Which countries are leading the creation of AI innovations? Which countries are leading in the commercialization of AI innovations?</p> <p>2. What are the technological specialisations in AI that are distinct for the countries that globally lead the development of AI innovations?</p>	Patents proxied by patent office (from Patstat 2017)	Indicator-based analysis considering a specialisation index (RCA) and two novel patent-based indicators.	From 1979 to 2015
3	Selected countries (United States, Japan, South Korea, and China)	Identify how successful technological trajectories with AI emerge in leading countries.	<p>3. Which technologies are used to create AI innovations? Are there any “AI-core technologies” essential to the development of AI innovations?</p> <p>4. Do innovation patterns and “AI-core technologies” change over time?</p> <p>5. How do changes in AI’s global innovation patterns affect the development of “AI-specific specialisations” by countries leading AI development?</p> <p>6. How do the “AI-specific specialisations” and “General specialisations” of a country interact to generate new local technological trajectories in AI?</p>	Patents proxied by inventors' location (from Patstat 2019)	Technological space perspective coupled with specialisation indices.	From 1974 to 2018
4	Firms (MNEs)	Analyse how AI adoption affects the creation of innovations at the firm level.	<p>7. Does the introduction of an AI innovation into a firm technological portfolio affect the relatedness of its subsequent innovations?</p> <p>8. Do firms from knowledge domains close and distant from AI’s knowledge change the relatedness of their innovations in the same way following the introduction of an AI innovation?</p> <p>9. Does the introduction of an AI innovation into firms’ technological portfolio lead to a significant change in their innovative performance?</p>	Patents (from Orbis IP) matched to company data (from Orbis) with separately built historical ownership structures (from Orbis-Zephyr)	Matched-pair analysis combined with an extension of the Difference-in-Difference method to estimate the proposed effects.	From 2011 to 2019

Table 1: Overview of the dissertation’s Chapters by their main characteristics.

Furthermore, given the scientific call towards the reuse of scholarly data – formalised more strongly through the FAIR¹⁵ principles presented in Wilkinson et al. (2016) – the data¹⁶ and codes used in each of the three scientific papers are made publicly available. This is done through GitHub repositories¹⁷.

The three chapters are discussed in more detail in the next subsections, which outlines briefly the aims, methods, and empirical findings of each paper, as well as the links between them.

1.5.1. Patenting Patterns in Artificial Intelligence: Identifying National and International Breeding Grounds (Chapter 2)

Chapter 2 analyses technological trends and institutional differences in AI exploration with the main focus of identifying relevant players in the so-called “global AI race”. These players are identified according to their performance in two aspects, namely the production of AI innovations and the internationalisation in the commercialisation of these innovations. The use of patent registers as a proxy for innovation lies at the core of the analysis, so that both aspects are focused on the protection of AI innovations through Intellectual Property (IP) rights.

In addition to considering a classical indicator from the economic literature – namely the Revealed Comparative Advantage (RCA) index –, two novel indicators are proposed in the paper: i) the “National Breeding Ground” (NBG) indicator, for identifying countries that are specialised in AI and are also relevant markets for corresponding IP protection; and ii) the “International Breeding Ground” (IBG) indicator, for identifying countries that both attract and export¹⁸ AI innovations through IP protection. The considered indicators are also applied to

¹⁵ With FAIR standing for: Findable, Accessible, Interoperable, Reusable.

¹⁶ Within the legal possibilities. Due to property issues, the complete datasets from Papers 2 and 3 are not made publicly available (they are proprietary from the European Patent Office [EPO] and Bureau van Dijk [BvD], respectively). This legal limitation is partially addressed by explaining how to recreate the same datasets using the proprietary databases, or through the offering of simplified tables based on non-proprietary data (e.g., indicators calculated based on the said data). Moreover, the R codes shared allow both reproducing the complete empirical analyses. They also contain details about the description steps taken. They include comments thoroughly, made with the intent of facilitating accessibility and reproducibility.

¹⁷ The repositories are, following the papers’ order presented in this section, available at: <https://github.com/matheusleusin/Patenting-Patterns-in-Artificial-Intelligence> (Paper 1); https://github.com/matheusleusin/Paper-The_Emergence_of_Artificial_Intelligence (Paper 2); and [https://github.com/matheusleusin/AI and MNEs](https://github.com/matheusleusin/AI_and_MNEs) (Paper 3).

¹⁸ The “attract” and “export” of innovations, considered in the IBG indicator, refers to the possibility of extending the IP rights of a patent by registering it in other patent offices. Thereby, the IBG indicator identifies countries that have many AI patents registered nationally but were developed abroad (i.e., they “attract” international IP) and that extend the protection of their own AI patents abroad (i.e., they “export” their national IP).

identify technological trends in the use of distinct AI techniques in the creation of AI innovations.

In contrast to the other third chapter from this dissertation, which use inventors' location as a proxy for identifying the sources of knowledge creation, this investigation uses patent office information for identifying the origin of innovations¹⁹. This differentiation is justified by the focus on the IP rights of AI. Hence, the use of this patent office information allows differentiating the markets where the inventions were first protected (through priority filings), from the markets where the protection of these inventions was extended. This allows understanding AI's global development under two distinct perspectives, namely one identifying the main countries in the creation of AI knowledge (NBGs), and one comprehending the main countries in the international commercial exploration of this knowledge (IBGs).

The results show that NBGs and IBGs only overlap to a limited extent. The USA, together with Australia and selected European countries, are found to be the main actors in the internationalisation of AI innovations (IBGs). China, in turn, leads in terms of national development of AI innovations (NBGs), followed by the United States as well as some Asian and Eastern and Southern European countries. China appears as a strong leader from an NBG perspective, whereas the United States leads from an IBG perspective. These findings signal major structural differences in the international IP protection of AI, with some countries focusing exclusively on developing their domestic markets, while others focus on a broader international exploration. China leads by promoting AI development with a major focus on IP protection in its domestic market, whereas the United States sustains its AI progress in the international context as well.

Considerable changes also exist in the leadership of AI development globally. The results show that latecomers such as China and the United States have largely overtaken incumbents like Japan and selected European countries, increasing their lead mainly in the last two decades. The analysis of AI's technological trends also reveals that, despite the fact that a few specific AI techniques declined in use in the last decades due to their technical limitations (e.g., Expert

¹⁹ This is done by identifying the patent office in which the first register related to an invention was made (i.e., the priority filing). All following registers (also known as non-priorities) refer to an extension of the legal protection of the patent in the patent offices in which they are registered. Put together, all these registers constitute a patent family linked through the priority filing.

Systems and Rule-based learning), these techniques are still applied in the bulk of AI innovations generated recently by some latecomers. This pattern is specially seen for NBGs, i.e., for countries that focus on deploying AI particularly in their domestic markets.

1.5.2. Break-through or Break-in? How AI Becomes a Part of National Technological Trajectories (Chapter 3)

The findings from Chapter 2 highlight interesting dynamics related to AI's development. Not only have technological leaders changed drastically in the last two decades, but also AI itself evolved. This last finding, in particular, is indicated by the increasing use of biological models over time to create AI innovations. This indicates that the historical technological progress of AI – marked in the 1990's by a change from rule-based learning to machine learning methods (Li & Jiang, 2017) – was also reflected in the innovation patterns related to this technology. These changes are not uniformly reflected in how latecomer countries develop AI innovations. Some of these latecomers still innovate in AI using “obsolete” techniques, despite their technical limitations²⁰.

These findings motivate Chapter 3. They indicate the possibility of AI's technological progress being reflected in its innovation patterns. They also indicate that these innovation patterns are not entirely reflected in the way that countries develop AI innovations. Hence, Chapter 3 tries to understand i) how the technological change of AI is reflected (or not) on how countries explore this technology, and ii) the role that countries' specific specialisations play in this exploration. Based on the premise that countries leading AI development are to be the ones more sensible to adopting its newest technological developments, the paper analyses the technological trajectories²¹ of the four most prolific producers of AI innovations - namely the United States, Japan, South Korea, and China.

²⁰ This was shown to be the case for Eastern and South-eastern European countries which, despite entering the global AI race in later periods, still base a large proportion of their AI innovations in techniques that have long been declining in relevance due to their particular technical limitations. E.g., Belarus and Lithuania are shown to innovate in techniques related to Expert Systems, and Romania in techniques related to Fuzzy Logic.

²¹ There are two distinct possible interpretations for the term technological trajectory. The first one, proposed in Dosi (1982), refers to the idea of technological trajectories within the context of technological paradigms. The second one, commonly adopted in the Economic Complexity literature (e.g., see Petralia et al. [2017] and Hidalgo [2021]) refers to specific technological paths that are built due to cumulative, path dependent, and interactive process that is typical of knowledge production. The latter, therefore, focuses on the differences between technological trajectories across distinct geographies, whereas the former focuses particularly on the long path of development of a specific technology. This second paper focuses on the technological trajectory perspective typically adopted in the Economic Complexity literature.

Using a technological space perspective (Hidalgo et al., 2007), the paper differentiates between distinct dynamics occurring at the local and technological levels. Given the fact that the latter dynamics are particularly neglected in the current literature (Hidalgo, 2021; Juhász et al., 2021), the paper proposes a novel way of analysing technological dynamics. The proposed approach separates innovations related to a specific technology (in this case, AI) to analyse how innovation patterns change over time. The assumption is that these innovation patterns represent distinct technical routes taken to solve the problems typically related to a technology. These distinct routes are assumed to reflect the general progress of a technology. The innovation patterns are then compared, for the same intervals, to how the four selected countries developed new specialisations²² in AI. In this way, the new AI specialisations developed by the considered countries can be related to their “general”²³ technological trajectories, so that the relationship between the two can be better understood.

Chapter 3 also identifies a set of core technologies linked to AI innovations. This set of technologies has been shown to be dynamically changing over time, affected by AI’s technological progress. Importantly, this technological progress has little association with how AI was developed locally. Instead, AI’s local development followed countries’ existing knowledge bases (i.e., their “general” technological trajectories). This applied even in circumstances when countries’ specialisations were weakly related to AI. This pattern was shown to be reinforced as more knowledge about AI was accumulated, so that new AI specialisations developed locally increasingly coincided with countries’ existing general specialisations.

The alternative, in which local AI specialisations emerge in fields in which there are no local general specialisations, is shown to produce comparative advantages that are short-lived. Therefore, countries’ general specialisations are found to be the most relevant explanatory

²² Please note the distinct uses of the term “capability” in this paper. It refers to a specialisation measured through the RCA index. This measurement is made using two distinct datasets: For identifying local capabilities, it is used as in the classical implementation by considering all patents registered; for identifying technological-specific capabilities (i.e., the ones related to the technical routes taken with AI, here called AI’s “core technologies”), it considers only the AI patents identified, without any consideration for location information.

²³ Please note the distinction made between local AI-specific capabilities and General capabilities. Whereas the former refers to countries’ performance when AI patents are considered, the latter considers countries’ performance when all patent registers are considered.

variables of the areas in which new specialisations in AI are developed and maintained at the country level.

These findings indicate the powerful influence of the cumulative, path dependent, and interactive process that is typical of knowledge production (Dosi, 1982; Nelson & Winter, 1982). Particularly, they highlight how this process also influences the emergence of new AI-specific specialisations at the local level. They show that the production of AI knowledge becomes mostly concentrated in the technological areas in which countries hold some existing technological advantage. By revealing that this pattern holds even in the cases that local knowledge is not directly related to AI, this third chapter highlights a somewhat counterintuitive finding. This finding contrasts, for example, with empirical analyses linking the emergence of new technologies and/or industries exclusively to the presence of related local knowledge²⁴. The findings from Chapter 3 show that unrelated knowledge plays a larger role in the local emergence of AI when a country holds a competitive advantage on this unrelated knowledge. AI breaks-in locally by being combined with technological areas where countries hold competitive advantages, (largely) unaffected by their relatedness to AI.

1.5.3. The Development of AI in Multinational Enterprises – Effects upon Technological Trajectories and Innovation Performance (Chapter 4)

Chapter 4 moves from a country-level to a firm-level perspective to explore how AI adoption affects the technological trajectory of firms. By understanding how AI affects firms, one can better discern how some of the counterintuitive effects seen at the macro level in Chapter 3 emerge. In this context, I argue that AI adoption has a causal effect on firms' learning and innovative performance. This is arguably due to some of AI's technological affordances, which include a potential to be recombined with other technologies and the potential to enhance learning. I argue that AI adopters increase their learning about related knowledge. The further accumulation of related knowledge is to lead to the creation of increasingly related innovations. Related innovations are less costly to create than unrelated innovations, which is to contribute to an increase in innovative performance. AI's characteristics that make it a

²⁴ One example is Tanner (2016), who shows that knowledge in areas related to fuel cell industries explains the local emergence of this type of industry. According to the authors, the higher the variety of specialisations of regions in these related fields, the more likely it is that they develop this type of industry.

potential “Invention of a Method of Invention” may also contribute to an increase in the innovative performance of AI adopters.

Hence, the argument of Chapter 4 is that AI adoption affects firms’ technological trajectories in at least two distinct ways: It favours the emergence of increasingly related innovations and increases innovative performance. The first hypothesis is tested through an indicator of relatedness, so that the innovations created by a firm after AI adoption can be compared to its previous innovations. Thereby, it is possible to identify whether subsequent innovations are different (regarding the technologies adopted to create them), and if so, in which way (i.e., more or less related). The second hypothesis is tested by considering three distinct indicators linked to innovative performance, namely: Number of innovations owned by a firm, number of innovations owned per unit of Turnover, and R&D expenditures per innovation.

The considered effects of AI adoption are analysed by applying a matching procedure to a novel dataset combining firm and patent data for the observation period from 2011 to 2019. Because the current wave of AI adoption is skewed towards large and established firms (Zolas et al., 2021), Multinational Enterprises (MNEs) are chosen as the focus of analysis. Firms that adopted AI are matched to control firms from the same industry, company size, and with a similar age and innovation output throughout the observation period. AI effects are estimated considering the year that it was adopted by each firm using the extension of the Difference-in-Difference (DiD) method proposed in Callaway and Sant’Anna (2018, 2020).

The estimated effects indicate that AI adoption is indeed linked to a significant change in firms’ technological trajectories. AI adopters develop innovations that are increasingly related to their knowledge bases, strengthening their existing technological trajectories. The significant positive effect on relatedness is also seen for firms that are technologically distant from AI knowledge. This result is arguably linked to AI’s technological potential to enhance learning, which is an addition to AI’s value as a new piece of knowledge.

It is also noted that AI adopters increase their innovative performance significantly in terms of innovation output measured by the number of patents owned. The other two indicators considered for innovative performance support this affirmation through their measured effects, but the differences are insignificant. Altogether, the results indicate that AI influences the absorptive capacity of firms. Particularly, given the consistency of AI’s effects on relatedness across distinct sectors, it is argued that firms not only benefit from absorbing AI

as a new piece of knowledge, but also from using AI as a learning tool. This has implications for the concept of absorptive capacity, linking it to a changing ability that can be influenced by firms' technological choices. Therefore, not only knowledge dimensions, as suggested by Breschi et al. (2003), but also technologies hold the potential to influence knowledge relatedness.

1.6. Contributions, implications and future research

1.6.1. Main findings

The findings from Chapter 2 show that the leadership in the global AI race has changed in the last decades, with China and the United States surpassing Japan and Europe. The new global leaders present major structural differences in the international IP protection of their AI innovations: China promotes AI development mainly in its domestic market, whereas the United States promotes a domestic market that both attracts AI innovations from abroad and exports AI innovations created locally. Asian and some Eastern and Southern European countries resemble the IP strategy of China, whereas Western European countries follow the patterns observed for the United States (Research Question 1). The vast majority of AI specialisations deployed by countries leading AI development are related to data-driven techniques such as Machine learning and Deep Learning. These techniques reflect the latest AI developments and are arguably linked to better-performing AI innovations²⁵. But there are exceptions primarily in countries that focus on their domestic markets. Specialisations in rule-based techniques are particularly noticeable in these countries when the most recent period of analysis (2003-2015) is considered (Research Question 2 and the first part of Research Question 3).

It is found evidence in Chapter 3 indicating the existence of "AI-core technologies". These technologies changed slowly over time as AI progressed technologically, generating specific innovation patterns. In particular, AI technologies linked to "Basic communication processes" lost relevance, whereas technologies linked to the "Analysis of biological materials" gained it (second part of Research Question 3 and Research Question 4). These patterns of technological development were not reflected in how leading countries explored AI (Research Question 5). Countries deployed AI specialisations mostly in the areas in which they held a

²⁵ I.e., in comparison to techniques that are linked to earlier stages of AI development, like rule-based learning.

general comparative technological advantage (Research Question 6). When that was not the case, the newly deployed AI capability was short-lived.

Chapter 4 provides evidence that the introduction of an AI innovation into a firm's technological portfolio leads to the creation of increasingly related innovations (Research Question 7). This is true for firms across distinct technological sectors (Research Question 8). Moreover, AI's introduction leads also to an increase in firms' innovative performance (Research Questions 9).

The findings outlined above shed some light on AI's current development. Generally speaking, the findings indicate a certain gap between AI's verified potential and the expectations about this technology. It seems that AI is not yet disrupting entire sectors. Currently, AI is applied to further extend the particular affordances created by digital technologies. At the micro-level, AI allows firms to produce new related knowledge. Due to its digital aspect, resources that favour the use of ICTs (e.g., access to digital data, large datasets, computational power, access to technologies that allow the exchange of data in real-time, etc.) also empower AI. The knowledge generated by firms that adopt AI fills their existing knowledge gaps, enabling AI adopters to develop more coherent technological trajectories. As exploring related knowledge is more efficient than exploring unrelated knowledge, it is understandable that AI adopters also increase their innovative performance. Moreover, firms use AI to improve learning in the knowledge domains in which they typically innovate (i.e., to learn related knowledge).

At the country level, AI allows the renewal of existing local technological trajectories. The fact that the majority of AI's best-performing techniques are related to the use of digital resources partially explains the most recent changes in global leadership. When AI moved from rule-based learning to data-based learning, it created a window of opportunity for newcomers. Being able to combine AI with existing local specialisations was shown to be a strong factor for the emergence of new leaders. When new specialisations in AI were not combined with existing local specialisations, they were seen to be short-lived.

The complementary between the macro and micro level perspectives allows understanding the distinct impacts of AI. The micro perspective shows that AI adopters become increasingly coherent in their technological trajectories, creating innovations that are progressively related to their existing knowledge bases. For the macro level, there is evidence that AI allows countries to renew their existing technological trajectories. The literature complements these

findings with evidence showing that countries become less coherent in their technological trajectories as they progress technologically (Petrulia et al., 2017), which is explained by technological complexity. As countries become technologically more developed, the exploration of less valuable (i.e., less complex) technological areas is discontinued. Countries move towards technological sectors that increase their overall economic development. Thereby, local technological trajectories become less coherent because knowledge creation changes are based on the economic performance of technological sectors, rather than based on their relatedness to previous knowledge. This poses a great challenge to incumbent firms located in sectors with low economic potential. At the macro level, this indicates the importance of deploying AI in areas that are both part of countries' existing specialisations and bear potential for future economic development.

1.6.2. Theoretical contributions

This dissertation makes relevant theoretical contributions to the evolutionary view on learning and technological trajectories (Dosi, 1982; Nelson & Winter, 1982). The concept of absorptive capacity (Cohen & Levinthal, 1990) is adopted as the theoretical construct for understanding how learning about a specific technology takes place for two distinct types of units (i.e., countries and organisations). The analysis conducted allows improving the conceptual definitions of the mentioned construct. Particularly, two novel mechanisms are found affecting firms' absorptive capacity: i) firms' technological choices, and ii) time-dependent knowledge commonalities. Next, these two novel mechanisms are explained in detail. This includes discussing how and why they emerge, and their possible effects. It is proposed that the first mechanism emerges due to particular learning affordances provided by specific technologies. The second mechanism, in turn, emerges due to dynamics generated by technological change.

The first mechanism links changes in firms' absorptive capacity to the adoption of specific technologies. It helps to understand the broad impacts that the adoption of specific technologies can have on firms' technological trajectories. It is known that some technologies affect the innovation processes of firms by providing specific technological affordances (Yoo et al., 2012). This is an indication that there may be technologies with particular affordances that influence how firms absorb knowledge. Yet, evidence linking learning about a new technology to a change in firms' absorptive capacity was still missing.

The analysis presented in Chapter 4 shows that the adoption of AI causes firms to alter the similarity of their subsequent innovations. This happens because AI allows firms to better address their existing knowledge gaps. Importantly, a significant increase in relatedness also materializes for firms that are technologically distant from AI knowledge. Thus, AI supports firms in addressing their knowledge gaps even when AI is unrelated to firms' prior knowledge. Knowledge proximity, commonalities, and complementarities considered in the indicator of relatedness – which is the proxy used here for measuring absorptive capacity – can explain why relatedness is significantly affected in firms that are technologically closer to AI²⁶. But only one of these three knowledge dimensions can arise when one considers firms that are technologically distant to AI. This dimension alone, namely knowledge complementarity, cannot explain how the introduction of unrelated knowledge leads to an immediate and significant increase in the production of related knowledge in these firms.

The knowledge complementarity dimension, which arises from the need to use distinct technologies together, emerges in the considered case when firms combine AI with another piece of their existing knowledge. If firms are very distant from AI, they have greater difficulties in making this combination once AI is then unrelated to their existing knowledge. In this case, the relatedness of firms' following innovations should remain unchanged or even decrease as firms move closer to the AI cluster²⁷. It was found that this doesn't happen. The causal effects of AI on relatedness are increasingly positive in the first years after adoption for firms distant to AI (i.e., the longer the exposure to AI, the higher the effect on relatedness in these first years), and then decline albeit staying positive.

²⁶ I.e., the three knowledge dimensions – namely knowledge proximity, knowledge commonalities, and knowledge complementarities – are to play out for related firms. However, due to particular characteristics of these three dimensions, the only effects for unrelated firms are to come from AI being used as a knowledge complementarity.

²⁷ This statement can be better understood when one considers how the technological relatedness measure works. The implementation of the indicator, via the relatedness density index, captures the average relatedness between the technologies in which a firm is specialised. The implementation makes it possible to develop specialisations only in a maximum of half of all activities explored by a firm. Considering an increase in the density index (i.e., increase in relatedness), this may happen for two distinct reasons: The firm may have developed an additional specialisation that is on average closer to the technological cluster in which the firm is, or the firm may have switched one of its previously existing specialisations for a specialisation in a technology that is closer to its technological cluster. In any of these cases, firms technologically distant from the AI cluster could not increase their relatedness by specialising in any AI-related technology. Conversely, they could do it if, after AI adoption, they explore their related cluster without developing any AI-related specialisation. This situation could only occur if firms were using AI to innovate rather than to create AI innovations (i.e., rather than combining AI with their existing portfolio, once this would invariably lead to the emergence of a new unrelated specialisation as a technological code linked to AI is repeated over time in the firm subsequent innovations). More details about this are presented in Chapter 4.

This effect is puzzling when one observes only the knowledge dimensions considered in the concept of relatedness. However, AI is also a technology for discovering patterns in large amounts of digital data. This is at the core of the idea of AI being an “Invention of a Method of Invention” (Cockburn et al., 2018). When used as a learning technology, AI can help firms to learn from digital data and to find new commercial opportunities (which arguably leads to the creation of patents). Yet, the ability to recognise what data and opportunities are valuable to explore is dependent on firms’ existing knowledge bases. The data available to a firm is also to be specific to the firm’s current operations, which also determines what kind of innovations can be deployed. This specificity of data and the path dependence of firms in judging the value of opportunities explains why the introduction of AI is linked to an increase in knowledge relatedness even for firms unrelated to this technology. That is, despite AI’s potential to find commercial opportunities in a variety of technological domains (Cockburn et al., 2018), firms still use it for innovating in their typical domains.

AI’s technological potential to be a learning tool creates additional knowledge spillovers of related knowledge. Spillovers generated through learning processes are typically linked to the dimension of knowledge proximity (Breschi et al., 2003). However, in the case of AI, these spillovers are generated from using AI as a technology, rather than due to the learning process generated through its absorption as a new piece of knowledge. This use of AI as a learning technology explains why the results on relatedness are consistent even across sectors that are unrelated to AI. This effect cannot be explained by other technological particularities of AI, such as digital generativity or convergence.

This insight allows conceptualising absorptive capacity as an ability that can be changed by firms’ technological choices. This is in addition to changes linked to learning about AI as a new piece of knowledge. Breschi et al. (2003) highlight that firms extend their innovative activities in knowledge-related domains as a consequence of their learning processes and due to specific features of knowledge and its links. This thesis adds to that view by showing that the adoption of some specific technologies may also impact firms’ ability to explore knowledge-related domains. This finding is particularly interesting because it shows that even unrelated technologies can be used to enhance the creation of related knowledge.

The second mechanism highlights that technological change affects what is related or not to a firm’s knowledge base in a given moment in time. This is to say that technological

development affects the formation of distinct knowledge commonalities between technologies over time²⁸. It affects absorptive capacity by making distinct kinds of knowledge being related to firms' knowledge bases in different moments of time. This increases the uncertainty of the environments in which firms act, making it harder for them to forecast future technological opportunities and risks. For example, AI may reshape firms' competitive advantages by potentially enhancing or turning obsolete the existing knowledge and skills held by them (Paschen et al., 2020). Thus, AI poses uncertainty in the sense that companies do not know to which extent their current accumulation of knowledge may be affected in the future by AI.

One of the few previous works to consider this dynamic aspect of relatedness is Juhász et al. (2021). They find that the co-location of technologies and relatedness not only change over time but also affect each other. The more two technologies overlap within spatial distributions, the greater is the change in their relatedness. As relatedness between two technologies increases, so does the probability of them being co-located in the same geographical space (*ibid*).

These dynamic changes play a large role in firms' technological trajectories, since commonalities between the same pieces of knowledge may be different across locations and time. That is, what is related for a firm in one location could be less related or even unrelated if this firm was situated in a different location. The same goes for distinct periods of time. What creates these differences across locations are the networks available locally, which also affect learning. The differences across time, in turn, are to be created by technological development.

The evidence presented in Juhász et al. (2021) showing that these differences are dynamically influenced by local innovation means that they also grow over time. Taken to the extreme, this would mean that what is related to the existing knowledge of a firm is increasingly dependent on where this firm is located, and when. Hence, firms' location plays a large role

²⁸ An example of this, cited in Chapter 3, refers to cell phones and computers. There were no obvious knowledge commonalities between cell phones and computers in the '90s. Accordingly, distinct firms developed one or another. This is not the case two decades later. Computers and cell phones assimilated similar technologies, which made them more alike: Both now have incorporated the internet, user-friendly operating systems, faster microprocessors, larger data storage, high-resolution colourful displays, etc. This similarity of built-in technologies, in turn, leads to knowledge commonalities that favour their current development in tandem: If a company has the capabilities needed to develop a computer, little learning effort may be needed to develop a smartphone.

in the development of their technological trajectories. Arguably, one important aspect is still missing for a complete picture of how relatedness changes over time. It relates to the fact that some of the technological solutions deployed in innovations are not geographically bounded. This is reflected by the global adoption of specific technologies and by the recombination of these technologies into further innovations. This is the case, for example, of the internet, which started as a technology implemented in a limited network of computers and evolved until it was recombined into modern smartphones.

This geographically unbounded technological development is further explored in Chapter 3. Its existence is explained by the mechanism proposed in Arthur (2009) in which successful innovations that advance technological development are retained and disseminated globally. In this way, although local knowledge becomes increasingly specific over time, successful innovations break through these specificities by entering locally regardless of differences in geographical-based relatedness.

The extension of the dynamic aspect of relatedness through geographically unbounded technological diffusion allows better explaining the local emergence of technological lock-ins. The analysis presented in Chapter 3 shows how successful technological trajectories with AI emerged in countries leading the development of this technology. It is demonstrated that local “fitness” criteria select AI specialisations that coincide with existing general specialisations. This means that AI specialisations deployed locally in areas in which there was no such general knowledge are more likely to disappear. This is true even for areas that are technologically relevant for AI development. This selection criterion leads to the emergence of a pattern in which the majority of countries’ AI specialisations appear in areas in which they also have general specialisations. In this pattern, countries still deploy AI-specific specialisations in fields where they have no previous general specialisations, but these new AI specialisations are short-lived. Therefore, general specialisations act to “preserve” AI-specific specialisations, leading to the emergence of the pattern described above. In fact, this pattern becomes stronger as more knowledge about AI is accumulated locally. This pattern is also shown to be consistent across distinct institutional frameworks.

The mentioned pattern can be explained when one considers that learning is complemented by firms’ interactions (Dosi & Nelson, 1994). The fact that AI patents were being developed locally means that AI knowledge was being created by local actors. Firms were intentionally

learning about AI e.g., through deliberation routines. Yet developing AI knowledge was not enough to produce lasting AI-specific competitive advantages if this knowledge wasn't connected to the local technological trajectory (i.e., "AI" knowledge not connected to the "General" local knowledge). This is to say that locally created AI knowledge was likely to vanish if it didn't interact with the general local knowledge. Hence, the "fitness" criteria seen in the geographical perspective – which selects AI specialisations that coincide with general knowledge – is to be generated largely by the other dimension of learning, i.e., through learning via interactions.

Technological sectors holding a global competitive advantage are to have interactions between local institutions and actors that are more effective – in regards to the emergence of innovations – than sectors in which there are no such advantages (Bergek et al., 2008; Hekkert et al., 2007). These competitive advantages also indicate the existence of a successful local technological trajectory in the sectors linked to them. Moreover, having a global technological advantage in a given field is also indicative of a greater number of actors interacting locally in this field²⁹. All these factors contribute to the stronger preservation of a new piece of knowledge that is considered valuable. Therefore, the existence of well-integrated and larger local networks in successful technological trajectories may be at the root of why new AI specialisations were sustained in some sectors (i.e., in the sectors in which there was a "general capability") whereas in others not. The deployment of AI knowledge in such a context is also to contribute to the creation of successful technological variety with AI. This is to say that the recombination of AI with other knowledge in which the country has a competitive advantage is more likely to be successful (in comparison to the combination of AI with a piece of knowledge in which the country has no competitive advantage at all). Successful innovations are to create economic returns to local firms, which is to influence their decisions in further innovating with the selected technology.

The further local development of AI knowledge affects local relatedness. This is due to the already mentioned mechanism of change due to co-location (Juhász et al., 2021). Thereby, AI's local knowledge increasingly reflects the local context, which affects AI innovations created

²⁹ From the functional approach perspective to Technological Innovation Systems, it could be argued that larger networks are to be more successful in acting to create legitimacy or counteract resistance to change for a technology. This is often identified by analysing the rise and growth of interest groups and their lobby actions (Hekkert et al., 2007).

locally. These local-specific AI innovations can be linked to the idea of speciation (Adner & Levinthal, 2002) and branching (Frenken & Boschma, 2007), which define that new technological varieties emerge as the result of new combinations made with existing knowledge. This leads to the emergence of AI innovations that are increasingly specific to the local conditions, and therefore, increasingly distinct from AI's "unbounded" technological development. This partially explains why geographically unbounded technological development has little to no effect on how this technology is explored locally.

This perspective also allows better explaining why some once-promising and successful local industries eventually decline and fail. The stronger role of local technological trajectories over global (i.e., "unbounded") technological development may lead to lock-in by such industries into local versions of a technology. If a relevant technological development occurs abroad and is not successfully translated to the local context, existing industries may become less competitive and decline. This finding allows linking the idea of technological lock-in to differences created on local relatedness. It also helps explain why some firms and industries leave specific locations despite the fact that they are still well integrated technologically in these locations (Boschma, Balland, et al., 2014; Rigby, 2015).

Therefore, dynamic changes due to technological development affect both firms' ability to learn (i), and what firms learn (ii). The first mechanism is particularly helpful to explain differences in firms' technological performance after learning about specific technologies. This is linked to the fact that some technologies may provide particular learning affordances. The second mechanism is useful for discussing the emergence of local lock-in effects. These effects arguably emerge because local relatedness becomes increasingly disconnected from the global development of a technology. Thus, the adaption of state-of-the-art developments of a technology to the local context requires greater efforts than the creation of further innovations that match the local context.

1.6.3. Policy implications

The findings outlined above provide implications for innovation policies targeting the local development of AI. In particular, they highlight the need for a mix between bottom-up and top-down policies. Top-down policies are designed based on a specific demand (e.g., developing locally a specific technology or industrial sector) and bottom-up policies are planned based on supply (e.g., incentivising the further development of existing local

specialisations). While top-down policies may overlook the role of relatedness by neglecting existing local specialisations, they have shown remarkable success in mission-oriented projects and in developing new industries e.g., through techno-industrial policies. Bottom-up policies, in turn, are currently being introduced in several countries through “smart specialisation” policies (Balland, Boschma, Crespo, & Rigby, 2019; Hidalgo, 2021; Montresor & Quatraro, 2020; Uyarra, Zabala-Iturriagagoitia, Flanagan, & Magro, 2020; Whittle, 2020). These policies incorporate the idea of relatedness, arguing that locations should focus their innovation efforts on developing sectors that are related to existing local knowledge.

The finding of this thesis that combining AI with existing local specialisations is an effective way to preserve AI knowledge is an indication of the potential success of such bottom-up policies aiming at “smart specialisation”. The finding that AI adoption creates benefits even for firms from sectors unrelated to AI knowledge is an additional motivation for the adoption of smart specialisation policies. The main priority should be at identifying technological sectors in which there is already a locally established capability. Amongst such sectors, the ones better connected to others should be preferred, so that more local actors can benefit from positive externalities linked to the development of AI knowledge. Targeting highly connected activities is the optimal strategy to minimize the total time to diversify an economy (Alshamsi et al., 2018).

Another factor that should influence AI innovation policies is the economic potential of sectors. Extant literature indicates that local specialisations in less promising sectors vanish over time as countries improve economically. Hence, top-down policies need to support AI’s development in particularly promising sectors. This should be especially important for less developed economies that expect to grow significantly in the next decades. In this case, considering that local general specialisations would also need to be deployed to preserve the new AI knowledge, longer-term support is needed. Selected promising sectors should be incentivised as a whole, with additional tailoring towards the deployment of AI.

These two perspectives complement Goldfarb and Trefler (2018) and refute Buarque et al. (2020). The latter suggests that local computing-related specialisations are potentially a necessary condition to develop AI and should, therefore, be targeted in innovation policies. Conversely, the findings from this thesis highlight that AI can become part of local technological trajectories in multiple sectors, not limited just to the ones related to

computing. Moreover, combining AI with local specialisations is suggested to be more efficient than developing AI-specific capabilities.

Regarding the former, the authors discuss which kinds of regulation would favour the creation of an AI-friendly environment. Goldfarb and Trefler (2018) point out the need for regulations related to the acquisition, treatment, and sharing of data. This includes domestic privacy policies, data localization rules, access to government data, the establishment of regulation for AI application industries, and the protection of source codes. The authors also argue that economies of scale (due to data availability and AI specialisations) and scope (due to the availability of better software and hardware, and talents in AI) along with knowledge externalities are critical to the national development of AI.

This view complements the findings from this thesis. A comprehensive implementation of regulations that favour the ethical collection, use, and sharing of data should allow a wider set of actors to deploy AI innovations. This would induce the creation of additional knowledge externalities related to AI that can reduce the costs of learning it. This cost reduction should favour particularly actors with knowledge bases more distant to AI to manage the adoption of this technology, which would enhance a broad diffusion of AI knowledge in the local economy.

In sum, three main goals should be considered for the design of effective innovation policies towards local AI development: i) incentivise AI development in sectors in which there are existing specialisations already deployed locally; ii) long-term focus on developing new specialisations in economically promising sectors, with parallel incentives to AI deployment; iii) creation of a friendly environment to the deployment of AI innovations through a comprehensive framework that incentivises the ethical collection, use, and sharing of data.

1.6.4. Limitations and future research

The findings presented in this thesis provide a new point of departure for future research regarding the role of absorptive capacity. For example, further research could analyse whether technologies other than AI could also affect the absorptive capacity of firms. The analysis of unbounded innovation patterns of other technologies is also to be a promising extension. The identification of such patterns for alternative technologies is to be particularly helpful for better understanding commonalities in the emergence of lock-in effects. Such insights would contribute to identifying when an unbounded technological trajectory is divided into local-specific trajectories. Insights about this phenomenon could be helpful to

further understanding the emergence of technological speciation (Adner & Levinthal, 2002) and industrial branching (Frenken & Boschma, 2007), besides potentially contributing to the design of innovation policies by suggesting how such policies could be used to avoid lock-in.

Moreover, the findings presented in this thesis could be extended by further research on complexity. Together with relatedness, complexity is another measure from the toolkit of Economic Complexity that strongly contributed to understanding geographical aspects of innovation in the last decade (see Hidalgo [2021] for an overview). Whereas relatedness is linked to the idea of affinity between pieces of knowledge, complexity connects to the idea of economic value. An implementation of complexity measures should be particularly helpful for identifying whether companies adopting AI are also significantly increasing the value of their subsequent innovations, or for identifying by how much geography matters in the creation of AI's economic value. In any case, such analyses could provide additional insights into the creation of AI-specific innovation policies. They could assist in identifying the sectors in which AI is currently creating more economic value, or which sectors will be particularly promising for the deployment of this technology in the near future.

The research provided in this thesis has several limitations. The first refers to the exclusive use of patent data as a proxy for innovations. This approach neglects a large share of AI innovations. Actors may not patent their AI innovations for several reasons, e.g., to avoid disclosing technical aspects of their inventions and/or prevent copying, due to difficulties in proving the invention's novelty, due to costs related to patenting, etc. Conversely, actors may strategically patent applications that do not represent relevant innovations, e.g., via defensive patenting. In addition, the empirical investigations rely upon patent applications rather than granted patents, which potentially introduces a quality bias.

The second limitation comes from relying exclusively upon a keyword-based strategy to identify AI innovations. There are several alternative approaches in the extant literature³⁰, each with specific limitations. Keyword-based strategies suffer from biases generated by the

³⁰ Typical alternatives include the use of classification codes (e.g., Fujii and Managi (2018); Tseng and Ting (2013) for the identification of AI patents) and expert opinions (e.g., USPTO (2021) for the identification of promising digital technologies, among others). Classification-based strategies suffer from a lack of specificity or clear delimitation between distinct technologies, once classification codes are used with the intent of broadly highlighting technical mechanisms instead of intentionally distinguishing boundaries between them. Opinion-based approaches, in turn, are biased by the subjectivity generated through the particular perception of the experts.

“trending” of distinct terms and keywords over time. Therefore, it is likely that this strategy misses very recent technological developments and domain-specific applications related to AI.

The third limitation concerns the methods adopted. The implementation of the indicator of relatedness adopted in this thesis (in Chapters 3 and 4) uses technological codes for identifying which technologies are related to each patent considered. These codes are applied by experts in patent offices to classify the technical fields related to each patent. Yet, the boundaries between technical fields are hard to identify, and the subjective interpretation of specialists may lead to the same patent being classified with distinct codes by different specialists. The list of classification codes applied also changes slowly over time, meaning that these codes potentially miss recent technological developments. These factors are to create variances that are unaccounted for by the selected relatedness measure.

The fourth limitation comes from the fact that the thesis focuses exclusively on AI. This restricts the reliability of generalising the findings for AI to theoretical contributions linked to the evolutionary view on learning and technological trajectories.

Some of these limitations also suggest possible directions for future research. Particularly, there are promising alternatives that allow overcoming the limitations linked to the exclusive use of patent data and keywords for the identification of AI-related technologies. For the latter, there are alternative methods adopting concepts of technological similarity to identify specific technologies. Text-based methods, in particular, offer a promising alternative to the typical identification approaches based on keywords, codes, or specialists’ opinions. Hain, Jurowetzki, Konda, and Oehler (2020), for example, propose a vector space modelling approach to improve the identification of the technological features of a patent. Based on unstructured text data from patents’ titles and abstracts, the authors create a measure of similarity that is independent of time. This allows better capturing the life-cycle of a technology, besides analysing the novelty, impact, and technological change of patents over time.

Once these identification methods are based on unstructured text data, they can also be applied to other data sources than patents. For the AI case, much of the current AI progress comes from open-source software development. Although algorithms can be patented (Sterne & Bugaisky, 2004), one can see how much more difficult it may be to do so. Proving

novelty and commercial application of a mathematical model is quite challenging³¹. Besides, much of the value of AI applications comes from using existing algorithms over novel datasets. Thus, the novelty may come from the data (or how this data is used) and not the algorithm, making this kind of AI novelty technically not an invention and therefore, not patentable. Due to the openness in the AI community (Bostrom, 2017), some datasets arguably allow identifying this kind of innovation. Open-source public repositories linked to software development, like GitHub, are particularly promising examples. These repositories can show not only relevant actors that create and interact with repositories, but also important aspects of AI's technical innovation. This is because GitHub allow the introduction of text by users, a feature that is applied for describing repositories, their technical implementations, problems, and proposed solutions. Therefore, further research using this kind of data is particularly promising to complement the views presented in this thesis.

³¹ The challenge is to be even more difficult when one considers that this proof must consider the existing mathematical models that could be used as alternatives. The performance of AI algorithms is context-based, and these contexts can be very specific. How to prove that there is value in a model that performs e.g., 0.1% better than an alternative under a very specific circumstance? The value of such an improvement can be easily recognised in some specific situations (e.g., an algorithm that allows this improvement in the detection of early signs of cancer), but not for the large majority of applications.

References

- Aarstad, J., Kvitastein, O. A., & Jakobsen, S.-E. (2016). Related and unrelated variety as regional drivers of enterprise productivity and innovation: A multilevel study. *Research policy*, 45(4), 844-856. doi: <https://doi.org/10.1016/j.respol.2016.01.013>
- Adner, R., & Levinthal, D. A. (2002). The emergence of emerging technologies. *California Management Review*, 45(1), 50-66. doi: <https://doi.org/10.2307/41166153>
- Alcácer, J., Cantwell, J., & Piscitello, L. (2016). Internationalization in the information age: A new era for places, firms, and international business networks? In: Springer.
- Alshamsi, A., Pinheiro, F. L., & Hidalgo, C. A. (2018). Optimal diversification strategies in the networks of related products and of related research areas. *Nature communications*, 9(1), 1-7. doi: <https://doi.org/10.1038/s41467-018-03740-9>
- Antonelli, C., Krafft, J., & Quatraro, F. (2010). Recombinant knowledge and growth: The case of ICTs. *Structural Change and Economic Dynamics*, 21(1), 50-69. doi: <https://doi.org/10.1016/j.strueco.2009.12.001>
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The economic journal*, 99(394), 116-131. doi: <https://doi.org/10.2307/2234208>
- Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. United States of America: Simon and Schuster.
- Balland, P. A. (2016). Relatedness and the geography of innovation. In *Handbook on the geographies of innovation*: Edward Elgar Publishing.
- Balland, P. A., Boschma, R., Crespo, J., & Rigby, D. L. (2019). Smart specialization policy in the European Union: relatedness, knowledge complexity and regional diversification. *Regional studies*, 53(9), 1252-1268. doi: <https://doi.org/10.1080/00343404.2018.1437900>
- Beraja, M., Yang, D. Y., & Yuchtman, N. (2020). *Data-intensive innovation and the state: evidence from AI firms in China*.
- Bergek, A., Jacobsson, S., Carlsson, B., Lindmark, S., & Rickne, A. (2008). Analyzing the functional dynamics of technological innovation systems: A scheme of analysis. *Research policy*, 37(3), 407-429.
- Boschma, R., Balland, P. A., & Kogler, D. F. (2014). Relatedness and technological change in cities: the rise and fall of technological knowledge in US metropolitan areas from 1981 to 2010. *Industrial and corporate change*, 24(1), 223-250. doi: <https://doi.org/10.1093/icc/dtu012>
- Boschma, R., & Frenken, K. (2018). Evolutionary economic geography. *The new Oxford handbook of economic geography*, 213-229.
- Boschma, R., Heimeriks, G., & Balland, P.-A. (2014). Scientific knowledge dynamics and relatedness in biotech cities. *Research policy*, 43(1), 107-114. doi: <http://dx.doi.org/10.1016/j.respol.2013.07.009>
- Bostrom, N. (2017). Strategic implications of openness in AI development. *Global policy*, 8(2), 135-148.
- Breschi, S., Lissoni, F., & Malerba, F. (2003). Knowledge-relatedness in firm technological diversification. *Research policy*, 32(1), 69-87. doi: [https://doi.org/10.1016/S0048-7333\(02\)00004-5](https://doi.org/10.1016/S0048-7333(02)00004-5)
- Bresnahan, T. (2010). General purpose technologies. In *Handbook of the Economics of Innovation* (Vol. 2, pp. 761-791): Elsevier.
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies 'Engines of growth'? *Journal of econometrics*, 65(1), 83-108. doi: [https://doi.org/10.1016/0304-4076\(94\)01598-T](https://doi.org/10.1016/0304-4076(94)01598-T)
- Bresson, C. D. (1987). The evolutionary paradigm and the economics of technological change. *Journal of Economic Issues*, 21(2), 751-762.
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). *What Can Machines Learn, and What Does It Mean for Occupations and the Economy?* Paper presented at the AEA Papers and Proceedings.

- Buarque, B. S., Davies, R. B., Hynes, R. M., & Kogler, D. F. (2020). OK Computer: the creation and integration of AI in Europe. *Cambridge Journal of Regions, Economy and Society*, 13(1), 175-192. doi: <https://doi.org/10.1093/cjres/rsz023>
- Burnham, T. C., Lea, S. E., Bell, A., Gintis, H., Glimcher, P. W., Kurzban, R., . . . Teschl, M. (2016). Evolutionary behavioral economics. In: MIT Press.
- Callaway, B., & Sant'Anna, P. H. (2018). Difference-in-differences with multiple time periods and an application on the minimum wage and employment. *arXiv preprint arXiv:1803.09015*, 1-47.
- Callaway, B., & Sant'Anna, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of econometrics*.
- Cantner, U., & Vannuccini, S. (2012). *A new view of general purpose technologies*.
- Cantner, U., & Vannuccini, S. (2017). Pervasive technologies and industrial linkages: Modeling acquired purposes. *Structural Change and Economic Dynamics*. doi: <https://doi.org/10.1016/j.strueco.2017.11.002>
- Cantwell, J., & Andersen, B. (1996). A statistical analysis of corporate technological leadership historically. *Economics of innovation and new technology*, 4(3), 211-234.
- Castaldi, C., Frenken, K., & Los, B. (2015). Related variety, unrelated variety and technological breakthroughs: an analysis of US state-level patenting. *Regional studies*, 49(5), 767-781. doi: <https://doi.org/10.1080/00343404.2014.940305>
- Cockburn, I. M., Henderson, R., & Stern, S. (2018). *The Impact of Artificial Intelligence on Innovation*.
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative science quarterly*, 128-152. doi: <https://doi.org/10.2307/2393553>
- Colombelli, A., Krafft, J., & Quatraro, F. (2013). Properties of knowledge base and firm survival: Evidence from a sample of French manufacturing firms. *Technological Forecasting and Social Change*, 80(8), 1469-1483.
- Colombelli, A., Krafft, J., & Quatraro, F. (2014). The emergence of new technology-based sectors in European regions: a proximity-based analysis of nanotechnology. *Research policy*, 43(10), 1681-1696. doi: <https://doi.org/10.1016/j.respol.2014.07.008>
- Cowan, R. (1990). Nuclear power reactors: a study in technological lock-in. *The journal of economic history*, 50(3), 541-567. doi: <https://doi.org/10.1017/S0022050700037153>
- Delgado, M., Porter, M. E., & Stern, S. (2014). Clusters, convergence, and economic performance. *Research policy*, 43(10), 1785-1799.
- Dosi, G. (1982). Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change. *Research policy*, 11(3), 147-162. doi: [https://doi.org/10.1016/0048-7333\(82\)90016-6](https://doi.org/10.1016/0048-7333(82)90016-6)
- Dosi, G. (1988). Sources, procedures, and microeconomic effects of innovation. *Journal of economic literature*, 1120-1171.
- Dosi, G., & Grazzi, M. (2006). Technologies as problem-solving procedures and technologies as input-output relations: some perspectives on the theory of production. *Industrial and corporate change*, 15(1), 173-202. doi: <https://doi.org/10.1093/icc/dtj010>
- Dosi, G., & Nelson, R. R. (1994). An introduction to evolutionary theories in economics. *Journal of evolutionary economics*, 4(3), 153-172.
- Ejdemo, T., & Örtqvist, D. (2020). Related variety as a driver of regional innovation and entrepreneurship: A moderated and mediated model with non-linear effects. *Research policy*, 49(7), 104073. doi: <https://doi.org/10.1016/j.respol.2020.104073>
- Ellis, H. C. (1965). The transfer of learning.
- European Commission. (2017). *Harnessing the economic benefits of artificial intelligence*.
- Feldman, M. P., Kogler, D. F., & Rigby, D. L. (2015). rKnowledge: The spatial diffusion and adoption of rDNA methods. *Regional studies*, 49(5), 798-817. doi: <https://doi.org/10.1080/00343404.2014.980799>
- Freeman, C. (1987). *Technology policy and economic performance: Lessons from Japan*: Pinter Publishers Great Britain.
- Freeman, C. (1995). The 'National System of Innovation' in historical perspective. *Cambridge Journal of economics*, 19(1), 5-24.

- Frenken, K., & Boschma, R. A. (2007). A theoretical framework for evolutionary economic geography: industrial dynamics and urban growth as a branching process. *Journal of Economic Geography*, 7(5), 635-649. doi: <https://doi.org/10.1093/jeg/lbm018>
- Frenken, K., Van Oort, F., & Verburg, T. (2007). Related variety, unrelated variety and regional economic growth. *Regional studies*, 41(5), 685-697. doi: <https://doi.org/10.1080/00343400601120296>
- Fujii, H., & Managi, S. (2018). Trends and priority shifts in artificial intelligence technology invention: A global patent analysis. *Economic Analysis and Policy*, 58, 60-69. doi: <https://doi.org/10.1016/j.eap.2017.12.006>
- Furman, J., & Seamans, R. (2018). *AI and the Economy*.
- Goldfarb, A., & Trefler, D. (2018). *AI and International Trade*.
- Granstrand, O. (1998). Towards a theory of the technology-based firm. *Research policy*, 27(5), 465-489.
- Hain, D. S., Jurowetzki, R., Konda, P., & Oehler, L. (2020). From catching up to industrial leadership: towards an integrated market-technology perspective. An application of semantic patent-to-patent similarity in the wind and EV sector. *Industrial and corporate change*, 29(5), 1233-1255.
- Harhoff, D., Heumann, S., Jentzsch, N., & Lorenz, P. (2018). Outline for a German Strategy for Artificial Intelligence.
- Hekkert, M. P., Suurs, R. A., Negro, S. O., Kuhlmann, S., & Smits, R. E. (2007). Functions of innovation systems: A new approach for analysing technological change. *Technological Forecasting and Social Change*, 74(4), 413-432.
- Hidalgo, C. A. (2021). Economic complexity theory and applications. *Nature Reviews Physics*, 3(2), 92-113.
- Hidalgo, C. A., & Hausmann, R. (2009). The building blocks of economic complexity. *Proceedings of the national academy of sciences*, 106(26), 10570-10575. doi: <https://doi.org/10.1073/pnas.0900943106>
- Hidalgo, C. A., Klinger, B., Barabási, A.-L., & Hausmann, R. (2007). The product space conditions the development of nations. *science*, 317(5837), 482-487. doi: <https://doi.org/10.1126/science.1144581>
- Hilgard, E. R., & Bower, G. H. (1966). Theories of learning.
- Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, 313(5786), 504-507. doi: <https://doi.org/10.1126/science.1127647>
- Ivarsson, I., Alvstam, C., & Vahlne, J.-E. (2015). Global technology development by collocating R&D and manufacturing: the case of Swedish manufacturing MNEs. *Industrial and corporate change*, dtw018.
- Jin, G. Z. (2018). *Artificial Intelligence and Consumer Privacy*.
- Juhász, S., Broekel, T., & Boschma, R. (2021). Explaining the dynamics of relatedness: The role of collocation and complexity. *Papers in Regional Science*, 100(1), 3-21. doi: <https://doi.org/10.1111/pirs.12567>
- Kim, J., Lee, C.-Y., & Cho, Y. (2016). Technological diversification, core-technology competence, and firm growth. *Research policy*, 45(1), 113-124.
- Klinger, J., Mateos-Garcia, J., & Stathoulopoulos, K. (2018). Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology. doi: <https://arxiv.org/pdf/1808.06355.pdf>
- Kogler, D. F., Rigby, D. L., & Tucker, I. (2013). Mapping knowledge space and technological relatedness in US cities. *European Planning Studies*, 21(9), 1374-1391.
- Kuusk, K., & Martynovich, M. (2021). Dynamic Nature of Relatedness, or What Kind of Related Variety for Long-Term Regional Growth. *Tijdschrift voor economische en sociale geografie*, 112(1), 81-96. doi: <https://doi.org/10.1111/tesg.12427>
- Le Bas, C., & Sierra, C. (2002). 'Location versus home country advantages' in R&D activities: some further results on multinationals' locational strategies. *Research policy*, 31(4), 589-609.

- Leten, B., Belderbos, R., & Van Looy, B. (2007). Technological diversification, coherence, and performance of firms. *Journal of Product Innovation Management*, 24(6), 567-579.
- Li, X., & Jiang, H. (2017). Artificial Intelligence Technology and Engineering Applications. *Applied Computational Electromagnetics Society Journal*, 32(5).
- Lipsey, R. G., Carlaw, K. I., & Bekar, C. T. (2005). *Economic transformations: general purpose technologies and long-term economic growth*: OUP Oxford.
- Lundvall, B.-Å. (1992). *National systems of innovation: Toward a theory of innovation and interactive learning*: Pinter Publishers.
- Montresor, S., & Quatraro, F. (2020). Green technologies and Smart Specialisation Strategies: a European patent-based analysis of the intertwining of technological relatedness and key enabling technologies. *Regional studies*, 54(10), 1354-1365. doi: <https://doi.org/10.1080/00343404.2019.1648784>
- Neffke, F., Henning, M., & Boschma, R. (2011). How do regions diversify over time? Industry relatedness and the development of new growth paths in regions. *Economic geography*, 87(3), 237-265. doi: <https://doi.org/10.1111/j.1944-8287.2011.01121.x>
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Cambridge, Massachusetts, and London, England: The Belknap Press of Harvard University Press.
- Nesta, L. (2008). Knowledge and productivity in the world's largest manufacturing corporations. *Journal of Economic Behavior & Organization*, 67(3-4), 886-902.
- Nilsson, N. J. (2009). *The quest for artificial intelligence*: Cambridge University Press.
- Paschen, U., Pitt, C., & Kietzmann, J. (2020). Artificial intelligence: Building blocks and an innovation typology. *Business Horizons*, 63(2), 147-155.
- Petralia, S., Balland, P. A., & Morrison, A. (2017). Climbing the ladder of technological development. *Research policy*, 46(5), 956-969. doi: <https://doi.org/10.1016/j.respol.2017.03.012>
- Rigby, D. L. (2015). Technological relatedness and knowledge space: entry and exit of US cities from patent classes. *Regional studies*, 49(11), 1922-1937. doi: <https://doi.org/10.1080/00343404.2013.854878>
- Righi, R., Samoilu, S., Cobo, M. L., Baillet, M. V.-P., Cardona, M., & De Prato, G. (2020). The AI technological complex System: Worldwide landscape, thematic subdomains and technological collaborations. *Telecommunications Policy*, 101943. doi: <https://doi.org/10.1016/j.telpol.2020.101943>
- Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*: Malaysia; Pearson Education Limited.
- Santoalha, A., Consoli, D., & Castellacci, F. (2021). Digital skills, relatedness and green diversification: A study of European regions. *Research policy*, 50(9), 104340. doi: <https://doi.org/10.1016/j.respol.2021.104340>
- Schumpeter, J. A. (1939). *Business cycles: a theoretical, historical, and statistical analysis of the capitalist process*: McGraw-Hill.
- Schumpeter, J. A. (1942). *Capitalism, socialism and democracy*: Harper & Brothers.
- Simon, H. A. (1985). What we know about the creative process. *Frontiers in creative and innovative management*, 4, 3-22.
- Solheim, M. C., Boschma, R., & Herstad, S. (2018). *Related variety, unrelated variety and the novelty content of firm innovation in urban and non-urban locations*. Retrieved from Papers in Evolutionary Economic Geography (PEEG) 1836:
- Solheim, M. C., Boschma, R., & Herstad, S. J. (2020). Collected worker experiences and the novelty content of innovation. *Research policy*, 49(1), 103856.
- Sterne, R. G., & Bugaisky, L. B. (2004). The Expansion of Statutory Subject Matter Under the 1952 Patent Act. *Akron L. Rev.*, 37, 217.
- Taddy, M. (2018). *The Technological Elements of Artificial Intelligence*.
- Tanner, A. N. (2016). The emergence of new technology-based industries: the case of fuel cells and its technological relatedness to regional knowledge bases. *Journal of Economic Geography*, 16(3), 611-635. doi: <https://doi.org/10.1093/jeg/lbv011>

- Teece, D., & Pisano, G. (1994). The Dynamic Capabilities of Firms: an Introduction. *Industrial and corporate change*, 3(3), 537-556. Retrieved from <https://doi.org/10.1093/icc/3.3.537-a>. doi:10.1093/icc/3.3.537-a
- Teece, D. J. (2018). Profiting from innovation in the digital economy: Enabling technologies, standards, and licensing models in the wireless world. *Research policy*, 47(8), 1367-1387. doi: <https://doi.org/10.1016/j.respol.2017.01.015>
- Teece, D. J., Rumelt, R., Dosi, G., & Winter, S. (1994). Understanding corporate coherence: Theory and evidence. *Journal of Economic Behavior & Organization*, 23(1), 1-30.
- Trajtenberg, M. (2018). *AI as the next GPT: a Political-Economy Perspective*.
- Tseng, C.-Y., & Ting, P.-H. (2013). Patent analysis for technology development of artificial intelligence: A country-level comparative study. *Innovation*, 15(4), 463-475.
- Tucker, C. (2018). Privacy, Algorithms, and Artificial Intelligence. In *The Economics of Artificial Intelligence: An Agenda*: University of Chicago Press.
- Tushman, M. L., & Anderson, P. (1986). Technological discontinuities and organizational environments. *Administrative science quarterly*, 439-465.
- USPTO. (2021). CPC Section Y. Retrieved from <https://www.uspto.gov/web/patents/classification/cpc/html/cpc-Y.html>
- Uyarra, E., Zabala-Iturriagoitia, J. M., Flanagan, K., & Magro, E. (2020). Public procurement, innovation and industrial policy: Rationales, roles, capabilities and implementation. *Research policy*, 49(1), 103844. doi: <https://doi.org/10.1016/j.respol.2019.103844>
- Van Den Berge, M., & Weterings, A. (2014). Relatedness in eco-technological development in European regions. *Papers in Evolutionary Economic Geography*, 14(13), 1-30.
- Van den Bergh, J. C. (2008). Optimal diversity: increasing returns versus recombinant innovation. *Journal of Economic Behavior & Organization*, 68(3-4), 565-580. doi: <https://doi.org/10.1016/j.jebo.2008.09.003>
- van Eck, N., & Waltman, L. (2009). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics*, 84(2), 523-538.
- Weber, C., & Weber, B. (2010). Social capital and knowledge relatedness as promoters of organizational performance: An explorative study of corporate venture capital activity. *International Studies of Management & Organization*, 40(3), 23-49.
- Whittle, A. (2020). Operationalizing the knowledge space: theory, methods and insights for Smart Specialisation. *Regional Studies, Regional Science*, 7(1), 27-34. doi: <https://doi.org/10.1080/21681376.2019.1703795>
- Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., . . . Bourne, P. E. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*, 3(1), 1-9.
- Winter, S. G. (1987). Natural selection and evolution. In *Allocation, Information and Markets* (pp. 214-222): Springer.
- WIPO. (2019). *WIPO Technology Trends 2019: Artificial Intelligence*. Retrieved from https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf
- Yoo, Y., Boland Jr, R. J., Lyytinen, K., & Majchrzak, A. (2012). Organizing for innovation in the digitized world. *Organization science*, 23(5), 1398-1408. doi: <https://doi.org/10.1287/orsc.1120.0771>
- Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., . . . Dinlersoz, E. (2021). *Advanced Technologies Adoption and Use by US Firms: Evidence from the Annual Business Survey*.

Authors: Matheus Eduardo Leusin, Jutta Günther, Björn Jindra, and Martin G.

Möhrle

Abstract: This chapter identifies countries at the forefront of Artificial Intelligence (AI) development and proposes two novel patent-based indicators to differentiate structural differences in the patterns of intellectual property (IP) protection observed for AI across countries. In particular, we consider (i) the extent to which countries specialise in AI and are relevant markets for corresponding IP protection ('National Breeding Ground'); and (ii) the extent to which countries attract AI from abroad for IP protection and extend the protection of their AI-related IP to foreign markets ('International Breeding Ground'). Our investigation confirms prior findings regarding substantial changes in the technological leadership in AI, besides drastic changes in the relevance of AI techniques over time. Particularly, we find that National and International Breeding Grounds overlap only partially. China and the United States can be characterised as dominant National Breeding Grounds. Australia and selected European countries, but primarily the United States, are major International Breeding Grounds. We conclude that China promotes AI development with a major focus on IP protection in its domestic market, whereas the United States sustains its AI progress in the international context as well. This might indicate a considerable bifurcation in the structural patterns of IP protection in global AI development.

Keywords: Artificial Intelligence; Digitalisation; Patent analysis; International comparison; Specialisation;

JEL Classification: D02; O14; O34; O57;

Publication

This chapter has been published as the following original research article:

Leusin, M.E., Günther, J., Jindra, B., Moehrle, M.G., 2020. Patenting patterns in Artificial Intelligence: Identifying national and international breeding grounds. *World Patent Information* 62, 101988. <https://doi.org/10.1016/j.wpi.2020.101988>

Formal changes, compared to the published version, have been made.

2.1. Introduction

The transformation from an analogue into a digitalised world economy has been underway for some decades, manifested primarily by the diffusion of information and communication technologies into the realm of business and society. In the current wave of digitalisation, the development and use of Artificial Intelligence (AI) represents a qualitatively new development. Despite its more than 60 years of existence, AI's potential and market relevance has been recognised worldwide only in the last decade, mainly due to the development of high-performance parallel computing chips and large datasets that have extended this technology's applicability (Klinger, Mateos-Garcia, & Stathoulopoulos, 2018; Li & Jiang, 2017). Nowadays, AI can be embedded in any technology (software, algorithm, a set of processes, a robot, etc.) that is able to function appropriately when endowed with the foresight of its environment (Nilsson, 2009).

Given its scope and potential impact, AI is currently considered a strategic technology for many countries. Accordingly, a "global AI race" was deployed in pursuit of its development, with countries increasingly investing in national AI strategies intended to gain advantages over global markets and industries (Klinger et al., 2018). However, it is known from Innovation Systems (IS) theory that, even if countries focus their efforts on the same direction, distinct national characteristics will affect technological development (Bassis & Armellini, 2018). Beyond national boundaries, it is also known that a technology, or the knowledge it embodies, is rarely embedded in just the institutional infrastructure of a single nation or region, since the relevant knowledge base for most technologies originates from distinct geographical areas (Hekkert, Suurs, Negro, Kuhlmann, & Smits, 2007; P Sharma & Tripathi, 2017). At the same time, the intrinsic characteristics of each type of technology affect both the development of their particular body of knowledge and the diffusion of its applications.

To understand the particularities of the "global AI race", we want first to identify which countries have been leading this development. Moreover, we agree with Hekkert et al. (2007) in the sense that, once innovation systems co-evolve with the process of technological change, it is important not only to identify a final and static picture of the technological development but the dynamics of its evolution over time. Thus we analyse, over time: i) which countries have been specialising in and increasingly adopting AI technologies; ii) which countries seek to enhance the attractiveness of their markets to foreign inventors and companies by

establishing strong local protection of AI-related Intellectual Property (IP) and/or seek to protect their own AI-related IP under the applicable laws in foreign markets; and iii) which corresponding trends we can observe for specific AI techniques across countries.

For this purpose, we use worldwide patent data and adopt a very broad concept of what an AI patent is: Besides searching for all patents with some direct reference to AI, we also search for AI techniques, which are advanced statistical and mathematical models used to implement AI functions such as vision, language, and decision-making. We use established indices to analyse the technological specialisation of countries and, as a novel contribution, we introduce two indexes: The National Breeding Ground (NBS) and the International Breeding Ground (IBG). The first orders countries on the basis of market favourability, that is, it assesses the extent to which the market is perceived as being exploitable by inventors and companies, no matter where they come from. The second orders countries on their ability to attract applications for domestic AI patents from foreign inventors/companies, as well as the degree to which AI patents registered in the country provide legal protection abroad.

We find that the number of AI-specialised countries has grown in the last decades, with a more pronounced growth seen mainly in the 90's. Despite this, the exploitation of AI-related techniques became much more intense after the turn of the century, pushed mainly by the use of biological models (neural networks, supervised and unsupervised machine learning, deep learning, etc.). In terms of international leadership, we can confirm a decline in terms of relevance of Japan and some European countries, conversely to a greater increase in the relevance of the United States and the emergence of China. These two leading countries manage AI exploration very distinctly, as described and explained in the next sections.

2.2. Literature review: Breeding Ground indicators related literature

Among the many indicators used in the literature for comparing country performances in e.g., FDI, number of patents produced, or amount of R&D spent at the national level, the Revealed Comparative Advantage (RCA) index stands out (Mansourzadeh, Shahmoradi, Dehdarirad, & Janavi, 2019). Introduced in Balassa (1965), the RCA indicator was proposed as a measure of the relative specialisation of a country in the production of a specific product. The simple idea behind the indicator is that when a country has a production level higher than the global

average, this country holds a comparative advantage in producing it. Equation 1 presents the formula for calculating the RCA of a given country “b” for a specific product “a”.

$$RCA_{Product\ a, Country\ b} = \frac{\frac{\text{Exports of product } a \text{ in country } b}{\text{Total of exports of country } b}}{\frac{\text{Total exports of product } a \text{ in the world}}{\text{Total exports in the world}}} \quad (1)$$

The RCA has been adapted to analyse, among others the comparative advantage of a given country for a specific technology/technological field, through the so-called Revealed Technological Advantage (RTA) index (Soete, 1987). Like the RCA, the RTA also has extensive application in the literature, with a recent example in Weresa (2019), in which the author analyses the comparative advantages of the European Union in digital technologies using patents as inputs for this index.

However, despite its popularity, the RCA (and its related adaptations, as the RTA) has some limitations, including the challenge of interpreting values above 1, highly asymmetric distribution, and the consequent impossibility of directly interpreting the resulting RCA values (Menzel & Maicher, 2017). The literature points to the prominence given by the RTA to less patent-active countries, and to the problem of statistical bias when the overall number of patents in a specific technology/technological field is too small (Soete, 1987). In this way, countries with a small national output (e.g., technological output as patents) tend to be highly favoured in cross-national comparisons. Namely, an above-one RTA value signals a country’s specialisation, even if the country’s absolute output is low. However absolute outputs are highly relevant for market comparisons, especially when discussing the adoption of a new technology: To attract relevant inventors and companies involved with the technology, a country must offer not only some kind of specialisation, but also a relevant market.

To assist technology holders with evaluating competing marketplaces and to enable countries to better understand where they stand, we propose “Breeding Ground” indexes. The basic premise behind them, supported by qualitative discussions in innovation system theory, is that particular characteristics of each innovation system affect innovators with regard to the markets they choose to exploit their technologies. Accordingly, we aim to identify countries that tailor their National Innovation System towards the adoption of a given technology by providing specific local conditions. In our case, the first favourable conditions considered are the availability of local specialisation advantages together with an attractive domestic market to exploit the technology, which we measure jointly with the proposed NBG index. We expect

that countries with a higher NBG number can explore this characteristic to develop the so-called national champions, as highlighted for example in Kroll (2011) in an analysis of China, which points that, powered by a strong domestic market, such companies might become competitive in foreign markets. In addition, we propose the IBG index, which we use to identify countries that are attractive to inventors/companies from abroad while also having their own national AI-patents recurrently extended to markets abroad (i.e., by the extension of their patents' legal protection to foreign countries). We argue that this index reflects not only the relevance (size) of markets for AI exploitation originating from abroad, but also efficient institutions reinforcing IP protection and cooperative behaviour (Gürerk, Irlenbusch, & Rockenbach, 2006). Furthermore, the indicator reflects the national production of AI technology considered relevant for protection on other foreign markets. A high country score on the NBG index but not on the IBG index might signal that strong technological development and exploitation, as well as corresponding patterns of IP protection, take place mainly in a closed domestic arena rather than in an international context.

2.3. Data collection

In 2019, WIPO (2019b) argued that despite the availability of information in patent documents, it can be difficult to identify exactly which patent families relate to AI because of the lack of a standardised definition; even non-standard definitions of AI change over time. The literature proposes a variety of strategies for identifying AI-related documents (e.g., patents or publications), including the use of predefined classes based on patent classification schemes (Fujii & Managi, 2018; Tseng & Ting, 2013), the use of specific keywords (Annoni et al., 2018; Aristodemou & Tietze, 2018; Huang, Miao, Zhang, Yu, & Wang, 2017; Niu, Tang, Xu, Zhou, & Song, 2016), or even both (Cockburn, Henderson, & Stern, 2018; Keisner, Raffo, & Wunsch-Vincent, 2015; WIPO, 2019b). Both strategies have pros and cons: The choice of keywords and IPC codes is inherently subjective – the first choice depending on which keywords are considered relevant and that of patent officers depending on the interpretation from IP specialists on the content provided in the application as they seek to classify them (see for example Meguro & Osabe (2019) for a recent related discussion for the use of IPC codes). Thus, our intention here is not to extend this discussion in the sense of defining AI precisely (even the syntagm “Artificial Intelligence” has only recently reappeared in industry, after some hypes and disillusionments, as noted in Lupu (2018)). Rather, with the aim of developing an

overview of AI patenting activities worldwide in mind, we seek to create a dataset that is strongly related to the core of AI development through the years. For this, we collected data on all patents with some direct reference to Artificial Intelligence or that are based on typical AI techniques.

To identify relevant AI techniques, we used the framework proposed in WIPO (2019b), which includes the analysis of AI experts on both the applied and the research domains on this topic. This framework, based on a computing classification scheme developed over the past 50 years by the Association for Computing Machinery, differentiates three main categories related to AI: i) AI techniques as advanced forms of statistical and mathematical models; ii) AI functional applications; and iii) AI application fields (i.e., areas or disciplines in which AI techniques or functions may find application). In our investigation, we analysed AI techniques as the technological core that enables diffusion into related functional applications and broader application fields. Thus, we adopted the 21 keywords related to AI techniques proposed in the WIPO report (see WIPO (2019b), p. 24), complementing these with their synonyms (to avoid losing relevant patents due to different wording) collected from Wikipedia. In addition to being the largest knowledge repository of the Web, Wikipedia offers multi-faceted and cross-linked classifications and concepts (Gabrilovich & Markovitch, 2006). Pawan Sharma, Tripathi, & Tripathi (2015) further emphasises Wikipedia as the best available option for gaining a comprehensive understanding of the many technical terms used in patents. The complete list of keywords and synonyms considered is presented in Table 2, together with the techniques' definitions given in WIPO (2019b).

AI Technique	Definition	Additional Wikipedia Synonyms
Bio-inspired approaches	A family of AI approaches inspired by biological systems, rather than a precise technique.	Bio-inspired computing, biologically inspired computing
Classification and regression trees	Predictive models that use tree-like representations of facts and their possible consequences.	Decision tree learning
Deep learning	A machine learning approach that tries to understand the world in terms of a hierarchy of concepts. Most deep learning models are implemented by increasing the number of layers in a neural network.	Deep structured learning, hierarchical learning
Description logistics	A form of programming used in Logic programming.	Keyword not found
Expert systems	A computer system that solves complex problems within a specialised domain, based on an expertise expressed manually by human experts in the form of a set of rules.	No additional synonym
Fuzzy logic	A decision-making approach that is not based on the usual “true or false” assessment, but rather on “degrees of truth” (where the “truth” value ranges between completely true and completely false).	No additional synonym
Instance-based learning	A family of machine learning algorithms that compare a new problem with cases seen in training and can adapt the model to previously unseen data.	Memory-based learning
Latent representation	The mathematical representation of variables that are inferred rather than directly observed. Latent representation is applied in natural language processing, for example, where it is usually inferred from the statistical distribution of words.	No additional synonym
Logic programming	Uses facts and rules to make decisions, without specifying additional intermediary steps, in order to achieve a particular goal.	No additional synonym
Logical and relational learning	It is a form of learning related to Machine Learning.	No additional synonym
Machine learning	An AI process that uses algorithms and statistical models to allow computers to make decisions without having to explicitly program it to perform the task.	No additional synonym

Multi-task learning	A machine learning approach where a single model is used to solve multiple learning tasks at the same time, exploiting commonalities and differences between the various tasks.	Multitask Learning
Neural networks	A learning process inspired by the neural structures of the brain, being the network generally organised in successive layers of functions, with each layer using the output of the previous one as an input.	No additional synonym
Ontology engineering	A set of tasks related to the methodologies for building ontologies, namely the way concepts and their relationship in a particular domain are formally represented.	No additional synonym
Probabilistic graphical models	A framework for representing complex domains using distribution of probabilities; the models use a graph-based representation for defining the statistical dependence or independence relationships between data.	Graphical model, structured probabilistic model
Probabilistic reasoning	An approach that combines deductive logic and probability theory to model logical relations under uncertainty in data.	Probability logic, probabilistic logic
Reinforced learning	An area of machine learning that uses a system of reward and punishment as it learns how to attain a complex objective.	Reinforce-ment learning
Rule learning	Machine learning methods which identify and generalise automatically a set of rules (which are usually simple conditional tests) to be used for prediction or classification of new, unseen data.	Rule induction
Supervised learning	The expected grouping of the information in certain categories (output) is provided to the computer through examples of data (input) that have been manually categorised correctly and comprise a training dataset. Based on these examples of input-output, the AI system can organise new, unseen data into predefined categories.	No additional synonym
Unsupervised learning	A type of machine learning algorithm that finds and analyses hidden patterns or commonalities in data that has not been labelled or classified.	No additional synonym
Support vector machines	A supervised learning algorithm that analyses labelled/grouped data, identifies the data points that are most challenging to group and, based on that, identifies how to separate the different groups and classify unseen data points.	Support vector networks

Table 2: AI techniques considered and definitions. Source: WIPO (2019b) (pp. 148–150).

Building on our previous definition, while some AI patents are related to the use of AI techniques, others reference AI directly. Hence, we also included the search term “Artificial Intelligence” and its Wikipedia synonym “Machine intelligence”.

The patent search was conducted in the autumn 2017 version of PATSTAT (PATSTATb). We identified 40,481 patent applications³² (each of which has a unique application ID³³, “appln_id”) whose title or abstract contains the above-outlined keywords. The full query used for retrieving these Application IDs is presented in Appendix A. Once the application IDs were identified, the remaining relevant data was retrieved from PATSTAT using these applications IDs as input.

As a test of robustness, we identified the number of patents related to each combination of the terms adopted within the collected sample. These results, as well as the exact keywords used for the collection and for this robustness analysis, are presented in Table 3.

³² An application is a request for patent protection of an invention. Applications are registered on PATSTAT whether or not they have been granted.

³³ The appln_id is a numerical technical identifier used in all PATSTAT databases to uniquely identify a patent application, allowing the identification of the same application across all editions of all PATSTAT databases (EPO, 2018).

AI Technique Keyword	Additional Wikipedia Synonym Keyword	No. of Patents
%neural network%	No additional synonym	19,784
%machine learn%	No additional synonym	5,228
%artificial intelligence%	%machine intelligen%	4,197
%expert system%	No additional synonym	3,838
%support vector machin%	%support vector network%	3,442
%fuzzy logic%	No additional synonym	2,883
%graphical model%	%structured probabilistic model%	806
%pervised learn%	No additional synonym	667
%deep learn%	%deep structured learn% OR %hierarchical learn%	663
%classification tree% OR %regression tree%	%decision tree learn%	415
%reinforced learn%	%reinforcement learn%	375
%logic programming%	No additional synonym	152
%rule learn%	%rule induction%	111
%probabilistic reason%	%probability logic% OR %probabilistic logic%	60
%task learn%	No additional synonym	56
%logical learn% OR %relational learn%	No additional synonym	30
%latent represent%	No additional synonym	10
%bio-inspired approach%	%bio-inspired comput% OR %biologically inspired comput%	7
%instance-based learn%	%memory-based learn%	7
%ontology engineer%	No additional synonym	5
%description logistic%	Keyword not found	0

Table 3: Keywords used for the patent search. Note: The double characters “%” at the beginning and end of each keyword are used to include variations before or after this character so that any patent that coincides with the term between these two characters is collected. The search in PATSTAT is not case-sensitive.

The inverse check resulted in 42,736 patents, which implies that at least 2,255 patents (or approximately 5.6% of the sample) have a combination of at least two of the selected terms. One might expect that keywords like "graphical model" would be too generic, but it is seen that this broad keyword generated only 860 (1.9%) of the patents. At the same time, 92.1% of the patents found concentrate around 6 of the 21 terms used (namely Neural networks, Machine learning, Artificial Intelligence or Machine intelligence, Support vector machines or Support vector networks, Expert systems, and Fuzzy logic).

For the final dataset, we excluded the patents of utility models. This is because, besides having a shorter protection terms and grant lags, utility models have been used by IP professionals as auxiliary tools in specific national contexts to overcome shortcomings of the patent system, as discussed in Radauer, Rosemberg, Cassagneau-Francis, Goddar, & Haarmann (2015). Furthermore, we also consider the distinction between priority and non-priority filings. In short, a priority filing (or priority patent) is the first patent application filed to protect an invention. It represents the total number of patent families, regardless of their spatial protection scope. After a priority filing, if the same patent is registered in other patent offices, the following registrations are called non-priorities, constituting a patent family linked through the priority filing. In PATSTAT, an application ID (`appln_id`) enables the retrieval of information about the first filling ID (`earliest_filing_id`) associated with this patent. If the application ID and the first filling ID are equal, this application ID is considered a priority filing; if not, this particular application ID is considered a non-priority filing.

We took into account the distinction between registrations made under the Patent Cooperation Treaty (PCT) and standard registrations filed only with national patent offices. The PCT, an international patent law treaty, provides a unified procedure for filing patent applications to protect an invention in each of its 152 signatory states (Gaétan De Rassenfosse, Dernis, & Boedt, 2014; WIPO, 2019a). The PCT registry doesn't grant nor examine patent applications. Instead, it allows the applicants of a patent to delay the expensive step of filing other foreign patent applications. Thus, the benefit of such PCT route is to allow applicants to seek patent protection simultaneously in a large number of countries. The possible following registers of the patent will have the filing date of the first application, and they cannot be invalidated through any acts that occurred during the interval allowed by the PCT (Lapenne, 2010). Furthermore, the PCT registry also extends the period to which a subsequent filing can be registered based on the priority filing from the regular 12 months to 31 months; in this

way, the applicant has more time to assess the potential of its invention and proceed or not with the patent application (Gaètan De Rassenfosse, Dernis, Guellec, Picci, & de la Potterie, 2013). Gaètan De Rassenfosse et al. (2013) and Thoma (2014) point out that although some bias might exist, empirical evidence suggests that the PCT route is associated with higher-value patents (or at least, inventions with high market potential, as also argued in Van Zeebroeck & Van Pottelsberghe de la Potterie (2011)). Accordingly, we assume that PCT registrations are chosen by assignees for the patents they consider more valuable on international markets, thus contributing more than other registrations to the international commercialisation of AI.

2.4. Comparison of strategies for identifying AI patents

As already mentioned, there is a wide variety of possibilities for searching patents on a given topic, which includes the use of keywords and the use of patent classification schemes. Given our use of a keyword-based strategy, we compared our results with those from two other AI-related papers, namely Fujii & Managi (2018) and Tseng & Ting (2013), that are instead based on the use of codes from the International Patent Classification (IPC)³⁴. The mentioned authors use a search based on the IPC codes presented in Appendix B.

In short, every code used by Fujii & Managi (2018) pertains to the subclass “Computer systems based on specific computational models” (G06N), while Tseng & Ting (2013) also includes codes related to the subclasses “Optical Computing Devices” (G06E) “Analogue Computers” (G06G), “Hybrid Computing Arrangements” (G06J), “Electric Digital Data Processing” (G06F), and one related to the subgroup of electric adaptive control systems (G05B 13/02). Excluding utility models for the sake of comparability, the search from Fujii & Managi (2018) and Tseng & Ting (2013) results in 23,599 and 146,049 priority filings, respectively. To compare the efficiency of these search strategies with ours, we manually selected the first³⁵ 100 patents from our dataset that i) had a title and abstract; and ii) were written in English. We then selected the first 100 records from the two competing datasets that, in addition to our criteria, iii) were not in our dataset. This yielded a total of 300 patents, which we classified individually based on title and abstract.

³⁴ Available in: <https://www.wipo.int/classifications/ipc/ipcpub>.

³⁵ First here means that, for reproductive purposes, each sample of patents was selected in ascending order in relation to `appln_id`.

Although our intention here is not to extend the discussion in the sense of defining what AI is, we do need to define what we consider an AI patent to make this classification possible. Thus, considering our broad view on AI, we classified as AI-related all patents that met at least one of the following criteria: i) can be used to generate some kind of prediction or classification useful to make some decision, or to interpret or summarise some type of knowledge; ii) enables the automation or optimisation of some task or parameter used in the patent to perform or improve some kind of selection; iii) enables the generation of useful and analysable data, or the autonomous correction of existing data; iv) is related to some kind of training, learning or dynamic adaptation based on data; v) enables the recognition or evaluation of objects or patterns of interest.

The results of this comparison are available in our public GitHub repository³⁶, which also includes all data used in this chapter and the associated R codes for reproducing it. In total, 13 patents in the sample (6 from our Query, 6 from Query 2 and 1 from Query 3) did not have enough information for being classified with certainty, so we categorised them as “Unclear”. Figure 2 summarises this comparison, highlighting the number of results that overlap in Queries 2 and 3 in relation to our query, and the accuracy of each query in relation to the total of 287 records analysed in what concerns being an AI patent or not.

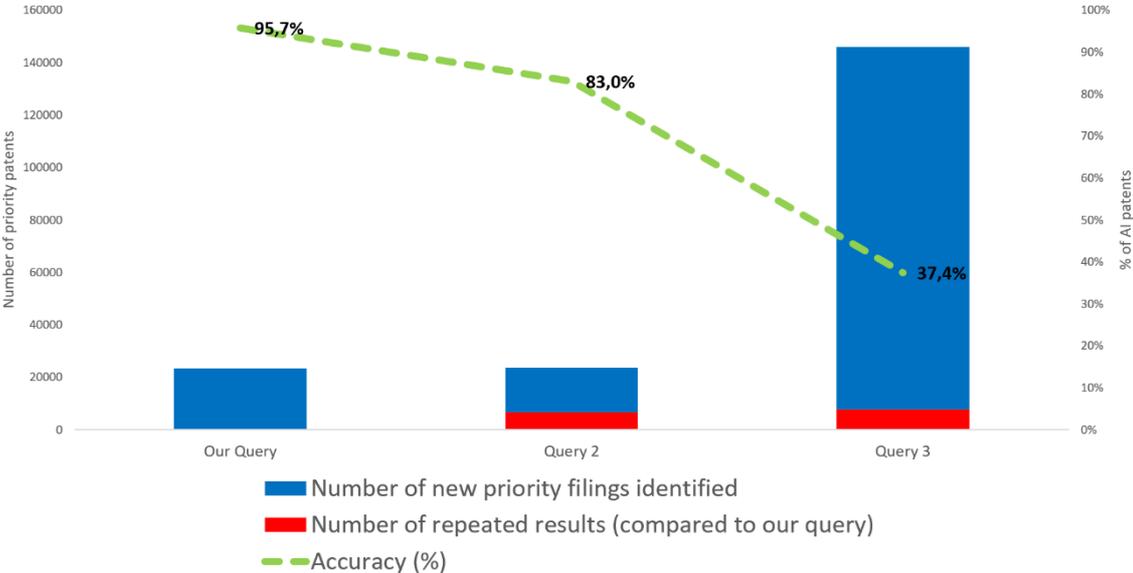


Figure 2: Comparison of results between the selected queries.

It can be seen from Figure 2 that there is a significant overlap of results between the two queries considered in relation to ours: 28.2% and 33.1% of our results are also found using

³⁶ <https://github.com/matheusleusin/Patenting-Patterns-in-Artificial-Intelligence>.

Queries 2 and 3, respectively. However, considering the mentioned definition of AI patents, those two queries are significantly less capable of detecting such patents. Query 3, in particular, expands greatly the number of results by including IPC codes less related to AI, which drastically reduces the quality of results generated by analysis of this dataset (AI patents comprised only 37.4%). Our query also performs significantly better than Query 2 (95.7% accuracy compared to 83.0%), thus generating a similarly sized data set of higher quality. One possible explanation for these outcomes, discussed in Meguro & Osabe (2019), is that IP offices tend to classify patents by their own standards, potentially losing important features of the invention in the translation into standardised IPC language. The authors even highlight that computer technology is one of the technical fields that demands especial attention by researchers who consider applications across IP offices, because patent applications are often relevant to a variety of classifications and quite-similar applications are coded differently in distinct offices. Thus, the use of a keyword-based search is shown to be a good choice for circumventing the challenges faced by other approaches to gaining a global perspective on far-reaching and still-emerging technologies like AI. In particular, the use of keywords related to AI techniques increases the probability that the patents found are associated with purposes typical of AI, which further improves the quality of the dataset.

2.5. Method and results

The patent dataset was separated into three periods of twelve years, thus excluding patent applications prior to 1979 (17 patents) and – since the version used of PATSTAT does not show the complete patents of the years 2016 and 2017 – applications after 2015 (4,662 patents), resulting in 34,679 priority and non-priority filings to be used as input for calculations related to the specialisation values, and to the attributes of National and International AI Breeding Grounds, discussed in the following sections. The basic idea of our method is using established indices for international comparisons, but developing novel combinations of them, as done in technology management to measure a technology's performance (Martino, 1993).

Before calculating these indexes, a choice had to be made regarding the identification of the countries of origin of the patents. Classically, there are two types of information that can be used for this purpose: The country of origin of the patent inventor, or the country of origin of the patent applicant (Zuniga et al., 2009). Other possible sources, used in Gaëtan De Rassenfosse et al. (2013) when more direct information was not available, is the country of

the patent office where the priority filing was registered. The use of this type of information is especially suitable for our analysis: Patent offices enable patent protection in the country where they are located (with some exceptions such as the European Patent Office (EPO), the International Bureau of the WIPO, and the Eurasian Patent Organisation (EAPO), for example, which enable protection across more than one country), and are therefore a relevant proxy for identifying the markets that inventors and companies prioritise when applying for patents. Thus, we chose to define the origin of each patent as the country whose patent office registered the priority filing related to the patent. A disadvantage of this choice is that the country of origin cannot be properly identified when the priority filing is registered in patent offices that cover more than one country, such as the EPO, WIPO or EAPO. Consequently, we may underestimate the performance of countries covered by wide-ranging patent offices.

Once the country of origin of each patent application is defined, the indicators used in this chapter could be easily calculated. To do so, we relied on the variables presented in Appendix C. Our subsequent analysis is divided into 3 steps (see Figure 3).



Figure 3: Description of the steps followed.

First, we referred to the well-known RTA index to analyse the specialisation of countries in the course of time. Second, we used the RTA index calculated on the previous step as one of the inputs of our indicator for National Breeding Grounds, which combines a country’s relative extent of specialisation (RTA value) with the absolute number of AI patent applications (unweighted and weighted) filed in that country. Finally, we calculated the International Breeding Ground index, which takes into consideration the country’s number of patents and their IP extension abroad, as well as the number of foreign patents registered in this given country (also considering unweighted and weighted patents).

2.5.1. AI specialists by Revealed Technology Advantage

We use the RTA index to measure the specialisation advantages of countries. An index result of 0 indicates that the country has no patent in the sector considered, 1 when the company’s share in the sector equals its share in all fields, and above 1 when the country has a positive

specialisation in the sector. We calculate the index for three 12-year periods (see equation (2)).

$$RTA_Country_p = \frac{1}{n} \sum_{t=1}^n \frac{\frac{Priority\ Patents\ AI_Country_{t,p}}{Total\ Patents_Country_{t,p}}}{\frac{Global\ Number\ of\ AI\ Patents_{t,p}}{Global\ Number\ of\ Patents_{t,p}}} \quad (2)$$

The variable “Priority Patents AI_Country_{t,p}” is defined as the number of AI priority filings whose application authority is that of the country considered during year t, which is in period p. Similarly, the variable “Total Patents_Country_{t,p}” represents the total number of priority filings registered by the application authority belonging to the country considered in year t and in period p. The “Global Number of AI Patents_{t,p}” and the “Global Number of Patents_{t,p}” represent the total number of AI priority filings registered and the total number of all priority filings registered, respectively, both at the global level, in year t and period p. As previously stated, t varies between 1 and 12, and p between 1 and 3.

The results for the 20 countries with the highest RTA values for each of the considered periods are presented in Figure 4. The three vertical lines indicate the average RTA value of these 20 countries for the periods considered.

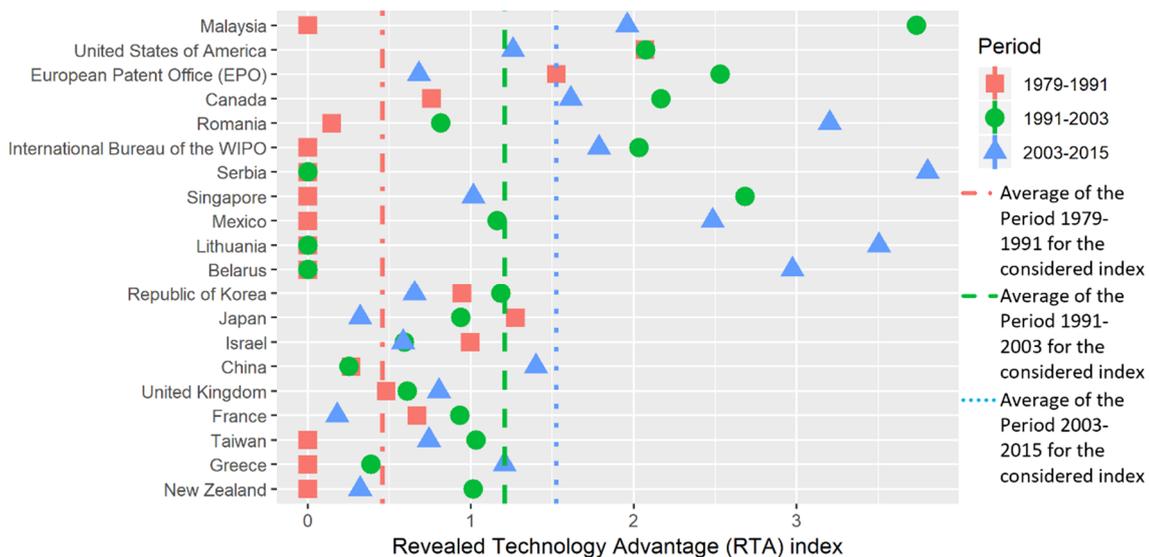


Figure 4: Top 20 Patent Offices according to (and sorted by) the sum of their RTAs over three periods.

Most remarkably, the WIPO number of patent offices with specialisation in AI technologies increases over time, going from two to ten and twelve patent offices with an RTA greater than 1 in periods one, two and three, respectively (Figure 4). The average RTA values increase from one

period to the next, although this increase is greater between periods one and two than between periods two and three.

Furthermore, two main clusters are identified concerning the evolution and RTA values: A European cluster and an Asian cluster. In the European cluster, the Eastern and Southeastern European patent offices (Serbia, Romania, Lithuania and Belarus and Greece) stand out, all sharing the characteristic of having achieved specialisation values only in the third period. The other two patent offices of the European cluster, United Kingdom and France, still present no specialisation. Patent offices of the European cluster still share the characteristic of having increased their RTA values over time, with the exception of France, which is the only patent office in this cluster whose RTA declined in the most recent period.

The Asian cluster includes three out of the four Asian tigers (Singapore, South Korea and Taiwan). Together with Malaysia, these countries share the characteristic of having reached the peak of their specialisation in the 1990s, at the pinnacle of their industrial spurt, decreasing in the following period. Japan and China also belong to this cluster, with very different patterns: Japan was an early specialist in AI patents, and was a leader in the first period (together with the EPO and the United States), but loses its specialisation in the following two periods, while China achieves specialisation status in the third period.

Finally, among the remaining patent offices, Canada and the United States stand out. Both present very similar patterns of evolution, with their RTAs increasing from period one to period two and, surprisingly, decreasing in period 3. Israel and New Zealand also show a decrease in their RTA values in the third period. Patent offices that facilitate the registration of a patent for more than one country, like the EPO and WIPO, also had lower RTAs values in the last period; the EPO lost its positive specialisation, WIPO retained it.

2.5.2. National AI Breeding Grounds

Next, we operationalise the idea of Breeding Grounds, already introduced, in a national context. We use two basic assumptions for our indicator, later adding a third assumption. The first basic assumption is that the characteristic of a country as a National AI Breeding Ground depends on the number of AI-related priority inventions registered directly by this country's patent office. We assume that such patents represent an intention from inventors or companies to prioritise the exploitation of this country's market. The second assumption is that a country's RTA correlates with the extent to which it can be described as an NBG. A high

RTA signals a specialisation into AI technologies in the country, which suggests the availability of relevant inputs, such as skilled labour, needed to exploit a technology. Combining both assumptions enables us to generate our indicator for NBGs, which is the product of the total number of patents related to AI with priority in a country in a period p , filed to the national patent office, and the country's RTA index for the given period p (see equation 3).

$$Nat\ Breeding\ Ground_p = RTA_Country_p \times Priority\ Patents\ AI_Country_p \quad (3)$$

This indicator reduces the effects of high patenting numbers made by countries highly active in patent registration (either by the use of more flexible rules for the registration of patents by the country's application office, or because the country has a higher population, for example). At the same time, it allows us to deal with some of the known problems of using the RTA index, as elaborated in Soete (1987).

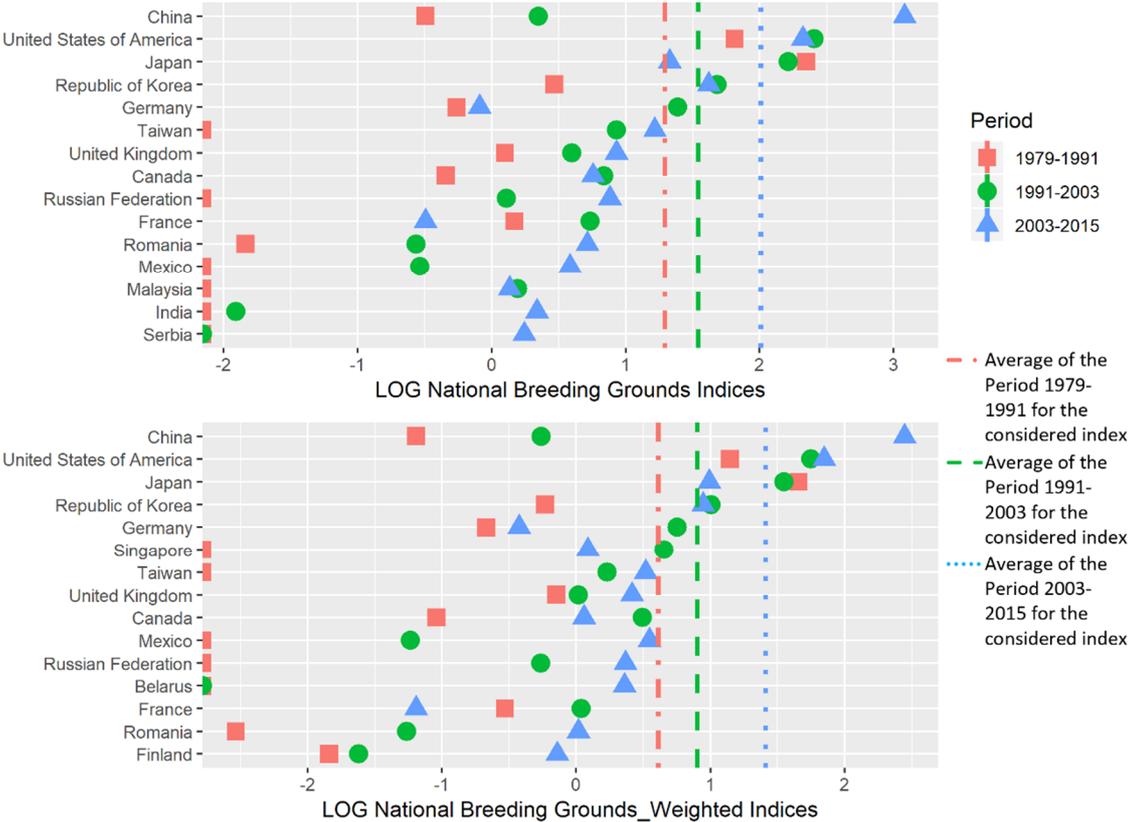
In addition to these two basic assumptions, we assume that PCT applications reflect more valuable inventions than non-PCT applications. Consequently, we calculate a weighted measure (see equation (4)).

$$\begin{aligned} & \text{Weighted Patents}_{AI_Country_p} = \\ & \frac{1}{n} \sum_{t=1}^n \left(\left(\frac{\text{Non - PCT Patents}_{AI_Country_{t,p}}}{5} \right) + \left(\text{PCT Patents}_{AI_Country_{t,p}} \times 5 \right) \right) \end{aligned} \quad (4)$$

Here, "Non – PCT Patents $AI_Country_{t,p}$ " means the number of AI priority filings of "A" type registered by the considered country at the year t for a given period p , while "PCT Patents $AI_Country_{t,p}$ " means the number of PCTs ("W" type) AI priority filings registered by the considered country during year t , which is in period p . With this measure, we calculated a slight variation from our previous equation (3), presented in equation (5).

$$\begin{aligned} & \text{Nat Breeding Ground Weighted Country}_p = \\ & RTA_Country_p \times \text{Weighted Patents}_{AI_Country_p} \end{aligned} \quad (5)$$

The results of equations (3) and (5) for the 15 patent offices³⁷ with the largest National AI Breeding Ground indicators are presented in Figures 5 and 6, respectively. Again, the three vertical lines indicate the average of the presented values for each of the three periods considered. For better visualisation, the Logs10 of the calculated values are presented (both for the indices and for the averages).



Figures 5 (top) and 6 (bottom): Top 15 Patent offices which are considered National AI Breeding Grounds, according to the Nat Breeding Ground_Countryp and the Nat Breeding Ground_Weighted_Countryp indicators, respectively.

Most remarkably, there is a substantial difference between China and the other patent offices, especially in the third period, as well as a large increase in both indicators for China between periods two and three, indicating that China became an NBG for AI more recently. At the same time, Japan and the United States had been the main National AI Breeding Grounds in the 80s and the 90s; Japan declined markedly during the third period, while the United States maintained its position as the second most relevant NBG.

³⁷ Patent offices that cover more than one single country (such as the EPO, the International Bureau of the WIPO and the EAPO) are removed from Figures 5 and 6, as they do not allow a particular country’s market to be identified.

Countries with high specialisation values but with an absolute low number of patents in AI, such as Serbia, Lithuania and Belarus (see Figure 4), cannot be considered NBGs for AI (see Figure 5). On the other hand, countries such as Germany, Russia and India come up in the ranking of NBGs. The increase in the average value of period three in relation to periods one and two is associated mainly with the increase in the values for China, India and Serbia in this period.

The high relevance of Asian countries is maintained (see Figure 5), with three of the four main National AI Breeding Grounds pertaining to Asian patent offices, while the Eastern European patent offices from the European cluster are not so relevant (although Romania and Serbia appear at the bottom of Figure 5). The patent offices from UK, Russia, Canada, Romania and Mexico can perhaps also be considered relevant National AI Breeding Grounds in the third period, but the relevance of Germany and France declines drastically.

The consideration of a higher weight for PCT patents reduces the mean of the patent offices' values (Figure 6). China remains in the first place, but now with a lower value and a smaller difference to the United States. There is an increase in the relevance of the Mexican patent office (regarding its position as an NBG). The consideration of the number of PCTs also favours Singapore, Belarus and Finland (comparing Figures 5 and 6). On the other hand, when the number of PCTs is considered, the patent offices from Taiwan, the UK, Canada, Russian Federation, France and Romania lose positions as NBGs, as do the patent offices of Malaysia, India and Serbia (see Figure 6).

2.5.3. International AI Breeding Grounds

The NBG indicators (weighted and not weighted) focus on the perspective of a single country. They deliberately neglect the registers of patents between countries. To take account of this international perspective, we developed an International Breeding Ground Index (weighted and not weighted). We make three assumptions for this index. First, countries with efficient IP protection and promising market potential for AI exploitation will attract a greater number of AI-inventions owned by inventors/companies from abroad. Second, countries with relevant AI development will seek the exploitation of their IP in promising foreign markets, reflected by their registration abroad of a higher number of national AI inventions. Finally, we assume that an IBG could be reflected by both types of international flows of patent registrations, which leads us to suggest a product rather than a sum for this index.

We introduce two new variables: “*AI patents coming from abroad_Country_{t,p}*” and “*AI patents going abroad_Country_{t,p}*”, which, respectively, represent the number of AI priority filings from patent offices located in other countries registered in the country considered, and the number of priority filings registered by the patent office of this particular country that have also been patented in other countries, during a given year t and in a given period p. To identify which countries are International AI Breeding Grounds, we calculated the product of the two newly introduced variables and related it to two additional variables. First, we related it to the priority AI patents of this country (see equation (6)).

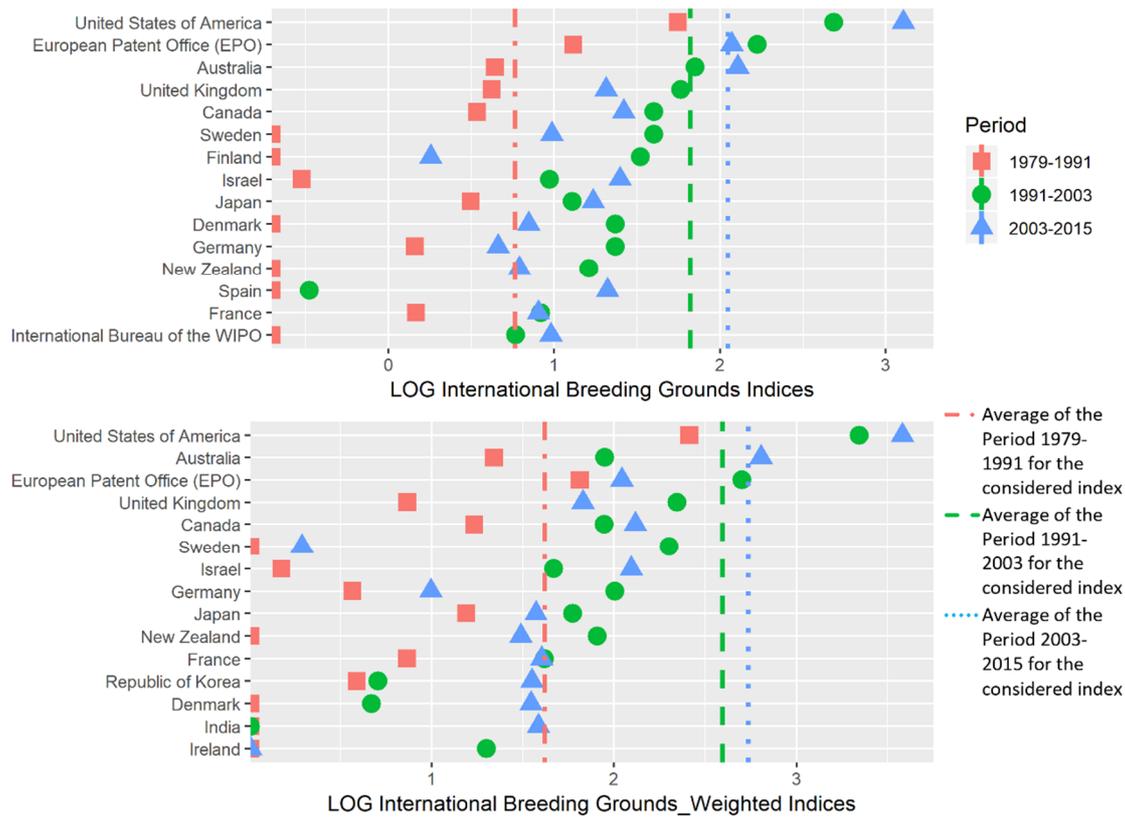
$$\begin{aligned}
 & \text{Int Breeding Grund}_{\text{country}_p} = \\
 & \frac{1}{n} \sum_{t=1}^n \frac{\text{AI patents coming from foreign_country}_{t,p} \times \text{AI patents going to foreign_country}_{t,p}}{\text{Priority Patents AI_Country}_{t,p}}
 \end{aligned}
 \tag{6}$$

Second, we related it to the weighted indicator for NBG (see equation (7)) in order to find out the relationship between National and International Breeding Grounds while also considering the differences between PCT and non-PCT applications.

$$\begin{aligned}
 & \text{Int Breeding Ground}_{\text{Weighted_country}_p} = \\
 & \frac{1}{n} \sum_{t=1}^n \frac{\text{AI patents coming from abroad_Country}_{t,p} \times \text{AI patents going abroad_Country}_{t,p}}{\text{Nat Greenhouse_Weighted_Country}_{t,p}}
 \end{aligned}
 \tag{7}$$

The results of these two measures are presented (see Figures 7 and 8) again for the 15 patent offices³⁸ with the largest values, both using vertical lines to indicate the mean and the Logs10 of the calculated values for better visualisation.

³⁸ For this indicator, international offices such as the EPO, WIPO and EAPO are not excluded. They can support International Breeding Grounds in the sense of attracting AI patents from abroad due to their economic potential as relevant AI markets. Similarly, patent offices without any priority filings are excluded.



Figures 7 (top) and 8 (bottom): Top 15 Patent Offices which are considered International Breeding Grounds, according to the Int Breeding Ground_Country and the Int Breeding Ground_Weighted_Country indicators, respectively.

We observe the United States as a dominant International AI Breeding Ground, with an increase in both of its indicators over all periods (see Figures 7 and 8). The United States is followed by Australia and the EPO, which are characterised by different patterns: Australia has a similar pattern to the United States, while the EPO shows a decrease during the latest period considered. This EPO pattern is also reflected in the trends of most European patent offices. Strikingly, while China is highly relevant as a National AI Breeding Ground, it has no relevance as an IBG. Furthermore, Sweden, Denmark and Spain, which are not considered National AI Breeding Grounds, now appear together with Finland as International AI Breeding Grounds. Moreover, Israel and Japan are shown to be constantly evolving as International AI Breeding Grounds.

In comparison with Figure 7, the consideration of a higher value for PCTs in Figure 8 increases the mean values. The consideration of PCTs also favours Australia, as well as Israel, Germany, New Zealand and France. Moreover, this consideration also shows that South Korea, India and Ireland have emerged as relevant International AI Breeding Grounds. On the other hand, the

position of the EPO, Finland, Denmark and Spain is negatively affected, and WIPO, Spain and Finland disappear from the picture.

2.6. AI Techniques: Evolution, specialisations and Breeding Grounds

To analyse AI techniques, we focused upon 12 techniques cited in more than 100 patents each. This time we made no distinction between the types of patents (PCTs and non-PCTs), thus looking only at the RTAs and at the (non-weighted) NBG and IBG values. As a general overview, the number of patents for each technique³⁹ is presented in Figure 9.

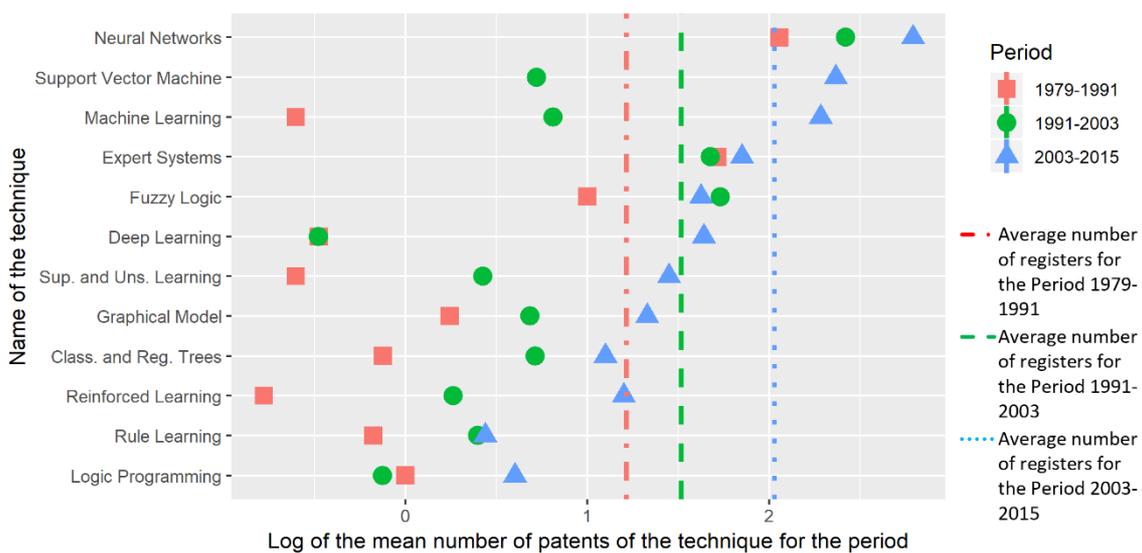


Figure 9: Evolution of the patents of each of the AI-related techniques considered, sorted by the total number of patents for each technique.

We observe a large increase in the number of AI patents associated with the considered techniques (see Figure 9). Far more patent applications related to Neural Networks, Expert systems and Fuzzy logic were registered in the second period than had been registered in the first. From the second to the third periods, however, the number of applications related to Fuzzy Logic systems decreased, while those associated with Machine Learning and Support Vector Machines further increased. By the end of third period, these two, together with Neural Networks, comprised the top three techniques. This movement at the top of the list is also seen farther down. In particular, Deep Learning had the most abrupt increase in this period; techniques related to Rule-based learning and Expert Systems had only a minimal growth in

³⁹ According to the same three periods and considering again the vertical lines for the mean values of each period as well as the Log10 of the real values for better visualisation.

the third period, which might indicate a possible decline in the use of patents related to these techniques in recent applications when the drastic increase of AI patents in this period is considered.

Next we looked at the RTA, as well as the (non-weighted) NBG and IBG values, to identify the two leading Patent offices for each AI technique (see Table 4). This time, no distinction between the periods was made (thus, Period = 1 and t = 36), and only those AI techniques that at least 1,5% of the total sample were considered (thus we excluded Graphical Models, Classification and Regression Trees, Reinforced Learning, Logic Programming and Rule Learning techniques). Together, the seven analysed AI techniques comprise 95.7% of the sample.

AI Technique	Position	RTA Index	National AI Breeding Ground leader	International AI Breeding Ground leader
Neural Networks	1 st	Malaysia	China	US
	2 nd	Serbia	Japan	EPO
Support Vector Machine	1 st	India	China	US
	2 nd	US	US	Australia
Machine Learning	1 st	New Zealand	US	US
	2 nd	US	China	Australia
Expert Systems	1 st	Belarus	Japan	US
	2 nd	Lithuania	US	EPO
Fuzzy Logic	1 st	Mexico	US	US
	2 nd	Romania	Germany	EPO
Deep Learning	1 st	Japan	China	US
	2 nd	Republic of Korea	Japan	Japan
Supervised and Unsupervised Learning	1 st	Singapore	US	US
	2 nd	Serbia	China	EPO

Table 4: National and International Breeding Ground leaders for each AI technique considered.

It turns out that China leads again as a National AI Breeding Ground, being on the top in three out of seven AI techniques considered, followed by the United States and Japan, which are the leaders of two and one AI techniques, respectively. When considering the International AI Breeding Ground index, the United States stands out, being the leader of all AI techniques considered. Furthermore, the United States can be considered the top NBG and IBG in Machine learning as well as in Supervised/Unsupervised Learning.

2.7. Discussion

The analysis of the presented data enables the identification of three perspectives on the evolution of AI. The first is related to the growth in the number of countries involved in AI in terms of specialisation. In this group, Asian and Eastern and Southern European countries stand out, with high levels of specialisation since the 90s and 2000s, respectively. Despite the fact that Eastern and Southeastern European countries specialised in AI later than other countries, a large proportion of their patents are related to AI techniques declining in global relevance. This is the case, for example, of the patent offices of Belarus and Lithuania in techniques related to Expert Systems, and of Romania in techniques related to Fuzzy Logic.

Furthermore, results from analysing NBG and IBG values suggest that Japan and most West European countries have lost their early vanguard status, while China and the United States have increased their leadership. China is leading in a cluster of National AI Breeding Grounds, whereas the United States is not only a significant NBG for AI but is also the leading International AI Breeding Ground for each and every analysed technique.

This signals a major structural difference in the international patterns for IP protection in AI after the 1990s: Some countries have focused on developing their domestic markets and have been less interested in the exploitation of foreign markets, while others have developed AI in an international context. Asian patent offices are by and large in the first group (China, India, Malaysia, Republic of Korea, Taiwan), as are the Russian Federation, Mexico, Romania and Serbia; whereas developed countries from the western hemisphere are the dominant IBGs.

Finally, our findings suggest substantial changes in the relevance of various AI techniques over time, as already prominently documented (Fujii & Managi, 2018). In particular, we confirm that the use of mathematical models (like fuzzy logic), as well as knowledge-based models (like Expert Systems) is decreasing, while Biological and Machine Learning Models (such as

neural networks, supervised and unsupervised machine learning, deep learning and Support Vector Machines) are increasing their relevance as AI techniques. Our investigation adds to these findings by considering the international competition in leadership as well as the above-explained structural differences between the United States and China at the level of AI techniques. Both the United States and China have top positions in terms of NBGs in the AI techniques considered, but only the United States is a top NBG and IBG in AI techniques. This indicates that the United States leads not only in these rapidly growing AI techniques at home but that this happens also in an international context of IP protection. In contrast, China seems to sustain its leading position by focusing on IP protection of its domestic market, almost in isolation from the international context.

Although, the increasing relevance of China in the global arena for AI development cannot be ignored, the structural features identified could be related to several underlying factors. On a very general level, China is characterised by a state-capitalistic approach: AI has become a high political priority in the last decade (Fujii & Managi, 2018; Harhoff, Heumann, Jentzsch, & Lorenz, 2018). It has been documented that in China, universities rather than corporate actors account for a large majority of AI-related patenting (Kroll, 2011; WIPO, 2019b). This applies also to Fujii & Managi (2018), which found that 98% of all registrations at SIPO in the sample originate from Chinese universities. On the other hand, the United States is more associated with a market-driven “Silicon Valley approach” to AI, which is more open and internationally connected. Furthermore, Chinese leadership as the top National AI Breeding Ground is reduced when considering PCT-applications, which signal that China-produced IP is less relevant for protection in international markets. Kroll (2011) points out further evidence suggesting that Chinese universities’ activities are reflected in their patenting behaviour but are not based on inventions marketable enough for international protection. Thus, it is very likely that the output in China-related AI patenting is connected to the incentive structure for Chinese researchers at universities. It is also likely that the extent of commercialisation of AI IP by universities and corporate actors differs, which might help to explain why China scores so low as an IBG, since the prospect of commercialisation motivates corporate actors to seek to extend IP protection in foreign jurisdictions. However, Kroll (2011) further highlights that Chinese MNC subsidiaries are indeed aiming to protect their local IP on the Chinese market, but that they are still not generating innovations relevant enough to file on international lead offices. Moreover, our results could indicate that foreign corporate actors do not register AI-

related IP with SIPO, which might be related to institutional barriers and/or limitations in terms of local enforcement (ibid).

2.8. Conclusions

This chapter analyses trends and structural differences in patenting patterns in AI-related technologies. We propose two novel patent-based indicators to differentiate structural differences in the patterns of IP protection observed for AI across countries. We considered (i) to what extent countries specialise in AI and are relevant markets for corresponding IP protection ('National Breeding Ground'); and (ii) to what extent countries attract AI from abroad for IP protection and extend IP protection of their own AI to foreign markets ('International Breeding Ground'). We demonstrate that NBGs and IBGs overlap only to a limited extent. Primarily, China and the United States can be characterised as dominant NBGs. Australia, selected European countries, but primarily the U.S., are major IBGs. We conclude that China promotes AI development with an almost-exclusive focus on IP protection in the domestic market, whereas the United States sustains its AI progress in an international context, too. This might indicate a considerable bifurcation in structural patterns of IP protection in global AI development. We discussed possible explanations related to the institutional particularities of the Chinese National Innovation System.

This chapter contributes to the broader debate by introducing and operationalising the concepts of NBGs and IBGs. The proposed approach, in general, can be used as a reference for further patent mining and technology innovation analysis of other technical or scientific fields. However, we acknowledge limitations in our approach, which include the data source, the method, and the indicators themselves. First, we based our analyses exclusively on the patent process. However, actors might use other means to disclose their inventions (e.g., by defensive publication in scientific journals, pre-print servers or platforms); attention exclusively on patents might miss valuable innovations in the area of AI. Nevertheless, we were interested in longer-term trends and dynamics, and it is here that patents offer a suitable source of internationally standardised information available in a longer time series. Second, we used priority filings as a proxy for the first market on which companies, other organisations and inventors aim to protect an invention. By doing this, we can indicate the market impact as well as the markets impacted by a given technology. However, this approach neglects the development aspects of AI, which could be captured by the location of the inventors.

Furthermore, we use patent applications rather than granted patents, which introduces potentially non-relevant IP. Some applicants are not granted patents, as many as seven years might be required before the patent office makes a decision and, most problematic, both the proportion of applications granted and the processing time vary from office to office. Third, we try to account for the quality of patent application by considering PCT registers. Yet, there exist more comprehensive quality measures for patents, such as citations, renewal rates or high-impact inventions. Incorporating these kinds of data could help to reduce the “noise” that surely appears in the comparative analysis presented here. Fourth, we pay limited attention to variations in the particular characteristics of patentability between patent offices (e.g., highlighted in Meguro & Osabe (2019)), which might also introduce a “quality bias” into our dataset. Finally, we restricted our keyword search to title and abstract. This could be improved by considering the claims in the whole patent document. The challenge would be to differentiate dependent and independent claims in the identification strategy, but if this challenge could be overcome, this approach could not only improve the identification strategy, but also reduce the potential risk of “fashionable labelling” trends that discourage the use of keyword-based searches.

CRedit authorship contribution statement

Matheus Eduardo Leusin: Formal analysis, Investigation, Data curation, Writing - original draft, Visualization. Jutta Günther: Conceptualization, Resources, Writing - review & editing, Supervision. Björn Jindra: Conceptualization, Methodology, Investigation, Validation, Writing - review & editing, Supervision. Martin G. Möhrle: Conceptualization, Methodology, Validation, Project administration, Investigation, Writing - original draft.

References

- Annoni, A., Benczur, P., Bertoldi, P., Delipetrev, B., De Prato, G., Feijoo, C., . . . Junklewitz, H. (2018). *Artificial Intelligence: A European Perspective*.
- Aristodemou, L., & Tietze, F. (2018). The state-of-the-art on Intellectual Property Analytics (IPA): A literature review on artificial intelligence, machine learning and deep learning methods for analysing intellectual property (IP) data. *World Patent Information*, 55, 37-51.
- Balassa, B. (1965). Trade liberalisation and “revealed” comparative advantage 1. *The manchester school*, 33(2), 99-123.
- Bassis, N. F., & Armellini, F. (2018). Systems of innovation and innovation ecosystems: a literature review in search of complementarities. *Journal of evolutionary economics*, 28(5), 1053-1080.
- Cockburn, I. M., Henderson, R., & Stern, S. (2018). *The Impact of Artificial Intelligence on Innovation*.
- De Rassenfosse, G., Dernis, H., & Boedt, G. (2014). An introduction to the Patstat database with example queries. *Australian Economic Review*, 47(3), 395-408.
- De Rassenfosse, G., Dernis, H., Guellec, D., Picci, L., & de la Potterie, B. v. P. (2013). The worldwide count of priority patents: A new indicator of inventive activity. *Research policy*, 42(3), 720-737.
- EPO. (2018). Data Catalog – PATSTAT EP Register – 2018 Autumn Edition. In.
- Fujii, H., & Managi, S. (2018). Trends and priority shifts in artificial intelligence technology invention: A global patent analysis. *Economic Analysis and Policy*, 58, 60-69. doi: <https://doi.org/10.1016/j.eap.2017.12.006>
- Gabrilovich, E., & Markovitch, S. (2006). *Overcoming the brittleness bottleneck using Wikipedia: Enhancing text categorization with encyclopedic knowledge*. Paper presented at the AAAI.
- Gürerk, Ö., Irlenbusch, B., & Rockenbach, B. (2006). The competitive advantage of sanctioning institutions. *science*, 312(5770), 108-111.
- Harhoff, D., Heumann, S., Jentzsch, N., & Lorenz, P. (2018). Outline for a German Strategy for Artificial Intelligence.
- Hekkert, M. P., Suurs, R. A., Negro, S. O., Kuhlmann, S., & Smits, R. E. (2007). Functions of innovation systems: A new approach for analysing technological change. *Technological Forecasting and Social Change*, 74(4), 413-432.
- Huang, L., Miao, W., Zhang, Y., Yu, H., & Wang, K. (2017). *Patent Network Analysis for Identifying Technological Evolution: A Case Study of China's Artificial Intelligence Technologies*. Paper presented at the Management of Engineering and Technology (PICMET), 2017 Portland International Conference on.
- Keisner, C. A., Raffo, J., & Wunsch-Vincent, S. (2015). *Breakthrough technologies—Robotics, innovation and intellectual property* (Vol. 30): WIPO.
- Klinger, J., Mateos-Garcia, J., & Stathoulopoulos, K. (2018). Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology. doi: <https://arxiv.org/pdf/1808.06355.pdf>
- Kroll, H. (2011). Exploring the validity of patent applications as an indicator of Chinese competitiveness and market structure. *World Patent Information*, 33(1), 23-33.
- Lapenne, J. (2010). Patent Cooperation Treaty (PCT). *J. Pat. & Trademark Off. Soc'y*, 92, 192.
- Li, X., & Jiang, H. (2017). Artificial Intelligence Technology and Engineering Applications. *Applied Computational Electromagnetics Society Journal*, 32(5).
- Lupu, M. (2018). Artificial Intelligence and Intellectual Property (Editorial). *World Patent Information*, 53, A1-A3. doi: <https://doi.org/10.1016/j.wpi.2018.06.001>
- Mansourzadeh, M. J., Shahmoradi, B., Dehdarirad, H., & Janavi, E. (2019). A note on using revealed comparative advantages in scientometrics studies. *Scientometrics*, 121(1), 595-599.
- Martino, J. P. (1993). *Technological forecasting for decision making*: McGraw-Hill, Inc.
- Meguro, K., & Osabe, Y. (2019). Lost in Patent Classification. *World Patent Information*, 57, 70-76.
- Menzel, K., & Maicher, L. (2017). A novel method for retrieving specialisation profiles—The case of patent agent firms. *World Patent Information*, 51, 46-56.

- Nilsson, N. J. (2009). *The quest for artificial intelligence*: Cambridge University Press.
- Niu, J., Tang, W., Xu, F., Zhou, X., & Song, Y. (2016). Global Research on Artificial Intelligence from 1990–2014: Spatially-Explicit Bibliometric Analysis. *ISPRS International Journal of Geo-Information*, 5(5), 66.
- Radauer, A., Rosemberg, C., Cassagneau-Francis, O., Goddar, H., & Haarmann, C. (2015). *Study on the economic impact of the utility model legislation in selected Member States: Final Report. A study tendered by the European Commission–DG Internal Market and Services in 2013*.
- Sharma, P., & Tripathi, R. (2017). Patent citation: A technique for measuring the knowledge flow of information and innovation. *World Patent Information*, 51, 31-42.
- Sharma, P., Tripathi, R., & Tripathi, R. (2015). Finding similar patents through semantic query expansion. *Procedia Computer Science*, 54, 390-395.
- Soete, L. (1987). The impact of technological innovation on international trade patterns: the evidence reconsidered. *Research policy*, 16(2-4), 101-130. doi: [https://doi.org/10.1016/0048-7333\(87\)90026-6](https://doi.org/10.1016/0048-7333(87)90026-6)
- Thoma, G. (2014). Composite value index of patent indicators: Factor analysis combining bibliographic and survey datasets. *World Patent Information*, 38, 19-26.
- Tseng, C.-Y., & Ting, P.-H. (2013). Patent analysis for technology development of artificial intelligence: A country-level comparative study. *Innovation*, 15(4), 463-475.
- Van Zeebroeck, N., & Van Pottelsberghe de la Potterie, B. (2011). Filing strategies and patent value. *Economics of innovation and new technology*, 20(6), 539-561.
- Weresa, M. A. (2019). Technological competitiveness of the EU member states in the era of the fourth industrial revolution. *Economics and Business Review*, 5(3).
- WIPO. (2019a). The PCT now has 152 Contracting States. Retrieved from https://www.wipo.int/pct/en/pct_contracting_states.html
- WIPO. (2019b). *WIPO Technology Trends 2019: Artificial Intelligence*. Retrieved from https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf
- Zuniga, P., Guellec, D., Dernis, H., Khan, M., Okazaki, T., & Webb, C. (2009). OECD patent statistics manual. *Francia: OECD Publications*.

Appendix A: Query for identifying AI application IDs.

Select appln_id from t1s202_appln_title

Where appln_title like '%Artificial intelligence%' OR appln_title like '%machine learn%' OR appln_title like '%Probabilistic reason%' OR appln_title like '%Fuzzy logic%' OR appln_title like '%Logic Programming%' OR appln_title like '%Ontology engineer%' OR appln_title like '%pervised learn%' OR appln_title like '%reinforced learn%' OR appln_title like '%task learn%' OR appln_title like '%neural network%' OR appln_title like '%deep learn%' OR appln_title like '%expert system%' OR appln_title like '%support vector machin%' OR appln_title like '%description logistic%' OR appln_title like '%classification tree%' OR appln_title like '%regression tree%' OR appln_title like '%logical learn%' OR appln_title like '%relational learn%' OR appln_title like '%probabilistic graphical model%' OR appln_title like '%rule learn%' OR appln_title like '%instance-based learn%' OR appln_title like '%latent represent%' OR appln_title like '%bio-inspired approach%' OR appln_title like '%machine intelligen%' OR appln_title like '%probability logic%' OR appln_title like '%probabilistic logic%' OR appln_title like '%reinforcement learn%' OR appln_title like '%multitask learn%' OR appln_title like '%Decision tree learn%' OR appln_title like '%support vector network%' OR appln_title like '%deep structured learn%' OR appln_title like '%hierarchical learn%' OR appln_title like '%graphical model%' OR appln_title like '%structured probabilistic model%' OR appln_title like '%Rule induction%' OR appln_title like '%memory-based learn%' OR appln_title like '%bio-inspired comput%' OR appln_title like '%biologically inspired comput%'

UNION

Select appln_id from t1s203_appln_abstr

Where appln_abstract like '%Artificial intelligence%' OR appln_abstract like '%machine learn%' OR appln_abstract like '%Probabilistic reason%' OR appln_abstract like '%Fuzzy logic%' OR appln_abstract like '%Logic Programming%' OR appln_abstract like '%Ontology engineer%' OR appln_abstract like '%pervised learn%' OR appln_abstract like '%reinforced learn%' OR appln_abstract like '%task learn%' OR appln_abstract like '%neural network%' OR appln_abstract like '%deep learn%' OR appln_abstract like '%expert system%' OR appln_abstract like '%support vector machin%' OR appln_abstract like '%description logistic%' OR appln_abstract like '%classification tree%' OR appln_abstract like '%regression tree%' OR appln_abstract like '%logical learn%' OR appln_abstract like '%relational learn%' OR appln_abstract like '%probabilistic graphical model%' OR appln_abstract like '%rule learn%' OR appln_abstract like '%instance-based learn%' OR appln_abstract like '%latent represent%' OR appln_abstract like '%bio-inspired approach%' OR appln_abstract like '%machine intelligen%' OR appln_abstract like '%probability logic%' OR appln_abstract like '%probabilistic logic%' OR appln_abstract like '%reinforcement learn%' OR appln_abstract like '%multitask learn%' OR appln_abstract like '%Decision tree learn%' OR appln_abstract like '%support vector network%' OR appln_abstract like '%deep structured learn%' OR appln_abstract like '%hierarchical learn%' OR appln_abstract like '%graphical model%' OR appln_abstract like '%structured probabilistic model%' OR appln_abstract like '%Rule induction%' OR appln_abstract like '%memory-based learn%' OR appln_abstract like '%bio-inspired comput%' OR appln_abstract like '%biologically inspired comput%'

Appendix B: IPC-search strategy adopted in papers used for comparison.

	Fujii & Managi, 2018 (Query 2)	Tseng & Ting, 2013 (Query 3)
IPC codes	G06N 3/00, G06N 3/02, G06N 3/04, G06N 3/06, G06N 3/063, G06N 3/067, G06N 3/08, G06N 3/10, G06N 3/12, G06N 5/00, G06N 5/02, G06N 5/04, G06N 7/00, G06N 7/02, G06N 7/04, G06N 7/06, G06N 7/08, G06N 99/00	G05B 13/02, G06E 1/00, G06E 3/00, G06F 9/44, G06F 15/00, G06F 15/18, G06F 17/00, G06F 17/20, G06G 7/00, G06J 1/00, G06N 3/00, G06N 3/02, G06N 3/04, G06N 3/08, G06N 3/10, G06N 3/12, G06N 5/00, G06N 5/02, G06N 5/04, G06N 7/00, G06N 7/02, G06N 7/04, G06N 7/06, G06N 7/08, G06N 99/00

Appendix C: Variables considered and their definitions, according to the authors.

Variable	Meaning	Source
Global Number of AI Patent _{St,p}	Total number of patents related to AI globally in year t in a period p	Patstat 2017
Global Number of Patents _{t,p}	Total number of patents related to AI globally in year t in a period p	Patstat 2017
Priority Patents_Country _{t,p}	Total number of patents with priority in a country in year t in a period p	Patstat 2017
Priority Patents AI_Country _{t,p}	Total number of AI patents with priority in a country in year t in a period p	Patstat 2017
No PCT Patents AI_Country _{t,p}	Total number of patents related to AI with priority in a country in year t in a period p ('A' type patents in PATSTAT), filed to the national patent office	Patstat 2017
PCT Patents AI_Country _{t,p}	Total number of patents related to AI with priority in a country in year t in a period p ('W' type patents in PATSTAT), filed to WIPO	Patstat 2017
Weighted Patents AI_Country _p	Indicator related to the total number of AI patents with priority in a country in a period p, but with the distinction of taking different weights for PCT and non-PCT applications.	Authors calculations
AI patents coming from abroad_Country _{t,p}	Total number of patents related to AI with priority in a different country in year t in a period p	Patstat 2017
AI patents going abroad_Country _{t,p}	Total number of patents related to AI going abroad with priority in a country in year t in a period p	Patstat 2017
RTA_Country _p	Revealed technology advantage of a country in period p	Authors calculations
Nat Breeding Ground_Country _p	Indicator for National Breeding Ground in a country in period p	Authors calculations
Nat Breeding Ground_Weighted_Country _p	Indicator for National Breeding Ground in a country in period p, but with the distinction of taking different weights for PCT and non-PCT applications.	Authors calculations
Int Breeding Ground Country _p	Indicator for International Breeding Ground in a country in period p	Authors calculations
Int Breeding Ground_Weighted_Country _p	Indicator for International Breeding Ground in a country in period p, but with the distinction of taking different weights for PCT and n non-PCT applications.	Authors calculations
t	Index for years in a period p (ranging from 1 to 12 for each period p)	According to definition
p	Index for periods (ranging from 1 to 3)	According to definition
n	Number of years in a period p (12 years in our case)	According to definition

Authors: Matheus Eduardo Leusin, Björn Jindra, and Daniel S. Hain

Abstract: The concept of relatedness has been crucial to understanding the geography of innovation. It helps us to explain local innovation patterns such as why a location develops a given technology. Even so, we lack an understanding of how technological relatedness changes over time, and of how these changes affect local technological development. This question is of particular importance for policies aimed at inducing the development of selected promising technologies. In this study, we investigate how Artificial Intelligence (AI) core technologies change over time, how these changes affect the development of AI at the country level, and how AI's new emerging technological trajectories have been incorporated into countries' existing technological trajectories. We focus on the United States, Japan, South Korea, and China, which led AI development during the observation period (1974 – 2018). Using patent data, we apply a technological space perspective coupled with specialisation indices to distinguish between dynamics occurring at the national level from the ones occurring at the technological level. We find that the core technologies used to create AI innovations changed over time, which affected how these technologies relate to each other. This technological evolution of AI is shown to have little association with how it was locally developed. Instead, AI's local development followed countries' existing knowledge bases, even in cases when it was weakly related to AI. This pattern is reinforced as more knowledge about AI is accumulated, so that new AI capabilities developed locally increasingly coincide with countries' existing "general" capabilities.

Keywords: Artificial Intelligence; technological space; evolutionary economic geography; technological relatedness; knowledge complexity;

JEL Classification: O14; O33; O57; D83

Publication

This chapter was submitted as an original research article to a journal on Oct. 2021 and is currently under review.

3.1. Introduction

In the late 1990's, a cell phone was used to tell a friend about one's first-time access to the internet on a personal computer. Back then, phones and computers had very little in common. The cell phone, maybe produced by the then-popular Finish company Nokia, was used mainly to make phone calls (especially when the internet line was occupied); the personal computer, maybe a PC from the American IBM, was used mainly for computing, word processing and perhaps playing games. In the early 2020's, it is hard to differentiate one from the other. Both devices do many of the same things and may be produced by the same company, which may also produce tablets, smartwatches, smart TVs, and other smart products.

This similarity between devices is captured by the concept of relatedness, which considers that some elements share commonalities that favour them to be developed in tandem. The concept, translated to a geographical context in Hidalgo et al. (2007), was highly influential, creating what became known as the "evolutionary turn" in the field of economic geography (Boschma and Martin, 2007). Subsequent research stressed the role of relatedness in explanations of a variety of local innovation phenomena, like the economic (Balland, 2016; Santoalha et al., 2021) and scientific (Boschma, Heimeriks, et al., 2014) activities that a geography enters or exits, regional patterns of entrepreneurship (Ejdemo and Örtqvist, 2020), innovation performance and productivity of firms (Aarstad et al., 2016), internationalisation patterns of start-ups (Naldi et al., 2020), and other geographically relevant developments.

Despite consistent empirical evidence confirming the role of relatedness on local innovation patterns, the fact that it changes over time is still largely neglected (Hidalgo, 2021; Juhász et al., 2021). In our illustrative example, this is similar to saying that the commonalities between cell phones and computers have not changed since the '90s. As it happened, the two devices assimilated similar technologies, which made them more alike: Both now have incorporated the internet, user-friendly operating systems, faster microprocessors, larger data storage, high-resolution colourful displays, etc. This similarity of built-in technologies, in turn, leads to knowledge commonalities that favour their current development in tandem: If a company has the capabilities needed to develop a computer, little⁴⁰ learning effort may be needed to develop a smartphone.

⁴⁰ I.e., when compared to technologies that don't share any commonalities.

Thereby we argue that relatedness is a feature shaped by two distinct components: A local one influenced by geographically bounded factors, and a technological one influenced by the unbounded development of technologies. Despite current neglect in explicitly recognising it, the technological component of relatedness is relevant, particularly for emerging discussions on the role of relatedness in science, technology and innovation (STI) policy. As relatedness changes, distinct technologies are combined and used within industries and distinct patterns of co-evolution emerge. A failure to recognise these patterns may mean wasting limited resources by promoting industries and technological alternatives that are not locally relevant anymore. There are also possible risks of lock-in or disruption of technological trajectories generated by self-reinforcing local co-evolution of technologies when this process is influenced by public policies (Hoppmann, 2021).

To advance the literature on the dynamic aspect of relatedness, it is necessary to understand how technologies evolve and are incorporated into a country's knowledge base. Or, more precisely, how technological relatedness changes and the possible relations of these changes on local technological trajectories. We analyse such dynamics by considering the evolution of Artificial Intelligence (AI) in countries that lead the development of this technology. AI is a suitable choice for our purpose because it allows for the creation of very distinct local technological trajectories, which is especially useful for understanding how local knowledge influences innovation patterns. AI is an umbrella term for technologies that incorporate some kind of human-like intelligence and are modular (Nilsson, 2009), transversal (Righi et al., 2020), and digital (Teece, 2018) in nature. These aspects create multiple entry points for AI to be incorporated into local knowledge. Modularity allows AI to be combined with other technologies, while its digital aspect means it can be coupled into objects with physical materiality to create digital innovations (Yoo et al., 2012). Transversality, in turn, refers to the possible use of AI across a variety of technological sectors. Combined, these characteristics potentially allow for the emergence of very distinct local technological trajectories related to AI: One can develop it by focusing on software-related innovations, or coupling AI with existing physical products from a variety of technology sectors to generate physical innovations.

We assume that countries leading AI development over the last decades had successful trajectories in incorporating this technology, so we focus on the United States, Japan, South

Korea, and China as the locations⁴¹ where AI progressed the most during the observation period (1974–2018). We analyse i) how the relatedness of global AI innovations changed over time, ii) how these changes affected the local exploration of AI in the four countries leading its development, and iii) how the emerging technological trajectories of AI were incorporated (or not) into these countries' existing "general" technological trajectories.

Besides highlighting the dynamic aspect of relatedness and its role on innovation patterns, this chapter also contributes to further understanding the emergence of local lock-in effects – which we attribute to the dominance of local knowledge over global technological development. Regarding relatedness, our findings indicate that the technological relatedness of global AI innovations is characterised by changing innovation patterns through its development. However, these innovation patterns were not simply replicated by local developments. Instead, countries developed AI by building upon their existing knowledge bases, i.e., by developing AI innovations in fields in which they already had a comparative advantage. As more countries accumulated knowledge about an AI element, this pattern was reinforced, so that AI's technological trajectories converged with existing local trajectories. Hence, we argue that AI "breaks-in" to local technological trajectories by being incorporated mostly in areas where a country holds an existing comparative advantage. Another possible situation, in which AI would "break-through" existing technological trajectories by emerging independently from local knowledge, is shown to produce comparative advantages that are short-lived.

This chapter is structured as follows. Section 2 presents the relevant theory on relatedness and formalises the research questions. Section 3 describes the data and method, and empirical findings are presented in Section 4. Finally, Section 5 discusses the main findings, outlines theoretical contributions, policy implications, and existing limitations of our approach.

⁴¹ In economic geography, location is used to refer to cities, regions, or countries. We focus on the latter and use the words "location" and "country" interchangeably throughout the paper. Please also note that the term "location" in this paper refers to political geographies, i.e., spatial structures which are also influenced by political processes. This means that when we refer, for example, to the South Korean "location", we are considering that its innovation outcomes are not shaped only by spatial but also political aspects (which is important for innovative outcomes; think about the difference between South Korea and its closely spatially located North Korea: Spatial differences are not enough to explain the drastic difference between the innovative outcome of these countries).

3.2. Theoretical background

3.2.1. The concept of relatedness

The consideration of how distinct technologies relate to each other has been crucial to explaining of variations in local technological development (Balland, 2016; Boschma, Balland, et al., 2014; Petralia et al., 2017; Rigby, 2015; Santoalha et al., 2021) and the emergence of new industries (Colombelli et al., 2014; Feldman et al., 2015; Neffke et al., 2011; Tanner, 2016). The underlying relationship between technologies is captured through the concept of relatedness, which is tied to the idea of absorptive capacity (Hidalgo, 2021), which in turn refers to the premise that a firm's ability to absorb new knowledge depends on its prior level of related knowledge (Cohen and Levinthal, 1990). In 2003, Breschi et al. proposed the concept considering the commonalities between distinct types of knowledge. Hidalgo et al. (2007) extended it by arguing that local factors also affect the creation of commonalities. These local factors include the institutions, infrastructure, physical factors, and technologies existing in a given location. These commonalities explain distinct patterns of innovation seen at the geographical level.

The “knowledge space”⁴² framework proposed in Hidalgo et al. (2007) is commonly depicted as a network in which nodes represent knowledge categories, such as technological or scientific fields, and links between them represent their degree of relatedness (Balland, 2016). This visual representation illustrates how technologies are related, and how local technological trajectories are shaped by these relations. This framework became a pillar of Economic Complexity (Hidalgo, 2021) and was rapidly incorporated into the location-based analyses of the Evolutionary Economic Geography (EEG) literature. Its consistency to measure the relation between the emergence and development of distinct forms of knowledge gave new tools to EEG for explaining historical patterns of spatial concentration of knowledge.

The empirical evidence produced in this literature shows that knowledge concentration patterns are strongly linked to existing local capabilities, including technological trajectories seen in cities (Boschma, Balland, et al., 2014; Rigby, 2015), regions (Buarque et al., 2020; Colombelli et al., 2014; Ejdemo and Örtqvist, 2020; Van Den Berge and Weterings, 2014), and

⁴² Although “knowledge space” in the broader generalization of the framework, the original term used in Hidalgo et al. (2007) was “product space”, once the focus was products exports; later variations like “technological space” and “scientific space” were used to refer to technologies and publications, for example, as highlighted in Balland et al. (2016).

countries (Hidalgo and Hausmann, 2009; Petralia et al., 2017). It also explains, for example, the emergence of knowledge related to radical technologies. In particular, Tanner (2016) shows that existing knowledge in areas relevant to fuel cell industries explains the local emergence of this disruptive technology. The higher the variety of specialisations of regions in related fields, the more likely it is that they develop this type of industry (ibid).

3.2.2. Relatedness and STI policies

STI policies can be divided into two broad categories regarding the verticality of interventions: Top-down policies that are planned based on a specific demand (e.g., developing locally a specific technology or industrial sector) and bottom-up policies that are planned based on supply (e.g., incentivising the further development of existing local capabilities).

While top-down policies may overlook the role of relatedness by neglecting existing local capabilities – and thus, lead to the emergence of technologies and industries that break through existing knowledge bases – they have shown remarkable success in mission-oriented projects (after all, they led the man to the moon (Uyarra et al., 2020)) and in developing new industries and innovation clusters through techno-industrial policies (see for example Breznitz (2007); Chen and Naughton (2016); Etzkowitz and Brisolla (1999)). Bottom-up policies, in turn, received more attention just recently and are currently being introduced in several countries through “smart specialisation” policies (Hidalgo, 2021; Uyarra et al., 2020)⁴³. By leveraging local knowledge, smart specialisation policies favour mechanisms of technological speciation (Adner and Levinthal, 2002) and industrial branching (Frenken and Boschma, 2007) – which consider that new technologies and industries may emerge as variations of the existing ones – and hence they favour the break-in of new technologies into existing knowledge bases.

Both top-down and bottom-up approaches are deeply affected by the dynamic nature of relatedness. The promising technology pursued this year may become irrelevant in the next, when an even more promising technological alternative emerges. Similarly, policies leveraging the further development of local knowledge may cause a “lock-in” in poor technological alternatives if the alternatives adopted locally are not as competitive as foreign ones.

⁴³ Relatedness has been also highlighted recently as relevant for the development of problem-oriented cluster policies. See for example Grashof (2021).

3.2.3. Research gap

Despite the central role of relatedness in the EEG literature and in the “knowledge space” framework, there is surprisingly little consideration of how it emerges and changes over time (Hidalgo, 2021). Relatedness is treated almost as an exogenous factor (Juhász et al., 2021). Only selected recent studies analyse relatedness as being dynamic and endogenously created. For example, Kuusk and Martynovich (2021) find that inter-industry relatedness changes considerably over time and that this change influences regional employment growth. Juhász et al. (2021) focus specifically on technological relatedness. They find that co-location of technologies and relatedness not only change over time but also affect each other: The more two technologies overlap within spatial distributions, the greater is the change in their relatedness. As relatedness between two technologies increases, so does the probability of them being co-located in the same geographical space (ibid).

Although geography – in terms of location-specific factors – plays a crucial role, relatedness is also affected by technological changes. In particular, Juhász et al. (2021) find that when two distinct technologies are combined once, the likelihood that they are combined again (and thus become more related) increases. Similarly, the literature on recombinant innovations stresses the possibility that previously unrelated technologies may become related through new technological combinations (Castaldi et al., 2015; Frenken et al., 2012), regardless of geography. The underlying mechanism is that successful technological variations are selected and disseminated globally (Arthur, 2009).

In summary, successful new variations are disseminated globally through the repetition of the more efficient technological combinations in further innovations, which increases the relatedness between the combined technologies. Local conditions, in turn, affect how these successful new variations are disseminated, generating distinct geographical patterns of adoption.

Albeit not yet recognising directly the existence of this mechanism, the EEG literature does find evidence of established firms leaving or declining in specific locations (Boschma, Balland, et al., 2014; Neffke et al., 2011; Rigby, 2015). As the possibility of “foreign” technological development is not yet addressed, these exits are presented without further discussions on why they happen. This effect has alternatively been discussed through the notion of path dependency and technological lock-in (Arthur, 1989; Cowan, 1990). The basic idea here is that

markets may get “stuck” in an inferior technological alternative due to network effects. In our geographically focused context, this would mean local technological trajectories not adopting a “better” technological alternative developed elsewhere. As a result, industries and regions that don’t cope with more efficient technological alternatives lose markets or fail, which would explain the local decline and/or exit of firms.

3.2.4. Research questions

The approach presented in Hidalgo et al. (2007) is, in our view, particularly suited to address the existing literature gap: They compare “unbounded” development to local patterns, arguing that the focus on an outcome-based measure (i.e., relatedness) allows capturing the factors that affect the emergence of geographical-related commonalities. We argue that the same reasoning applies to technological change: By focusing on an outcome-based measure of technological development (i.e., patents related to a given technology), relatedness captures the relevant technological trajectories linked to the emergence of a technology. These technological trajectories reflect innovation patterns and technical solutions taken over time with the technology. They are generated through the innovation process and occur both because the technology is used to solve distinct problems as it evolves, and because it is itself the result of distinct technological combinations that may vary over time. As a network perspective is especially useful to capture dynamic aspects of technological knowledge (Antonelli et al., 2010), we argue that the knowledge space framework is also particularly suited to identifying such technological dynamics.

Therefore, we focus on global innovations of a particular technology to identify how its use changes over time. Considering distinct intervals, we aim to identify technologies used to create AI innovations, the possible existence of technologies that are central to the development of AI innovations (which we term “AI-core technologies”), and possible changes over time in the use of these technologies to create AI innovations. We are interested in the following research questions:

1. Which technologies are used to create AI innovations, and are there any “AI-core technologies” central to the development of AI innovations?
2. Do innovation patterns and “AI-core technologies” change over time?

Using a network perspective that places related technologies close to each other and is composed of the technologies used to develop AI patents, one can then identify which

technologies are used and if any are central to the development of AI innovations. These technologies are to be the ones most often used to create AI innovations, and as a result they will be the most connected technologies. By capturing how innovation patterns related to a technology possibly change, this network perspective allows identifying the aspect of relatedness of dynamically changing due to technological evolution.

The other side of this possible dynamic nature of relatedness is how it affects local exploration. Specifically, countries leading the development of a given technology are to be the ones more sensible to adopting the newest technological developments. Thus, the local technological development of AI may lead to two outcomes: The change in the global innovation patterns of AI can be reflected on how countries innovate with AI; or these changes may have little or no influence, and innovation in a country continues to follow existing technological trajectories.

To understand these distinct outcomes, we need to disentangle how AI innovations are created in comparison to all other innovations that are developed nationally. This means separating AI local technological trajectories from “general” local technological trajectories. Accordingly, we make a distinction between two types of capabilities that can be created through technological development within a country: “General capabilities”, which refer to specialisations that a country has when all of its innovations are considered; and “AI-specific capabilities”, which refer to the specialisations that a country has when only its AI innovations are considered. We consider the dynamics between local AI-specific capabilities and global innovation patterns with AI in our third research question:

3. How do changes in AI’s global innovation patterns affect the development of “AI-specific capabilities” by countries leading AI development?

As pointed out in Hidalgo (2021), relatedness can anticipate changes in local specialisation patterns. Accordingly, we expect that any changes in global innovation patterns related to AI are reflected in the way leading countries create AI innovations, thus affecting their “AI-specific capabilities”. This is to say that successful recombinations developed worldwide with AI will be locally adopted by the countries that lead its development.

Finally, it is also important to consider how an emerging (and dynamically evolving) technology gets incorporated into countries’ knowledge bases. This aspect is linked to the local “General capabilities” and how they interact with the new capabilities developed locally with the

emerging technology. We suggest three main possibilities: AI's technological trajectories may become increasingly separated from countries' general technological trajectories (e.g., "AI-specific capabilities" rarely coinciding in the same fields as "General capabilities"), both technological trajectories may increasingly converge (i.e., "AI-specific capabilities" and "General capabilities" progressively coinciding in the same fields), or a mixed technological trajectory may emerge (e.g., "AI-specific capabilities" sometimes leading to the creation of new "General capabilities" in the same fields or vice-versa). Accordingly, our final research question is:

4. How do the "AI-specific capabilities" and "General capabilities" of a country interact to generate new local technological trajectories in AI?

Based on Tanner (2016) we expect that the leading countries considered have "General capabilities" in a variety of fields highly related to AI when they start exploring this technology, which could explain their technological leadership in AI. But as they accumulate AI knowledge, we expect them to create "AI-specific capabilities" in fields that are technologically relevant to the development of AI innovations, regardless of the direction of the "General" technological trajectory (i.e., AI starts to breakthrough existing general technological trajectory by creating its own, which follows global developments made in AI instead of local technological trajectories). This differs from the alternative in which "AI-specific capabilities" would emerge in fields following countries' existing capabilities (i.e., instead of breaking-in existing technological trajectories).

3.3. Data and methods

Following Balland et al. (2019); Boschma, Balland, et al. (2014); Feldman et al. (2015); Rigby (2015); Whittle (2020), we combine a technological space perspective with specialisation indices to analyse knowledge dynamics and the development of local capabilities. Specialisations indices are used in this literature as a proxy for local capabilities, whereas patents are used as a proxy for innovations. We follow this implementation and further differentiate between dynamics occurring at the local (i.e., "AI-specific capabilities" versus "General capabilities") and technological (i.e., "AI-core technologies") levels. Next, we describe in detail the dataset and method.

3.3.1. Data collection and identification

We use PATSTAT 2019 (Spring version) to identify all priority filings⁴⁴ applied for (i.e., granted or not) in the period of analysis considered. The creation of this patent dataset rests on three main choices: i) strategy for identifying AI patents, ii) assignment of patents to countries, and iii) definition of the overall period of analysis and three intervals within it used to highlight technological change.

Following Leusin et al. (2020), we used a keyword-based search strategy for (i) identifying as AI-patents all patents mentioning at least one typical AI technique in their title or abstract. These AI techniques are advanced statistical and mathematical models used to implement AI functions such as computer vision, natural language processing, or knowledge representation. They include, for example, keywords related to machine learning (e.g., deep learning, neural networks, classification and regression trees), logic programming (e.g., expert systems, logic programming), and probabilistic reasoning. The selection of these AI-techniques is based on the classification presented in WIPO (2019), which we complement with synonyms collected from Wikipedia. The resulting search strategy has a total of 36 keywords (see Appendix A).

Following De Rassenfosse et al. (2019)⁴⁵, we assigned (ii) patents to countries by using the inventors' address. In contrast to using the patent assignees' address, this approach captures the locus of knowledge creation (Squicciarini et al., 2013). Based on this location proxy, we discerned that more than 92% of all AI patents in our dataset were invented in one of only four countries: China, Japan, South Korea, and the United States. Japan and the United States were the early leaders in AI development, applying for AI patents since as early as 1975. South Korea and China emerged in the second interval, with a particular spike in AI registrations seen for China in the third interval. See Appendix B for more information on the registration of AI patents by these countries over time.

⁴⁴ A "priority filing" is the first patent application filed to protect an invention. If the same patent is applied for in other patent offices, the following applications are called non-priority filings that are linked into a patent family through the priority filing. From now on, when we use the term "patents", we refer to priority filings.

⁴⁵ Note that we do not use the dataset presented in De Rassenfosse et al. (2019a), since it ends in 2014 and would omit the most recent period, for which we can document a substantial number of AI-related inventions. Inspired by De Rassenfosse et al. (2019a), who used additional sources to increase the location information of inventors and applicants related to their dataset, we performed a cross-check validation with our dataset. Taking the AI patents as a reference, we found that 23,983 of our 42,971 AI patents were applied for between 2015 and 2018. Of the remaining 18,988 patents, 15,355 are also found in the dataset presented by De Rassenfosse et al. (2019b). These patents hold the location information for 28,324 inventors, from which our dataset contains the same information for 27,275 inventors, which suggests at least a 96.3% correspondence.

Finally, for (iii) we select the term from 1974 to 2018 as the overall period of analysis, which we divide into three intervals as we consider priority patent applications for AI-related inventions⁴⁶. This results in 29,935,041 patents⁴⁷ considered in our patent dataset. Defining 15-year intervals allows us both to create sub-periods with the same length within the considered period (i.e., 1974-1988, 1989-2003, and 2004-2018), and to separate early AI adopters (Japan and the United States) from latecomers (South Korea & China). The latter developed their first AI patents in 1989, which coincides with the first year of our second interval. Each interval also represents particular technological AI breakthroughs: The rise of knowledge-based expert systems, which took place from 1980 to 1987 (WIPO, 2019), the fast development of machine learning during the 1990's (Li and Jiang, 2017), and the culmination of machine learning through deep-learning, as proposed in Hinton and Salakhutdinov (2006).

3.3.2. Method

We apply a technological space approach and specialisation indices to analyse the above-described patent dataset. The technological space approach, proposed in Hidalgo et al. (2007) and thereafter adapted for patents (Balland, 2016; Boschma, Balland, et al., 2014), displays technologies in a network visualisation according to their particular measures of relatedness. These measures are based on the co-occurrence of two technologies within the same patent. If two technologies co-occur more often than what would be expected by chance under the assumption of statistical independence, the assumption is that they are more related to each other.

The measure of relatedness, first conceptualised in Breschi et al. (2003), considers the fact that patent examiners assign one or more classification codes to each patent. A symmetrical matrix of co-occurrences “C” is calculated to account for every possible combination between two distinct technologies (from technology “k” = 1 until the “n” last technology considered in the classification scheme considered, for example). The resulting symmetric matrix “C”, which has dimensions “n” x “n”, is then normalised to avoid the overestimation of knowledge links involving technologies that are largely used. Breschi et al. (2003) propose the use of the cosine

⁴⁶ The exceptions are two AI patents applied for by the United States in the year 1961, which were followed by a 13-year global hiatus of AI-registrations; the first AI patent after this hiatus was applied for in 1975.

⁴⁷ This number includes 72 patents for which the applied method of geographical proxying could not identify the inventors' location.

index “S” for this normalisation, which is calculated in pairs considering every co-occurrence between two generic “i” and “j” technologies, as presented below:

$$S_{ij} = \frac{\sum_{k=1}^n C_{ik} C_{jk}}{\sqrt{\sum_{k=1}^n C_{ik}^2} \sqrt{\sum_{k=1}^n C_{jk}^2}}$$

Hidalgo et al. (2007) extend this idea by representing a normalised symmetric matrix of products exported by countries in a network perspective, in which nodes represent products and links represent their “co-occurrence” in each country’s export basket.

We follow the technological space approach proposed by Hidalgo et al. (2007), but use patents as the unit of analysis and the cosine index proposed in Breschi et al. (2003) to normalise knowledge relatedness. To differentiate among technologies, we use the 3rd version of the IPC technological field classification, which includes a total of 35 technological fields. This classification overcomes some inconsistencies of earlier classifications by i) considering all extant IPC codes, ii) balancing the size of the considered fields, and iii) reducing the overlap between similar technologies (Schmoch, 2008).

We build two distinct technological spaces: “Global” and “AI-specific”. The Global Technological Space (GTS) is static and considers all patents identified in the considered period. It is used to highlight the local development of technologies (i.e., both AI technologies and all other technologies the considered country patented) by the United States, Japan, South Korea, and China. The AI-specific Technological Space (ATS), in turn, includes only AI patents and doesn’t make any distinction regarding location. Conversely to the GTS, the ATS is dynamic, varying from one interval to the other according to the number of AI patents applied for in each interval, as well as their technical characteristics, and is used to highlight AI’s development. As we are interested in changes in innovation patterns, the patents applied for in one interval are not considered in the next interval (i.e., no stocks of patents are considered between intervals) for any of the calculations.

Moreover, we measure the specialisation of every considered entity over each interval through the Revealed Comparative Advantage (RCA) index, presented in Balassa (1965):

$RCA_{Technology\ t, entity\ e}$

$$= \frac{\frac{\text{Number of patents in technology } t \text{ from entity } e}{\text{Total number of patents of entity } e}}{\frac{\text{Total number of patents in technology } t \text{ in the larger economy } E}{\text{Total number of patents in the larger economy } E}}$$

If an entity has an RCA⁴⁸ equal to or higher than one, it has a specialisation, whereas values below this threshold show an absence of specialisation. We use the RCA both to highlight specialisations of entities over the considered technological spaces and as an independent indicator. For the latter, we measure the specialisation of countries in ten selected IPC subclasses (i.e., 4-digits IPC codes), differentiating between each country's "General" specialisations (i.e., based on all patents), and "AI-specific" specialisations (i.e., based only on the AI patents). For the former, we consider the same classification used in the technological spaces, i.e., IPC technological fields. In the GTS, the RCA is used to highlight a country's existing capabilities (i.e., "AI-specific" and "General" specialisations), whereas the ATS reflects technologies that are used above their particular relative average in AI patents (i.e., "AI-core technologies")⁴⁹. A summary of the main characteristics of both types of technological spaces considered, including how the RCA is used in each of them, is presented in Table 5.

⁴⁸ Technically, we use the Revealed Technological Advantage (RTA) index, which is an extension of the RCA index to technologies, but the principle is the same.

⁴⁹ The RCA compares the use of a technology by an entity in comparison to how this technology is used in a larger economy. In the particular case of defining AI-core technologies, the "entity" is the bulk of AI patents applied for in a given interval, and the larger economy are all patents applied for worldwide (i.e., not only AI). In this way, AI-core technologies reflect technological fields that are relatively (i.e., in relation to all other techn. fields used in AI patents) overused in AI patents in comparison to their relative use worldwide.

Type of technological space	Description and purpose	Data used	RCA's use
AI-specific technological space (ATS)	The ATS is divided into the three considered intervals, and presents a dynamic view of the relatedness between classification codes linked to AI innovations over each interval. It is used to analyse the technological evolution of AI and the main innovation patterns that emerged in each of the considered intervals.	All AI patents identified (i.e., not only the patents related to the four considered countries), separated according to the interval in which they were applied for.	Highlights the most important classification codes to AI innovations (i.e., "AI-core technologies") in a given interval.
Global technological space (GTS)	The GTS presents a static view of the relatedness between all classification codes used in patents for the whole considered period. It is used to analyse and compare the technological trajectories of the four leading countries considered.	All patents applied for (i.e., not only the patents related to the four considered countries) in the considered period.	Highlights specialisations (i.e., General or AI-specific) of each country over each considered interval.

Table 5: Types of technological space considered.

The complete R implementation of the method, including the creation of all technological spaces, is available in our public GitHub repository⁵⁰.

3.4. Empirical analysis

3.4.1. Identifying technologies linked to AI in the ATS

We first look into how AI-related innovations evolved over the considered intervals. We aim at identifying the main structure of technological knowledge used in AI innovations in each considered interval, which shows the main technologies (i.e., "AI-core technologies") used to create such innovations.

Accordingly, in Figure 10, we present the technologies linked to AI over the three considered intervals. We use colours and node formats to highlight distinct technological sectors, and node size to emphasise nodes' connectivity (regarding the number of links) to other

⁵⁰ To avoid revealing the identity of the authors at this stage, the link is not shown here yet and the GitHub repository status is currently kept as private. The link and the change in status will be presented at the proper time before the publication of the paper.

technologies⁵¹. Specialisations, i.e., “AI-core technologies”, are highlighted by depicting the name of the technological field which has a specialisation in the respective interval.

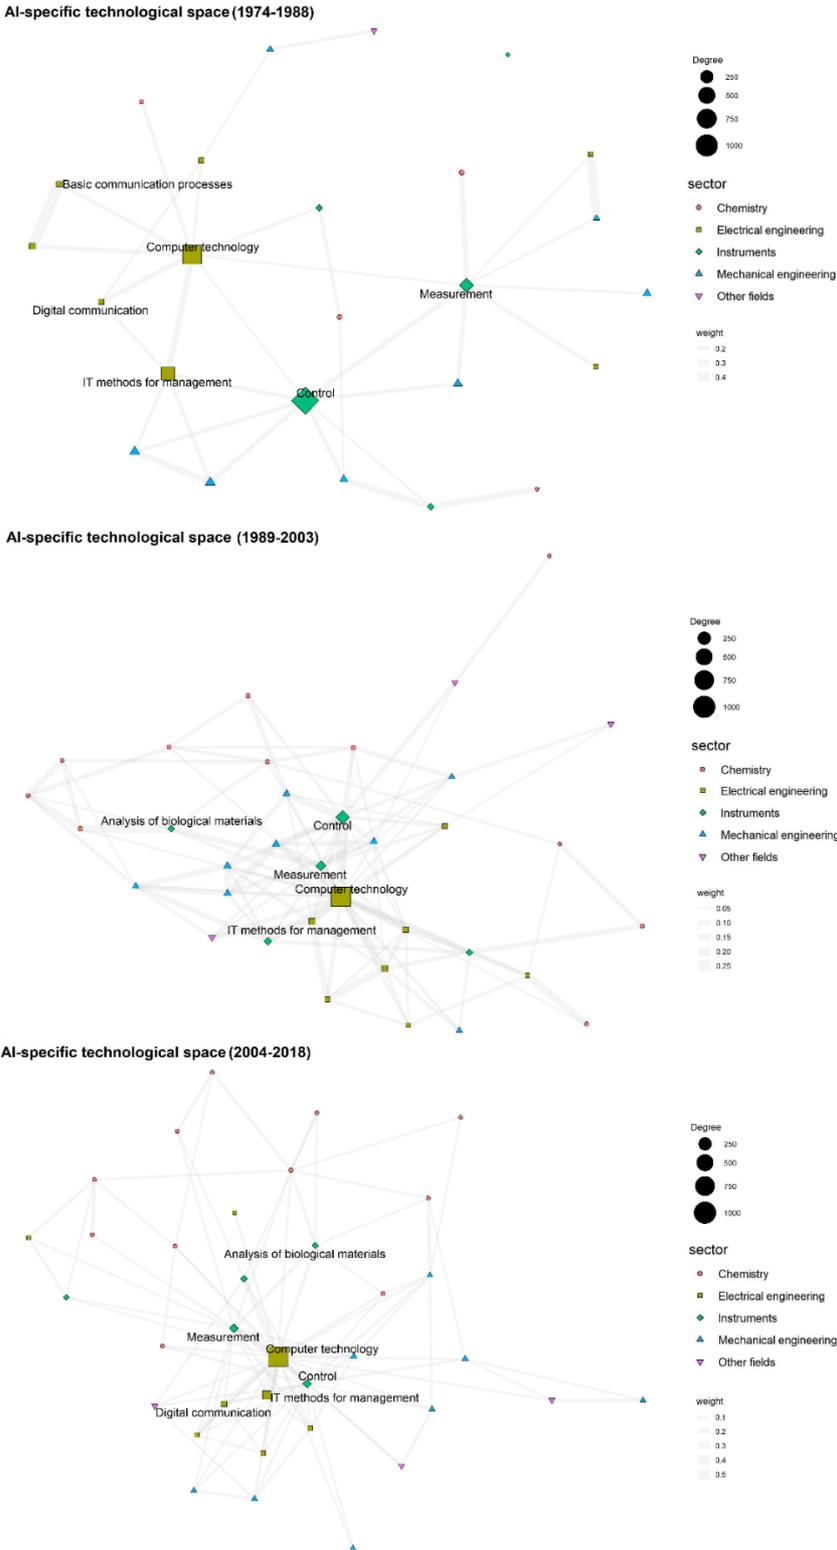


Figure 10: ATS and AI-core technologies over the considered intervals and IPC fields.

⁵¹ Depicted in the label “Degree”, which stands for Degree of Connectivity (i.e., higher values mean more connected technological fields), whereas the label “weight” stands for the weight of the links (i.e., higher values mean more connections between the two technologies related to the link).

AI started as a combination of technologies sparsely distributed in the network, which then evolved to become more densely connected (see Figure 10). The number of technological fields linked to AI innovations is 25 in the first interval, and reaches the maximum of 35 fields in the second and third intervals. Specialisations, i.e., the “AI-core technologies”, are seen through the whole period exclusively in the sectors of “Electrical engineering” and “Instruments”. These specialisations do change over time, although just to a certain extent: Three out of the seven identified “AI-core technologies” present some change. The less connected “Basic communication processes” and “Digital communication” lose their relevance after the first interval, although the latter presents again a specialisation in the third interval. The field of “Analysis of biological materials”, in turn, is highlighted with a specialisation only in the second and third intervals. The “stable” AI-core technologies, defined as those which show a specialisation in every interval, are “Computer technology”, “Measurement”, “Control” and “IT methods for management”. In a nutshell, these results show how, and how much AI technologies changed over time, and the effects of these changes on the relatedness between technologies used to create AI innovations. The persistence of technological change and its effects on technological relatedness are one of the arguments of this chapter.

Overall, the technological fields of “Control” and “Computer technology” are almost equally essential for AI innovations in the first interval. After this, the latter becomes increasingly important and central in the AI network. In the third interval, “Computer technology” is indisputably the most connected technological field for AI innovations. The fields of “Control” and “Measurement” also play important bridging roles. “Control” connects AI to technologies related to the sector of “Mechanical engineering”, while “Measurement” connects AI with technologies from the “Chemistry” sector.

3.4.2. General and AI-specific specialisations in the GTS

Next, we focus on all patents applied for in the considered period (1974 – 2018), which are used subsequently to analyse the local exploration of AI. This global technology space is presented in Figure 11.

Global technological space: IPC Technological fields

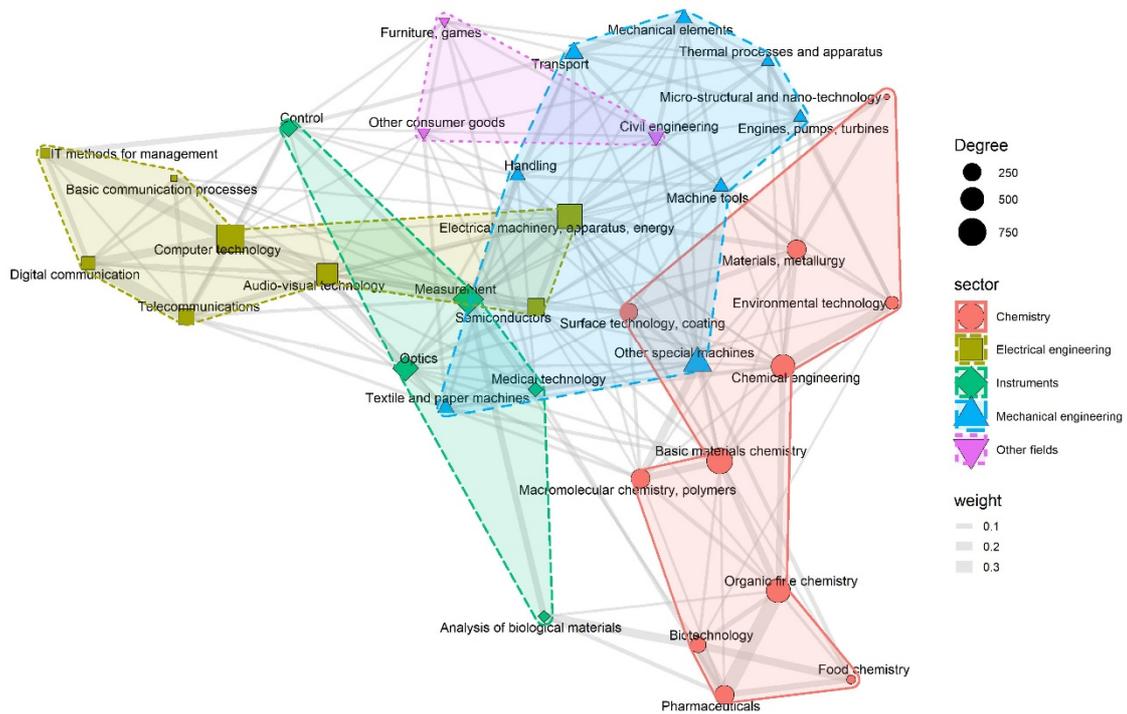


Figure 11: GTS for IPC fields.

One can see in the GTS that technological fields from the same sector are often close to each other, highlighting the expected stylised fact that the similarity between technologies within the same sector is higher than the similarity across distinct sectors. “Electrical engineering”-related fields are mostly placed jointly on the left side of the network, while “Mechanical engineering” and “Chemistry”-related ones are placed on the top and right of the figure respectively (see Figure 11). One can also see that AI-core technologies identified previously (see Figure 10) are placed close to each other in a cluster on the left side of the network (the only exception being the field of “Analysis of biological materials”, placed at the bottom). The field of “Computer technology”, which was the most connected in the AI network in the third interval, is also the most connected when all patents are considered in the Global perspective (see Appendix C for a complete list of the fields’ connectivity).

We use this GTS to highlight the technological trajectory of countries leading AI development. We identify local technological development by outlining the specialisations of Japan, the United States, South Korea, and China in the GTS over the three considered intervals. Thereby, we differentiate between two main kinds of specialisations: “General specialisations” refers

to the performance of a country considering all patents applied for⁵² in each considered interval; “AI-specific specialisations”, in turn, focusses on AI patents, measuring the performance of each country according to its registrations of AI patents. We use these two kinds of specialisations to create four labels capturing local exploration of AI, summarised in Figure 12. We use these labels to highlight country capabilities developed over the considered intervals, presented in Figure 4. Furthermore, we use now the fields’ names to highlight the previously identified technologies at each interval (i.e., in Figure 1) as AI-specialised (i.e., “AI-core technologies”). Tabular data for each of these specialisations is presented in Appendix D.

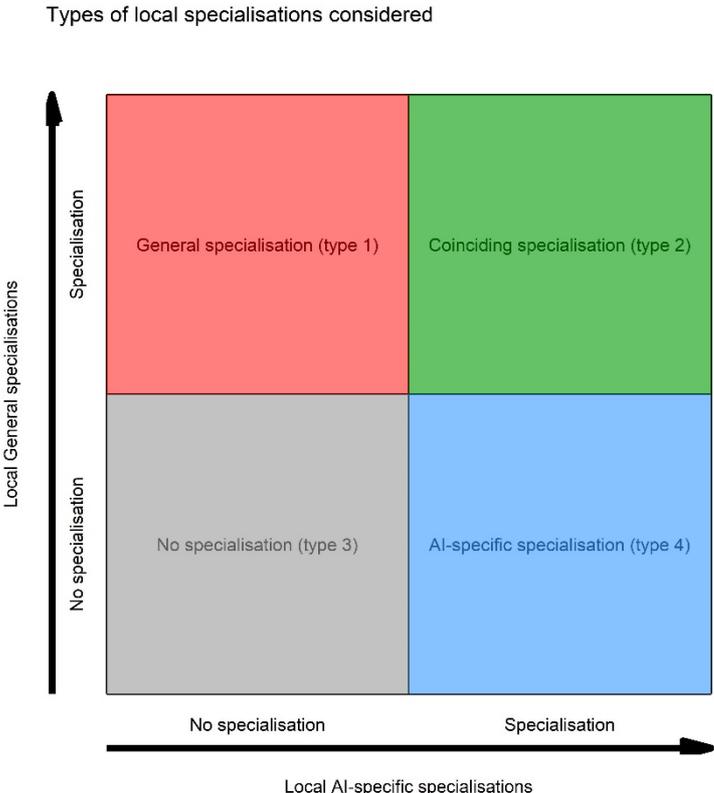
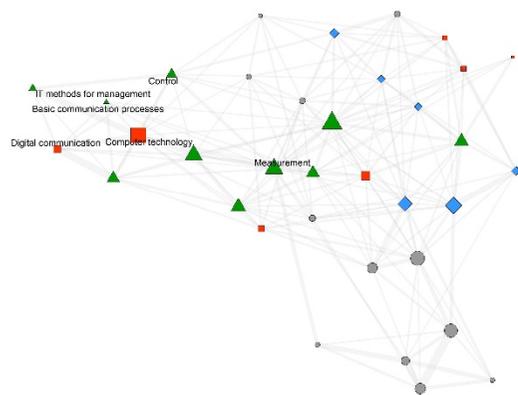


Figure 12: Four types of specialisations considered.

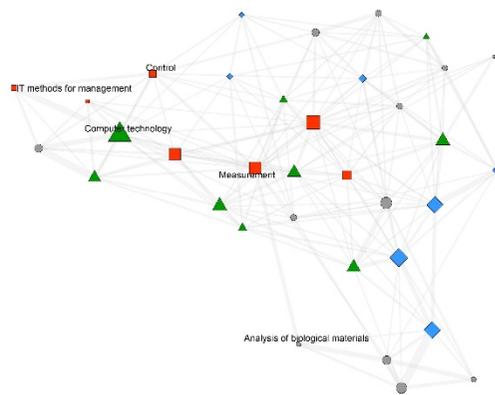
⁵² The calculation for both kinds of specialisations considers the patents applied for by all countries, not only the four countries in which the analysis is focused on.

a) Japan

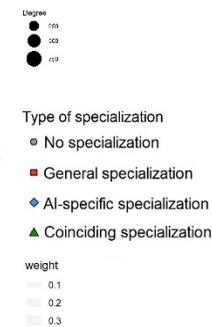
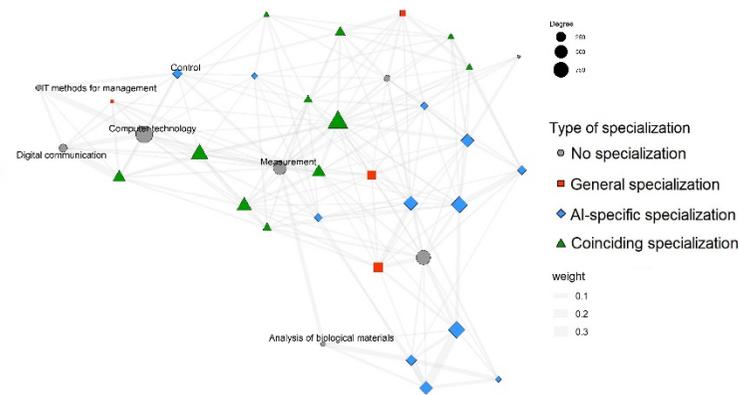
Global technological space: Japan (1974-1988)



Global technological space: Japan (1989-2003)

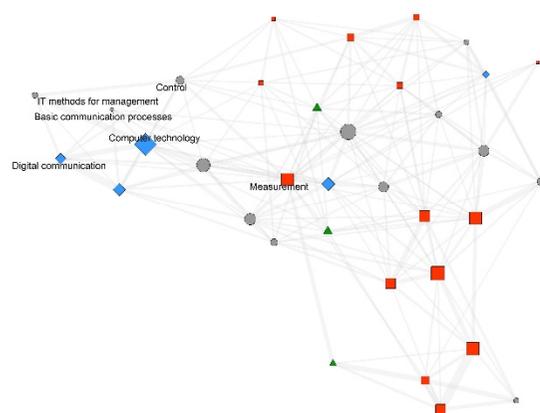


Global technological space: Japan (2004-2018)

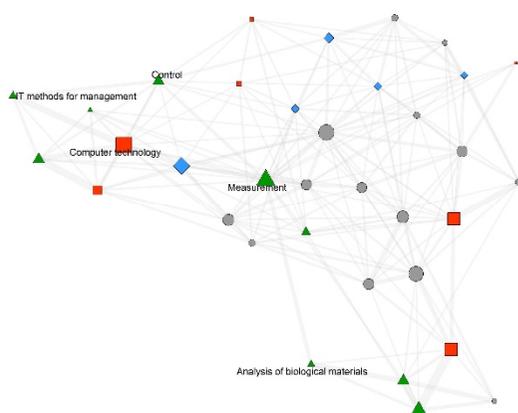


b) United States

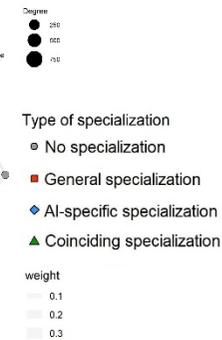
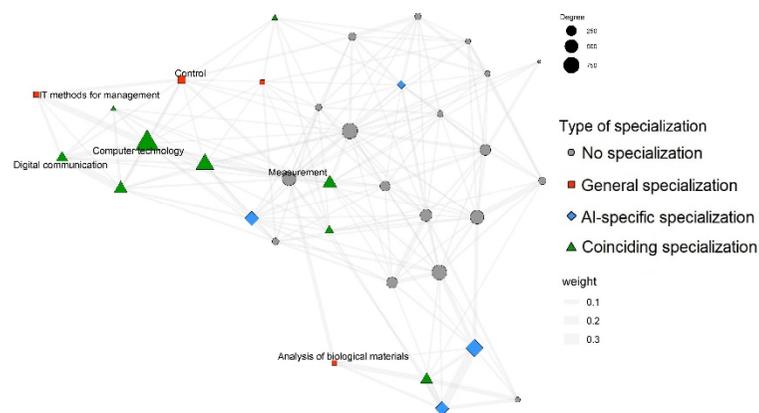
Global technological space: USA (1974-1988)



Global technological space: USA (1989-2003)

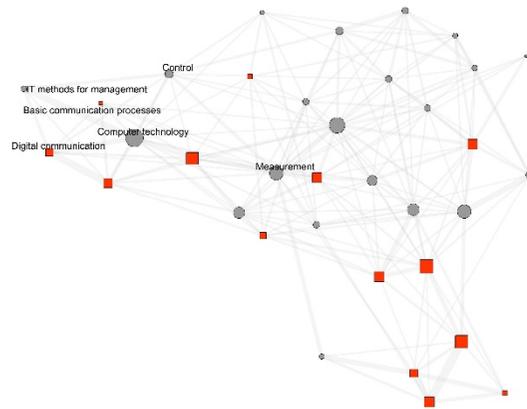


Global technological space: USA (2004-2018)

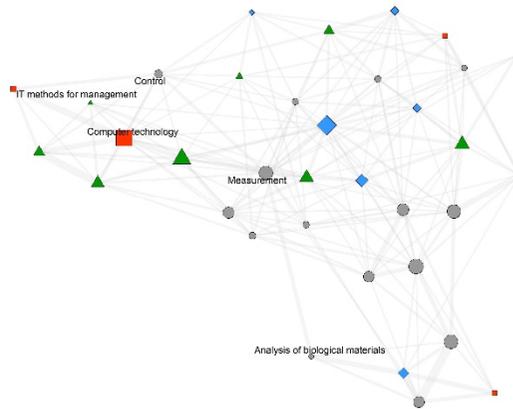


c) South Korea

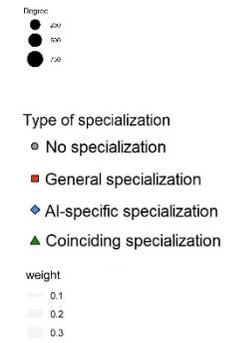
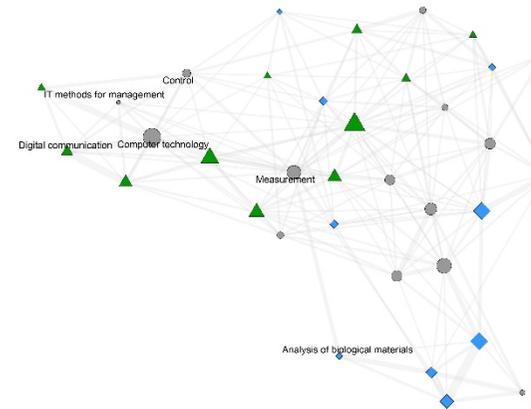
Global technological space: South Korea (1974-1988)



Global technological space: South Korea (1989-2003)

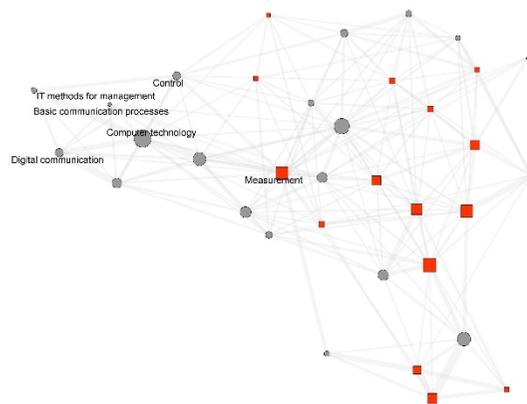


Global technological space: South Korea (2004-2018)

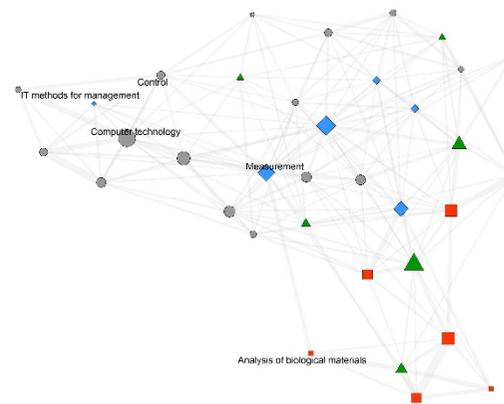


d) China

Global technological space: China (1974-1988)



Global technological space: China (1989-2003)



Global technological space: China (2004-2018)

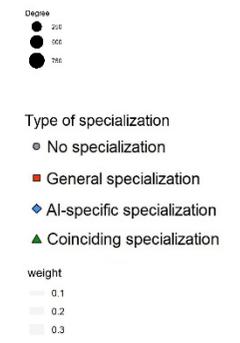
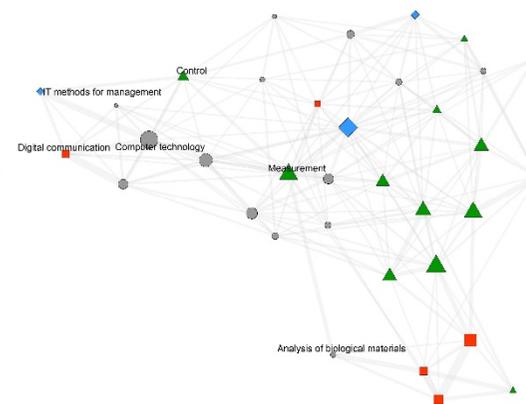


Figure 13: Specialisations in the considered technological fields for Japan (13a), United States (13b), South Korea (13c), and China (13d).

We detect a trend in Figure 13, which shows that the share of coinciding specialisations (type 2) steadily increases over time in comparison to the number of general specialisations (type 1). This indicates that countries increasingly combine their General and AI-specific specialisations in the same technological fields. As a result, general-only specialisations become less common, whereas coinciding specialisations increase (red squares are overtaken by green triangles). Figure 14 quantifies the mentioned trend by displaying the share of coinciding specialisations in relation to the total number of general specialisations⁵³.

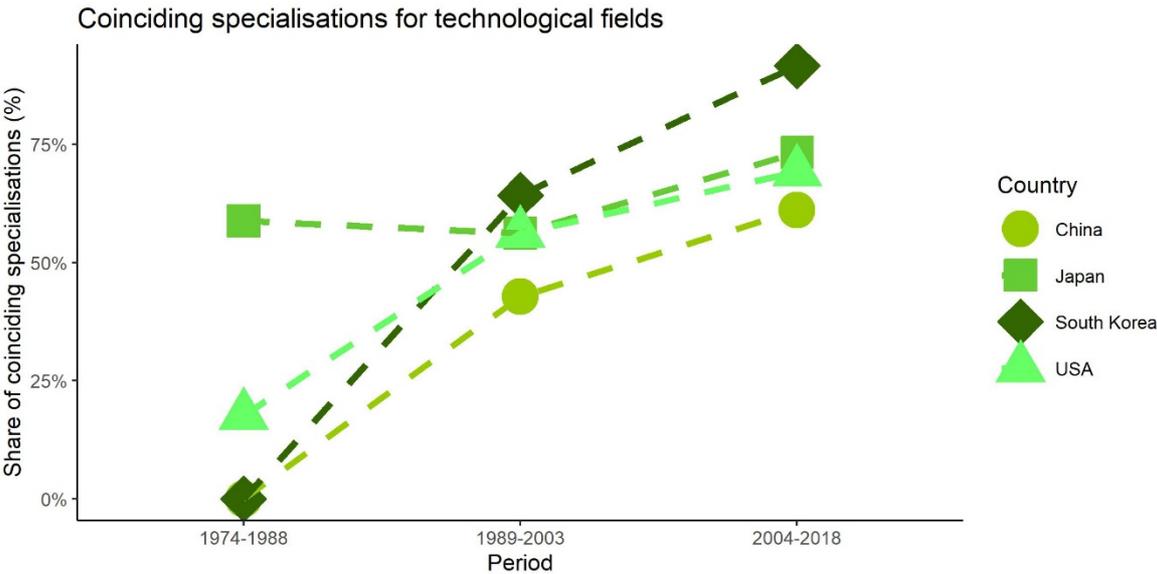


Figure 14: Share between coinciding and general specialisations for the four considered countries at the technological field level.

Surprisingly, specific trends related to “AI-core technologies” do not seem to affect how countries develop AI-specific specialisations (see Figure 13). The latecomers South Korea and China develop both an AI-specific specialisation in the technological field “Basic communication processes” when they start exploring AI in the second interval, although this field is not particularly relevant to AI anymore. Notably, China does not develop any kind of specialisation (i.e., neither General nor AI-specific) in the field of “Computer technology” over the whole considered period, despite its importance as an AI-core technology. The country instead is the only one holding a coinciding specialisation in the bridging field of “Measurement” in the third interval, which might be better suited for China given its particular capabilities in fields related to the “Chemistry” sector. South Korea, in comparison, starts exploring AI by developing a General specialisation in the “Computer technology” field, which

⁵³ Calculated as: $Share\ of\ coinciding\ specialisations = \frac{Number\ of\ coinciding\ spec.}{Number\ of\ coinciding\ speci.+Number\ of\ general\ spec.}$

happens to be close to fields in which the country already had General specialisations in the first interval. Conversely to China, South Korea's coinciding specialisations remain close to the AI cluster on the left side of the network. The early comers Japan and the United States also present different specialisation patterns over time. Japan starts with some type of specialisation in every field related to AI, but loses both General and AI-specific specialisations in "Computer technology" and other fields around it in the third interval. The USA instead starts exploring the AI-side of the network mainly through AI-specific specialisations and doesn't lose its "Computer technology" General specialisation over time.

3.4.3. General and AI-specific specialisations in technological classes at the 4-digits IPC level

We interpret our results based on the RCA cautiously, since this index compares the relative share of patents produced by a country to a relative global average. This implies that countries with few patents may present specialisations due to a low total number of patents (see also Soete (1987)). None of the considered countries has a low patent output. However, few AI patents may have been applied for in some technological fields, so a country with a small number of AI patents in these fields might be identified as having AI-specific specialisations in them without necessarily having a relevant development.

Therefore, we refine our analysis by considering only the technologies more often used in AI patents in absolute numbers. In this way, we minimise the risk of having AI-specific specialisations linked to a country just because the number of AI patents registered in a technological category is too low. As we have a limited number of fields strongly linked to AI at the technological field level, we broaden our analysis of specialisation indices to the 4-digit IPC level (subclasses). These subclass codes are considerably more specific than the technological field codes, which allows us to better detail the technologies most used in AI patents. There are 645 subclasses available in the considered 2019 scheme⁵⁴, from which 456 are referred to at least once in one of our identified AI patents. We avoid the network analysis at this stage (which is reported separately in Appendix E), and focus instead on the ten IPC subclasses most used in AI patents. These ten codes concentrate 77% of the subclasses used in the identified AI patents. The majority of them are related to the technological field of "Computer technology" (namely the subclasses "Electric digital data processing" (G06F),

⁵⁴ <https://www.wipo.int/classifications/ipc/en/ITsupport/Version20190101/transformations/stats.html>.

“Recognition and presentation of data” (G06K), “Computer Systems” (G06N), “Image data processing or generation” (G06T), and “Speech analysis or processing” (G10L)⁵⁵.

We focus on the same four countries, and differentiate once again between General (see upper panel, Figure 15a) and AI-specific specialisations (see lower panel, Figure 15b). The main focus now is to disentangle if AI-specific specialisations coincide with General specialisations for the technologies in which there is a high concentration of AI patents. Moreover, this time we use a non-binary RCA index⁵⁶. By doing so, we can see in detail the small variations in countries’ specialisations that precede the consolidation or not of a particular specialisation. Additionally, we use the Log 10 values of the specialisation indexes for better visualisation. In this way, a specialisation is present for values above 0 and absent for values below this threshold.

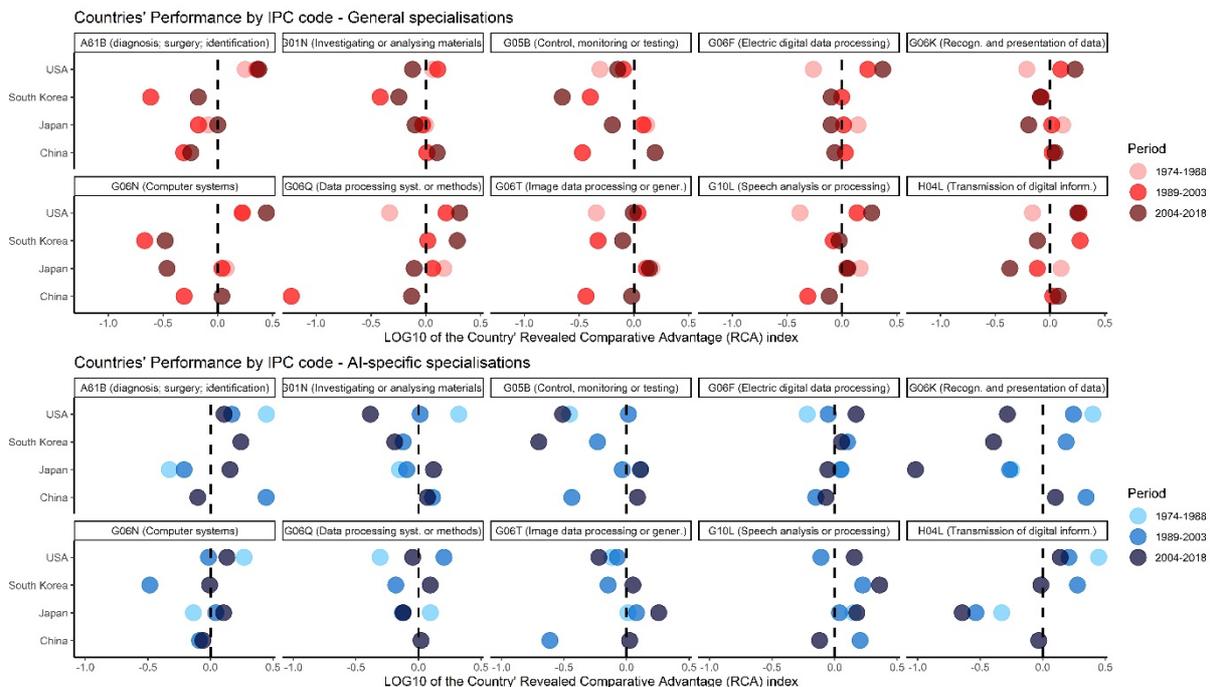


Figure 15: General (15a) and AI-specific specialisations (15b) across patent subclasses for Japan, the United States, South Korea, and China.

We see again that a country’s AI-specific specialisations often appear in codes in which they hold a General specialisation (see Figure 15a). Conversely, we see that a specialisation in AI rarely precedes a general specialisation in the same subclass. Where this is the case, it does not last. For example, China loses very rapidly the A61B and G10L AI-specific specialisations it

⁵⁵ The remaining five subclasses are from the technological fields of: “Measurement” (G01N), “Control” (G05B), “IT methods for management” (G06Q), “Digital communication” (H04L), and “Medical technology” (A61B).

⁵⁶ Meaning that it can take values between 0 and 1, and also values beyond 1.

acquired in the second interval, before these subclasses were a part of the country's General specialisations. The same holds for the United States in the codes G05B, G06K, and H04L, as well as South Korea for G06K. The only exception to this is South Korea's sustained AI-specific specialisation in the subclass "Speech analysis or processing" (code G10L). We also find that despite Japan losing several of its specialisations in AI-core technologies in the third interval in the network perspective (see Figure 13a), the country maintains specialisation advantages in main AI subclasses (e.g., leading in the subclasses G05B and G06T as well as relevant in G10L, A61B, G01N, and G06N).

Furthermore, there are specialisation patterns in the considered codes that are not entirely captured by a binary consideration of RCAs as used in the prior analysis. For example, all countries move towards the subclass "Computer systems" (code G06N) in the AI-specific perspective. The USA, an early AI adopter, is the only country that reduced its level of specialisation in this subclass in the second interval, precisely when latecomers start developing AI. Japan and the United States are the only countries that do develop a specialisation in this subclass, but China and South Korea advance steadily towards it. Conversely, a decrease in the level of AI-specific specialisation is also seen for the subclass "Recognition and presentation of data" (code G06K). Japan and the United States move away from this specialisation in later intervals, followed by South Korea and China.

Having confirmed that general and AI-specialisations often coincide for the ten most used AI technologies, we measure the share of coinciding specialisations at the subclass level. We do so for the ten considered subclasses (see Figure 16a), and then for all subclasses available (see Figure 16b). The trend is very similar to what was seen previously at the technological field level: Over time, codes with a general specialisation are increasingly likely to show also an AI-specific specialisation. For the ten considered subclasses, leading countries lose some of their coinciding specialisations when latecomers start exploring AI. Countries' coinciding specialisations in all subclasses reach roughly a third of the scale seen previously at the technological field level in the third interval⁵⁷.

⁵⁷ Note that some staggering shares appear as a result of the limited number taken as reference in Figure 16a. These are shown for Japan and South Korea in the third interval, which appear with a 100% share of coinciding specialisations. These high shares are due to the limited number of general specialisations that these countries hold in this interval in the considered ten subclasses (Japan holds a general specialisation in the codes G06T and G10L, whereas South Korea does it in the code G06Q).

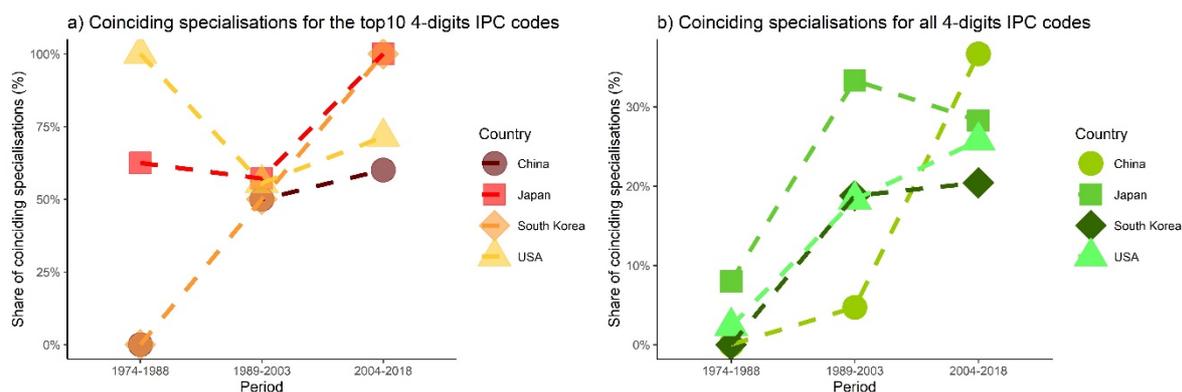


Figure 16: Share between coinciding and general specialisations for the four considered countries at the subclass level for ten codes (16a) and all codes (16b).

3.5. Discussion and concluding remarks

3.5.1. Summary of the main findings

This chapter analysed the technological evolution of AI and compared it to its local exploration by countries leading AI's development. The purpose was to understand how AI innovation patterns changed over time, whether the changing innovation patterns were reflected on how leading countries explored this technology, and how these countries incorporated AI's technological trajectories into their global technological trajectory. Using a technological space framework to visualise the changes in relatedness between technologies linked to AI innovations, we identified a considerable transformation over time. AI evolved from a dispersed network of weakly related technologies to a dense and larger network centred around the field of "Computer technology". In this process, some AI-core technologies initially linked to AI lost their importance (e.g., "Basic communication processes"), whereas others gained relevance (e.g., "Analysis of biological materials").

Our results indicate that these changes in AI-core technologies were not uniformly reflected in their local exploration by countries leading AI development. Instead, the existing local general knowledge was seemingly more correlated with the development of AI-specific capabilities. Interestingly, this local knowledge did not hinder countries from developing AI-specific specialisations in fields where there was no previous related knowledge, but it acted to "preserve" the specialisations that were linked to existing general knowledge. As a result, AI-specific specialisations developed in technologies that were not part of the local existing knowledge bases were very likely to vanish. This "selection" process was shown to lead to the emergence of a pattern where new local specialisations developed in AI increasingly coincide

with existing “general” specialisations. Hence, AI’s technological trajectory and local general technological trajectories increasingly converge.

3.5.2. Contributions

Our findings add to the scarce literature which recognises that technological relatedness is being created dynamically. In particular, we complement Juhász et al. (2021) by showing that geographically “unbounded” dynamics also contribute to shaping technological relatedness. This follows the mechanism proposed in Arthur (2009). in which successful innovations that advance technological development are repeated. The repetition of these successful innovations creates innovation patterns, which reflect the technical routes taken to solve the problems typical of a technology (Dosi and Grazzi, 2006). The repetition of these technical routes, in turn, leads to an increase in the relatedness between the technologies linked to them.

Combinations that significantly improve a technology’s performance are expected to be adopted and repeated in local innovations. The fact that this global technological progress is weakly reflected on local technological exploration offers an explanation for why some once-promising and successful local industries eventually decline and fail. The stronger role of path dependency over global technological development may lead such industries to get stuck in local versions of a technology. In this case, if a relevant technological development occurs abroad and is not successfully translated to the local context, existing industries may become less competitive and decline. This finding allows linking the idea of technological lock-in (Arthur, 1989; Cowan, 1990) to the EEG literature, offering an explanation to why firms and industries leave specific locations (Boschma, Balland, et al., 2014; Rigby, 2015). In particular, we documented that these dynamic innovation patterns occur over time. These dynamics are even greater when the evolution of all technologies is considered simultaneously.

Our findings also show that leveraging local knowledge is likely to open new, unrelated technological opportunities. While counterintuitive in the sense that this kind of leverage is expected to create only related technological variety, we show that the recombinatorial process allows new technologies to break-in locally. This happens through the recombination of a novel technology with the existing local knowledge base, which may lead to the creation of a new technological trajectory. This is a form of speciation (Adner and Levinthal, 2002) that occurs locally, in which a new technological variety emerges as the result of a new

combination made with existing knowledge. This new technological variety may diffuse locally, leading to industrial branching (Frenken and Boschma, 2007) which, in our analysis, was shown to converge over time towards a more coherent general technological trajectory. This “coherence” was shown to be correlated with the destruction of the new capabilities (and possibly their associated emerging industries) that diverge from existing ones.

Regarding methodology, we demonstrate a way of “disentangling” local innovations made with a particular technology from “global” technological innovations. This was done by analysing these innovations separately. Global technological innovations refer to all patents applied for worldwide for a given technology, and reflect unbounded innovation patterns. Local innovations refer to patent registrations made locally (i.e., in this particular case, location information is used as a proxy). By using the RCA index to highlight “global” and local specialisation patterns over technological spaces, one can compare “unbounded” technological development to local development. Although based on a simple idea and implementation, we find no similar approach in the literature. This disentanglement allows the measurement of distinct kinds of specialisations (i.e., technology-specific and general ones), which is crucial to understand how new technologies are incorporated into local knowledge bases. In this regard, we think our methodological approach is more intuitive than the ones considered so far, for example in Buarque et al. (2020)⁵⁸.

Regarding the AI literature specifically, our investigation contributes to the identification of the main technologies related to AI innovations in distinct intervals of time. We identified seven core technologies linked to AI, four of which remained stable for the last four decades. Interestingly, not all of them are related to the computing aspect of AI. Therefore, in contrast to earlier studies (Buarque et al., 2020; Fujii and Managi, 2018; Klinger et al., 2018; Righi et al., 2020), we manage to capture a more dynamic development of AI, including the use of distinct main technologies over time. The consideration of this distinct set of technologies contrasts with findings presented in Buarque et al. (2020) and Fujii and Managi (2018). Regarding the latter, the authors identify AI development by considering only patent registrations related to the “G06N” patent class, which refers to “computer systems based on

⁵⁸ Put simply, the authors examine how AI is integrated into the knowledge spaces of regions by looking at how these networks change when AI patents are excluded from them. In our opinion, this kind of exclusion is problematic. For one, it takes the assumption that technological efforts made towards developing AI wouldn't be used to develop other technologies in the case that AI wouldn't exist. Besides, a bias may be created towards regions with an overall low number of patents, making AI appear to be more important than it actually may be.

specific computational models”. Our approach suggests that this identification may be too restrictive and potentially misleading, since countries can become global leaders in AI innovations without developing an AI specialisation in computer-related fields. Similarly, Buarque et al. (2020) point out that local computing-related capabilities are potentially a necessary condition to develop AI. Our findings, stressed most clearly by the case of China, highlight that this condition does not seem to apply to the national development of AI. The existing specialisations of China when it started exploring AI stress the absence of general capabilities in technological fields apparently central to AI development, like “Computer technology”, “Measurement” and “Control”. Nevertheless, China managed to develop specialisations in AI through its existing capabilities in fields mostly related to the “Chemistry” sector.

3.5.3. Policy implications

Our empirical findings stress the benefits of bottom-up policies aimed at inducing technological development. Analysing the case of AI, we found evidence that technological specialisations developed in fields where there is also an existing general specialisation are likely to last longer than technology-specific specialisations developed in fields disconnected from local knowledge. By presenting evidence that the creation of capabilities based on local knowledge is more effective to generate long-lasting capabilities, we argue towards the leveraging of local capabilities as a way to induce further technological development. The concept of Smart specialisation, in particular, has been recently stressed in the EEG literature as particularly appropriate to that aim (Balland et al., 2019; Hidalgo, 2021; Montresor and Quattraro, 2020; Uyarra et al., 2020; Whittle, 2020).

Although looking particularly at AI innovations, we believe that the technological patterns of how countries incorporated AI reflect a broadly generalisable characteristic of technological development, in which technologies are recombined and incorporated in each other (Arthur, 2009). This is an important factor for creating technological commonalities, as highlighted in our illustrative case of computers and cell phones. We found this pattern occurring even in the two intervals before the recent breakthrough of AI, in which it crossed the barrier from research to the application domain by being applied in a variety of technological domains (WIPO, 2019). Conversely, if the identified incorporation pattern does not hold for all technologies, it should hold at least for technologies that have a high potential for

recombination, like modular and digital technologies (Yoo et al., 2012), enabling technologies (Teece, 2018), General Purpose Technologies (Bresnahan and Trajtenberg, 1995), or emerging technologies (Adner and Levinthal, 2002).

Nevertheless, the smart specialisation concept alone is not a silver bullet to innovation policy. Existing literature stresses the risks of unintended lock-in effects generated by self-reinforcing local co-evolution (see for example Hoppmann (2021) for an example of this risk associated with the co-evolution of science and industry). In our case, this could happen through local technological path dependencies that become so strong that they hinder local industries from adopting more efficient technological varieties. Thereby, smart policies would worsen the problem if aimed at reinforcing the existing capabilities that are too dependent on a less efficient technology. The concept of smart policy is a step towards the creation of more tailored bottom-up innovation policies – a characteristic increasingly recognised as relevant in recent literature (Grashof, 2021; Magro and Wilson, 2019) – but it does not replace the need for carefully balancing a mix of distinct policies (Del Rio, 2014; Magro and Wilson, 2019). A possible key point for its implementation may lay in the changing role of relatedness, which allows capturing technological evolution. Top-down policies may be relevant in a well-tailored policy mix as a tool to break or weaken harmful local path dependencies and, in this sense, relatedness may provide a useful tool to identify detrimental trajectories. The recognition of the dynamic nature of technological evolution should induce these policies to be used more as steering wheels toward technological development than as an end in themselves.

Finally, our findings also have implications for policies being created globally towards AI development in the so-called “AI global race” (Klinger et al., 2018). We highlight the multiple “entry points” possible for AI development, which should impact policy-making towards leveraging the deployment of AI in combination with existing local capabilities. This strategy seems appropriate even for technologies not very related to AI, which was notably seen in the case of China through its exploration of AI in technologies related to the Chemistry sector. Therefore, developing capabilities in the digital software-learning part of AI is not the only way to create AI innovations. Due to its modularity and transversality, AI is likely to lead to the emergence of very distinct technological trajectories across the globe, with some locations specialising in the development of software while others become specialised in combining AI software with physical products.

3.5.4. Limitations and future research directions

Finally, we need to acknowledge the limitations of our approach, of which the most explicit ones are linked to the data. Firstly, we identify AI innovations by just considering patents, although many innovations in this field are not patented. This might constitute a particular issue for AI innovations based on software and even more to open source software development. Hence, in a strict sense, our results refer to AI “inventions” rather than “innovations”. Secondly, we consider both granted and non-granted patents, which possibly introduce a quality bias on our dataset. Thirdly, we develop a keyword-based search to identify AI, which we prefer in our case over strategies based on classification codes. However, such an approach is inherently subjective and sensitive to the choice of keywords. We also kept the “global technological space” static, conversely to what we considered in the AI technological space. This was done to simplify the analysis and justified by our main focus on the technological dynamics of AI and how they interact with local knowledge.

We also highlight that the adopted methods do not allow conclusions regarding causality. In particular, this applies to the emergence and failure of the different types of specialisations identified. We cannot argue that the short-lived aspect of AI specialisations that did not match the local knowledge base was caused by this mismatch, although we believe that future research would support this claim. These aspects are subject to further inspection, which would need to apply causal inference strategies to test, whether the indicative findings revealed in this study are robust and beyond reasonable econometric doubt. This might also require a larger set of countries under investigation during the observation period. Our study focused on the four countries that account for the lion’s share of AI patents during the observation. Obviously, this choice limits our ability to generalise our findings. The exclusive focus on AI further hinders this generalisation aspect.

References

- Aarstad, J., Kvitastein, O. A., & Jakobsen, S.-E. (2016). Related and unrelated variety as regional drivers of enterprise productivity and innovation: A multilevel study. *Research policy*, 45(4), 844-856. doi: <https://doi.org/10.1016/j.respol.2016.01.013>
- Adner, R., & Levinthal, D. A. (2002). The emergence of emerging technologies. *California Management Review*, 45(1), 50-66. doi: <https://doi.org/10.2307/41166153>
- Antonelli, C., Krafft, J., & Quatraro, F. (2010). Recombinant knowledge and growth: The case of ICTs. *Structural Change and Economic Dynamics*, 21(1), 50-69. doi: <https://doi.org/10.1016/j.strueco.2009.12.001>
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The economic journal*, 99(394), 116-131. doi: <https://doi.org/10.2307/2234208>
- Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. United States of America: Simon and Schuster.
- Balassa, B. (1965). Trade liberalisation and “revealed” comparative advantage. *The manchester school*, 33(2), 99-123. doi: <https://doi.org/10.1111/j.1467-9957.1965.tb00050.x>
- Balland, P. A. (2016). Relatedness and the geography of innovation. In *Handbook on the geographies of innovation*: Edward Elgar Publishing.
- Balland, P. A., Boschma, R., Crespo, J., & Rigby, D. L. (2019). Smart specialization policy in the European Union: relatedness, knowledge complexity and regional diversification. *Regional studies*, 53(9), 1252-1268. doi: <https://doi.org/10.1080/00343404.2018.1437900>
- Boschma, R., Balland, P. A., & Kogler, D. F. (2014). Relatedness and technological change in cities: the rise and fall of technological knowledge in US metropolitan areas from 1981 to 2010. *Industrial and corporate change*, 24(1), 223-250. doi: <https://doi.org/10.1093/icc/dtu012>
- Boschma, R., Heimeriks, G., & Balland, P.-A. (2014). Scientific knowledge dynamics and relatedness in biotech cities. *Research policy*, 43(1), 107-114. doi: <http://dx.doi.org/10.1016/j.respol.2013.07.009>
- Boschma, R., & Martin, R. (2007). Constructing an evolutionary economic geography. In: Oxford University Press.
- Breschi, S., Lissoni, F., & Malerba, F. (2003). Knowledge-relatedness in firm technological diversification. *Research policy*, 32(1), 69-87. doi: [https://doi.org/10.1016/S0048-7333\(02\)00004-5](https://doi.org/10.1016/S0048-7333(02)00004-5)
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies ‘Engines of growth’? *Journal of econometrics*, 65(1), 83-108. doi: [https://doi.org/10.1016/0304-4076\(94\)01598-T](https://doi.org/10.1016/0304-4076(94)01598-T)
- Breznitz, D. (2007). Industrial R&D as a national policy: Horizontal technology policies and industry-state co-evolution in the growth of the Israeli software industry. *Research policy*, 36(9), 1465-1482. doi: <https://doi.org/10.1016/j.respol.2007.06.006>
- Buarque, B. S., Davies, R. B., Hynes, R. M., & Kogler, D. F. (2020). OK Computer: the creation and integration of AI in Europe. *Cambridge Journal of Regions, Economy and Society*, 13(1), 175-192. doi: <https://doi.org/10.1093/cjres/rsz023>
- Castaldi, C., Frenken, K., & Los, B. (2015). Related variety, unrelated variety and technological breakthroughs: an analysis of US state-level patenting. *Regional studies*, 49(5), 767-781. doi: <https://doi.org/10.1080/00343404.2014.940305>
- Chen, L., & Naughton, B. (2016). An institutionalized policy-making mechanism: China’s return to techno-industrial policy. *Research policy*, 45(10), 2138-2152. doi: <https://doi.org/10.1016/j.respol.2016.09.014>
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative science quarterly*, 128-152. doi: <https://doi.org/10.2307/2393553>
- Colombelli, A., Krafft, J., & Quatraro, F. (2014). The emergence of new technology-based sectors in European regions: a proximity-based analysis of nanotechnology. *Research policy*, 43(10), 1681-1696. doi: <https://doi.org/10.1016/j.respol.2014.07.008>

- Cowan, R. (1990). Nuclear power reactors: a study in technological lock-in. *The journal of economic history*, 50(3), 541-567. doi: <https://doi.org/10.1017/S0022050700037153>
- De Rassenfosse, G., Kozak, J., & Seliger, F. (2019). Geocoding of worldwide patent data. *Scientific data*, 6(1), 1-15. doi: <https://doi.org/10.1038/s41597-019-0264-6>
- Del Rio, P. (2014). On evaluating success in complex policy mixes: the case of renewable energy support schemes. *Policy Sciences*, 47(3), 267-287. doi: <https://doi.org/10.1007/s11077-013-9189-7>
- Dosi, G., & Grazzi, M. (2006). Technologies as problem-solving procedures and technologies as input-output relations: some perspectives on the theory of production. *Industrial and corporate change*, 15(1), 173-202. doi: <https://doi.org/10.1093/icc/dtj010>
- Ejdemo, T., & Örtqvist, D. (2020). Related variety as a driver of regional innovation and entrepreneurship: A moderated and mediated model with non-linear effects. *Research policy*, 49(7), 104073. doi: <https://doi.org/10.1016/j.respol.2020.104073>
- Etzkowitz, H., & Brisolla, S. N. (1999). Failure and success: the fate of industrial policy in Latin America and South East Asia. *Research policy*, 28(4), 337-350. doi: [https://doi.org/10.1016/S0048-7333\(98\)00077-8](https://doi.org/10.1016/S0048-7333(98)00077-8)
- Feldman, M. P., Kogler, D. F., & Rigby, D. L. (2015). rKnowledge: The spatial diffusion and adoption of rDNA methods. *Regional studies*, 49(5), 798-817. doi: <https://doi.org/10.1080/00343404.2014.980799>
- Frenken, K., & Boschma, R. A. (2007). A theoretical framework for evolutionary economic geography: industrial dynamics and urban growth as a branching process. *Journal of Economic Geography*, 7(5), 635-649. doi: <https://doi.org/10.1093/jeg/lbm018>
- Frenken, K., Izquierdo, L. R., & Zeppini, P. (2012). Branching innovation, recombinant innovation, and endogenous technological transitions. *Environmental Innovation and Societal Transitions*, 4, 25-35. doi: <https://doi.org/10.1016/j.eist.2012.06.001>
- Fujii, H., & Managi, S. (2018). Trends and priority shifts in artificial intelligence technology invention: A global patent analysis. *Economic Analysis and Policy*, 58, 60-69. doi: <https://doi.org/10.1016/j.eap.2017.12.006>
- Grashof, N. (2021). Putting the watering can away—towards a targeted (problem-oriented) cluster policy framework. *Research policy*, 50(9), 104335. doi: <https://doi.org/10.1016/j.respol.2021.104335>
- Hidalgo, C. A. (2021). Economic complexity theory and applications. *Nature Reviews Physics*, 3(2), 92-113.
- Hidalgo, C. A., & Hausmann, R. (2009). The building blocks of economic complexity. *Proceedings of the national academy of sciences*, 106(26), 10570-10575. doi: <https://doi.org/10.1073/pnas.0900943106>
- Hidalgo, C. A., Klinger, B., Barabási, A.-L., & Hausmann, R. (2007). The product space conditions the development of nations. *science*, 317(5837), 482-487. doi: <https://doi.org/10.1126/science.1144581>
- Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, 313(5786), 504-507. doi: <https://doi.org/10.1126/science.1127647>
- Hoppmann, J. (2021). Hand in hand to Nowhereland? How the resource dependence of research institutes influences their co-evolution with industry. *Research policy*, 50(2), 104145. doi: <https://doi.org/10.1016/j.respol.2020.104145>
- Juhász, S., Broekel, T., & Boschma, R. (2021). Explaining the dynamics of relatedness: The role of co-location and complexity. *Papers in Regional Science*, 100(1), 3-21. doi: <https://doi.org/10.1111/pirs.12567>
- Klinger, J., Mateos-Garcia, J., & Stathoulopoulos, K. (2018). Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology. doi: <https://arxiv.org/pdf/1808.06355.pdf>
- Kuusk, K., & Martynovich, M. (2021). Dynamic Nature of Relatedness, or What Kind of Related Variety for Long-Term Regional Growth. *Tijdschrift voor economische en sociale geografie*, 112(1), 81-96. doi: <https://doi.org/10.1111/tesg.12427>

- Leusin, M. E., Günther, J., Jindra, B., & Moehrle, M. G. (2020). Patenting patterns in Artificial Intelligence: Identifying national and international breeding grounds. *World Patent Information*, 62, 101988. doi: <https://doi.org/10.1016/j.wpi.2020.101988>
- Li, X., & Jiang, H. (2017). Artificial Intelligence Technology and Engineering Applications. *Applied Computational Electromagnetics Society Journal*, 32(5).
- Magro, E., & Wilson, J. R. (2019). Policy-mix evaluation: Governance challenges from new place-based innovation policies. *Research policy*, 48(10), 103612. doi: <https://doi.org/10.1016/j.respol.2018.06.010>
- Montesor, S., & Quatraro, F. (2020). Green technologies and Smart Specialisation Strategies: a European patent-based analysis of the intertwining of technological relatedness and key enabling technologies. *Regional studies*, 54(10), 1354-1365. doi: <https://doi.org/10.1080/00343404.2019.1648784>
- Naldi, L., Criaco, G., & Patel, P. C. (2020). Related and unrelated industry variety and the internationalization of start-ups. *Research policy*, 49(10), 104050. doi: <https://doi.org/10.1016/j.respol.2020.104050>
- Neffke, F., Henning, M., & Boschma, R. (2011). How do regions diversify over time? Industry relatedness and the development of new growth paths in regions. *Economic geography*, 87(3), 237-265. doi: <https://doi.org/10.1111/j.1944-8287.2011.01121.x>
- Nilsson, N. J. (2009). *The quest for artificial intelligence*: Cambridge University Press.
- Petralia, S., Balland, P. A., & Morrison, A. (2017). Climbing the ladder of technological development. *Research policy*, 46(5), 956-969. doi: <https://doi.org/10.1016/j.respol.2017.03.012>
- Rigby, D. L. (2015). Technological relatedness and knowledge space: entry and exit of US cities from patent classes. *Regional studies*, 49(11), 1922-1937. doi: <https://doi.org/10.1080/00343404.2013.854878>
- Righi, R., Samoil, S., Cobo, M. L., Baillet, M. V.-P., Cardona, M., & De Prato, G. (2020). The AI technological complex System: Worldwide landscape, thematic subdomains and technological collaborations. *Telecommunications Policy*, 101943. doi: <https://doi.org/10.1016/j.telpol.2020.101943>
- Santoalha, A., Consoli, D., & Castellacci, F. (2021). Digital skills, relatedness and green diversification: A study of European regions. *Research policy*, 50(9), 104340. doi: <https://doi.org/10.1016/j.respol.2021.104340>
- Schmoch, U. (2008). Concept of a technology classification for country comparisons. *Final report to the world intellectual property organisation (wipo)*, WIPO.
- Soete, L. (1987). The impact of technological innovation on international trade patterns: the evidence reconsidered. *Research policy*, 16(2-4), 101-130. doi: [https://doi.org/10.1016/0048-7333\(87\)90026-6](https://doi.org/10.1016/0048-7333(87)90026-6)
- Squicciarini, M., Dernis, H., & Criscuolo, C. (2013). *Measuring patent quality: Indicators of technological and economic value*. Retrieved from OECD Publishing, Paris:
- Tanner, A. N. (2016). The emergence of new technology-based industries: the case of fuel cells and its technological relatedness to regional knowledge bases. *Journal of Economic Geography*, 16(3), 611-635. doi: <https://doi.org/10.1093/jeg/lbv011>
- Teece, D. J. (2018). Profiting from innovation in the digital economy: Enabling technologies, standards, and licensing models in the wireless world. *Research policy*, 47(8), 1367-1387. doi: <https://doi.org/10.1016/j.respol.2017.01.015>
- Uyarra, E., Zabala-Iturriagoitia, J. M., Flanagan, K., & Magro, E. (2020). Public procurement, innovation and industrial policy: Rationales, roles, capabilities and implementation. *Research policy*, 49(1), 103844. doi: <https://doi.org/10.1016/j.respol.2019.103844>
- Van Den Berge, M., & Weterings, A. (2014). Relatedness in eco-technological development in European regions. *Papers in Evolutionary Economic Geography*, 14(13), 1-30.
- Whittle, A. (2020). Operationalizing the knowledge space: theory, methods and insights for Smart Specialisation. *Regional Studies, Regional Science*, 7(1), 27-34. doi: <https://doi.org/10.1080/21681376.2019.1703795>

WIPO. (2019). *WIPO Technology Trends 2019: Artificial Intelligence*. Retrieved from https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf

Yoo, Y., Boland Jr, R. J., Lyytinen, K., & Majchrzak, A. (2012). Organizing for innovation in the digitized world. *Organisation science*, 23(5), 1398-1408. doi: <https://doi.org/10.1287/orsc.1120.0771>

Appendix A: Search strategy applied in PATSTAT 2019 spring version for identifying AI patents.

The SQL query presented below comprehends all keywords used to identify AI patents. This search strategy is similar to the one presented and discussed in Leusin, M. E., Günther, J., Jindra, B., & Moehrle, M. G. (2020). Patenting patterns in Artificial Intelligence: Identifying national and international breeding grounds. *World Patent Information*, 62, 101988, with the exception that we exclude the generic keyword “Artificial Intelligence” and its related Wikipedia synonym “Machine Intelligence”, which are not included in the list of AI-techniques presented in WIPO (2019). For more details on each keyword and synonym, see Leusin et al (2020), Tables 1 and 2 (pg. 3 and 4, respectively).

Select appln_id from tls202_appln_title

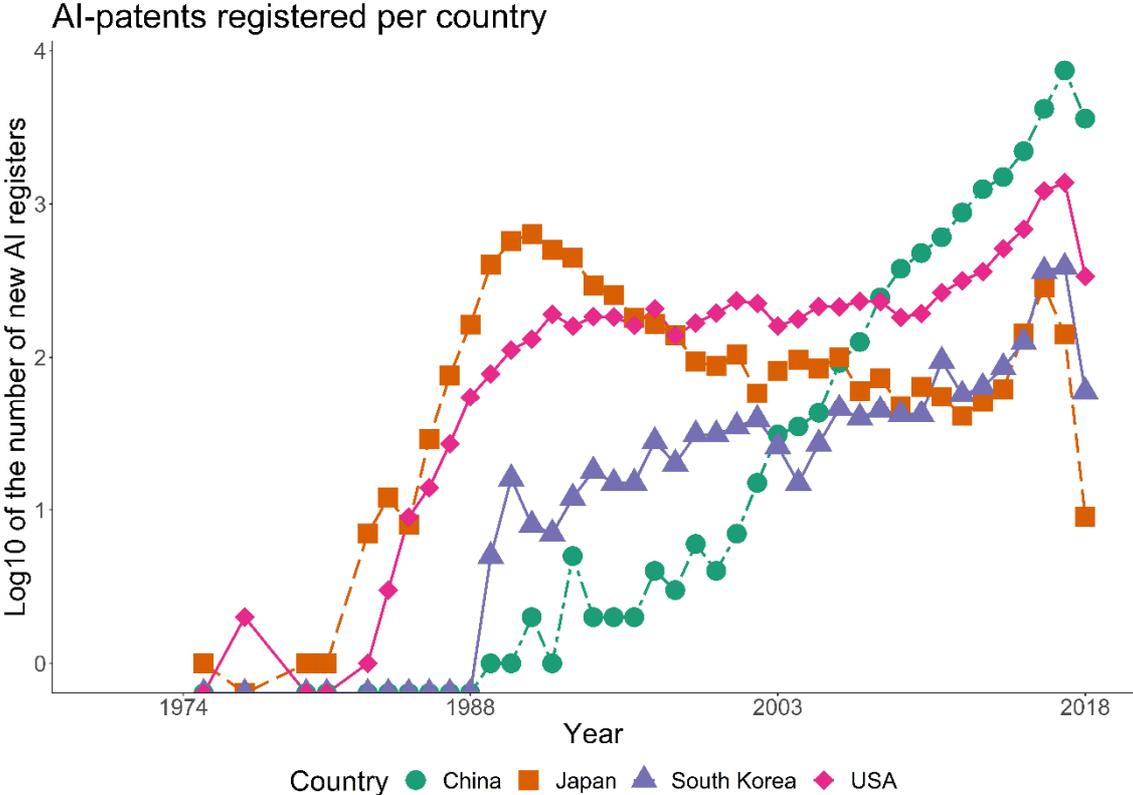
Where appln_title like '%machine learn%' OR appln_title like '%Probabilistic reason%' OR appln_title like '%Fuzzy logic%' OR appln_title like '%Logic Programming%' OR appln_title like '%Ontology engineer%' OR appln_title like '%pervised learn%' OR appln_title like '%reinforced learn%' OR appln_title like '%task learn%' OR appln_title like '%neural network%' OR appln_title like '%deep learn%' OR appln_title like '%expert system%' OR appln_title like '%support vector machin%' OR appln_title like '%description logistic%' OR appln_title like '%classification tree%' OR appln_title like '%regression tree%' OR appln_title like '%logical learn%' OR appln_title like '%relational learn%' OR appln_title like '%probabilistic graphical model%' OR appln_title like '%rule learn%' OR appln_title like '%instance-based learn%' OR appln_title like '%latent represent%' OR appln_title like '%bio-inspired approach%' OR appln_title like '%probability logic%' OR appln_title like '%probabilistic logic%' OR appln_title like '%reinforcement learn%' OR appln_title like '%multitask learn%' OR appln_title like '%Decision tree learn%' OR appln_title like '%support vector network%' OR appln_title like '%deep structured learn%' OR appln_title like '%hierarchical learn%' OR appln_title like '%graphical model%' OR appln_title like '%structured probabilistic model%' OR appln_title like '%Rule induction%' OR appln_title like '%memory-based learn%' OR appln_title like '%bio-inspired comput%' OR appln_title like '%biologically inspired comput%'

UNION

Select appln_id from tls203_appln_abstr

Where appln_abstract like '%machine learn%' OR appln_abstract like '%Probabilistic reason%' OR appln_abstract like '%Fuzzy logic%' OR appln_abstract like '%Logic Programming%' OR appln_abstract like '%Ontology engineer%' OR appln_abstract like '%pervised learn%' OR appln_abstract like '%reinforced learn%' OR appln_abstract like '%task learn%' OR appln_abstract like '%neural network%' OR appln_abstract like '%deep learn%' OR appln_abstract like '%expert system%' OR appln_abstract like '%support vector machin%' OR appln_abstract like '%description logistic%' OR appln_abstract like '%classification tree%' OR appln_abstract like '%regression tree%' OR appln_abstract like '%logical learn%' OR appln_abstract like '%relational learn%' OR appln_abstract like '%probabilistic graphical model%' OR appln_abstract like '%rule learn%' OR appln_abstract like '%instance-based learn%' OR appln_abstract like '%latent represent%' OR appln_abstract like '%bio-inspired approach%' OR appln_abstract like '%probability logic%' OR appln_abstract like '%probabilistic logic%' OR appln_abstract like '%reinforcement learn%' OR appln_abstract like '%multitask learn%' OR appln_abstract like '%Decision tree learn%' OR appln_abstract like '%support vector network%' OR appln_abstract like '%deep structured learn%' OR appln_abstract like '%hierarchical learn%' OR appln_abstract like '%graphical model%' OR appln_abstract like '%structured probabilistic model%' OR appln_abstract like '%Rule induction%' OR appln_abstract like '%memory-based learn%' OR appln_abstract like '%bio-inspired comput%' OR appln_abstract like '%biologically inspired comput%'

Appendix B: Log 10 of number of AI patents by Japan, the United States (USA), South Korea, and China.



Appendix C: Technological fields connectivity considered in the Global Technological Space (GTS) network.

Field name	Sector	Degree of connectivity
Computer technology	Electrical engineering	1,000
Electrical machinery, apparatus, energy	Electrical engineering	0,955
Basic materials chemistry	Chemistry	0,938
Organic fine chemistry	Chemistry	0,913
Measurement	Instruments	0,912
Chemical engineering	Chemistry	0,906
Audio-visual technology	Electrical engineering	0,901
Other special machines	Mechanical engineering	0,856
Optics	Instruments	0,839
Pharmaceuticals	Chemistry	0,825
Materials, metallurgy	Chemistry	0,822
Macromolecular chemistry, polymers	Chemistry	0,820
Semiconductors	Electrical engineering	0,804
Surface technology, coating	Chemistry	0,790
Telecommunications	Electrical engineering	0,782
Biotechnology	Chemistry	0,738
Digital communication	Electrical engineering	0,719
Control	Instruments	0,708
Transport	Mechanical engineering	0,688
Environmental technology	Chemistry	0,670
Textile and paper machines	Mechanical engineering	0,631
Mechanical elements	Mechanical engineering	0,615
Medical technology	Instruments	0,611
Civil engineering	Other fields	0,607
Machine tools	Mechanical engineering	0,602
Handling	Mechanical engineering	0,601
IT methods for management	Electrical engineering	0,566
Engines, pumps, turbines	Mechanical engineering	0,545
Analysis of biological materials	Instruments	0,510
Thermal processes and apparatus	Mechanical engineering	0,508
Other consumer goods	Other fields	0,499
Food chemistry	Chemistry	0,491
Furniture, games	Other fields	0,427
Basic communication processes	Electrical engineering	0,376
Micro-structural and nano-technology	Chemistry	0,346

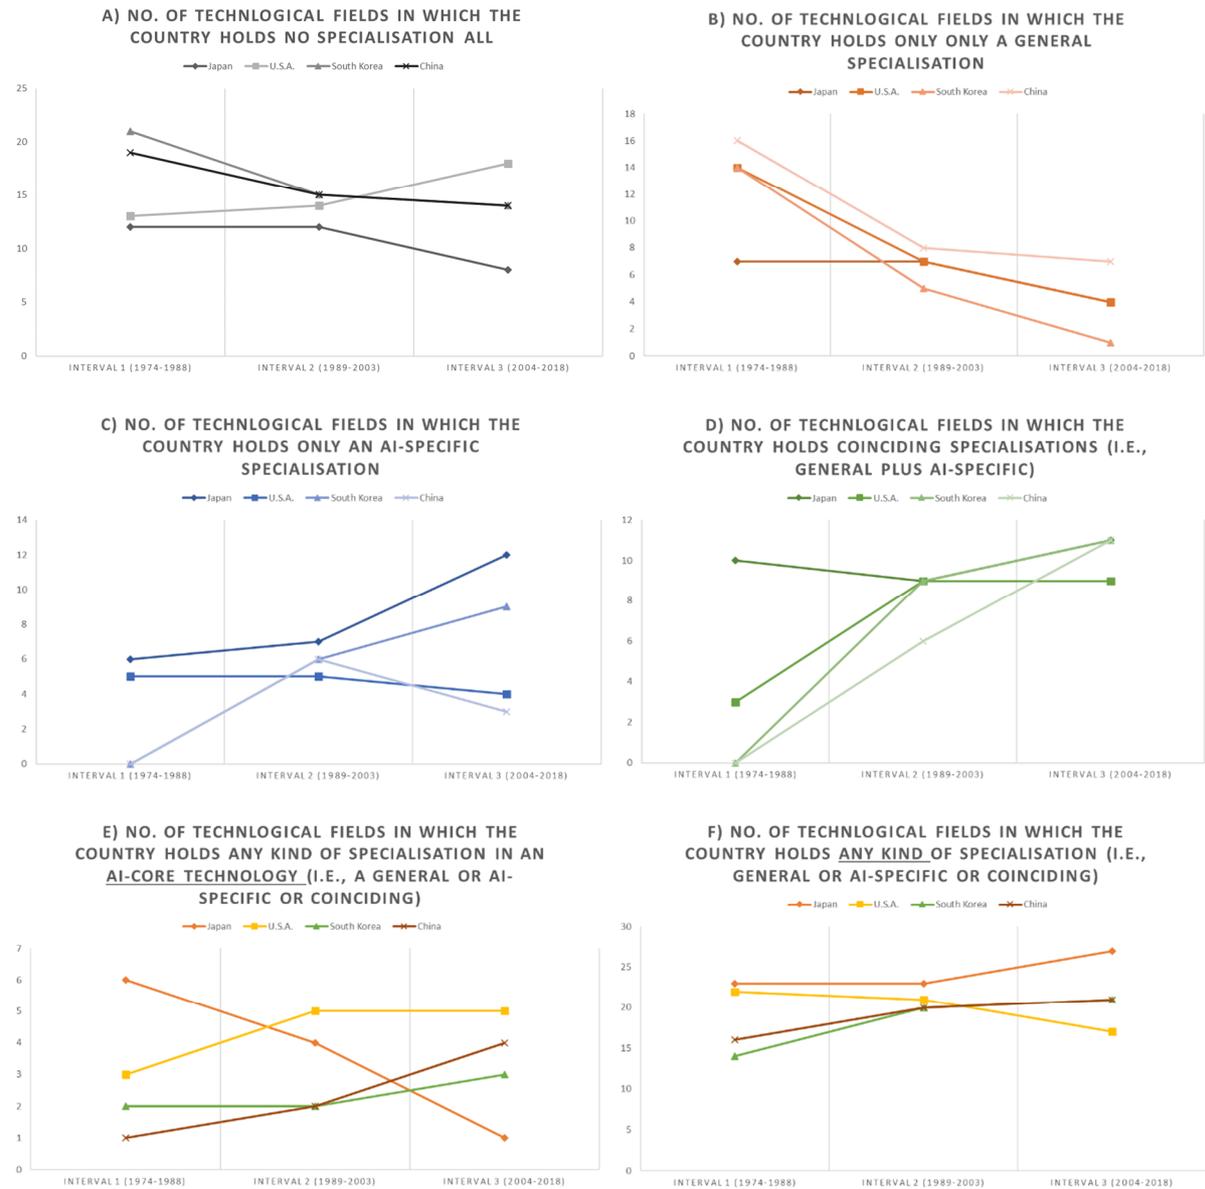
Appendix D: Tabular values of countries specialisations over the considered intervals (D1), followed by their respective plots (D2).

The tabular data of specialisations presented in Figure 13 is shown below, followed by the respective plots. Looking at the data, one can see no significant change in the number of specialisations that countries hold over time (Indicators A. and F.). Similarly, there is no clear pattern for AI-specific specialisations: The USA maintains a somewhat stable number, whereas Japan increase it considerably only in the third interval (Indicators C). The latecomers China and South Korea, which develop their first AI-specialisations in the second interval, decrease and increase, respectively, their number of specialisations of this type. Both China and South Korea increase their number of specialisations in AI-core technologies (Indicator E.) though, while Japan decrease it, and the United States (U.S.A.) shows a somewhat stable trend. A trend become apparent when we consider General and Coinciding specialisations (Indicators B. and D., respectively): All countries decrease their number of General-only specialisations, while their coinciding specialisations steadily increase.

D1. Tabular data of countries' number of specialisations over the considered intervals.

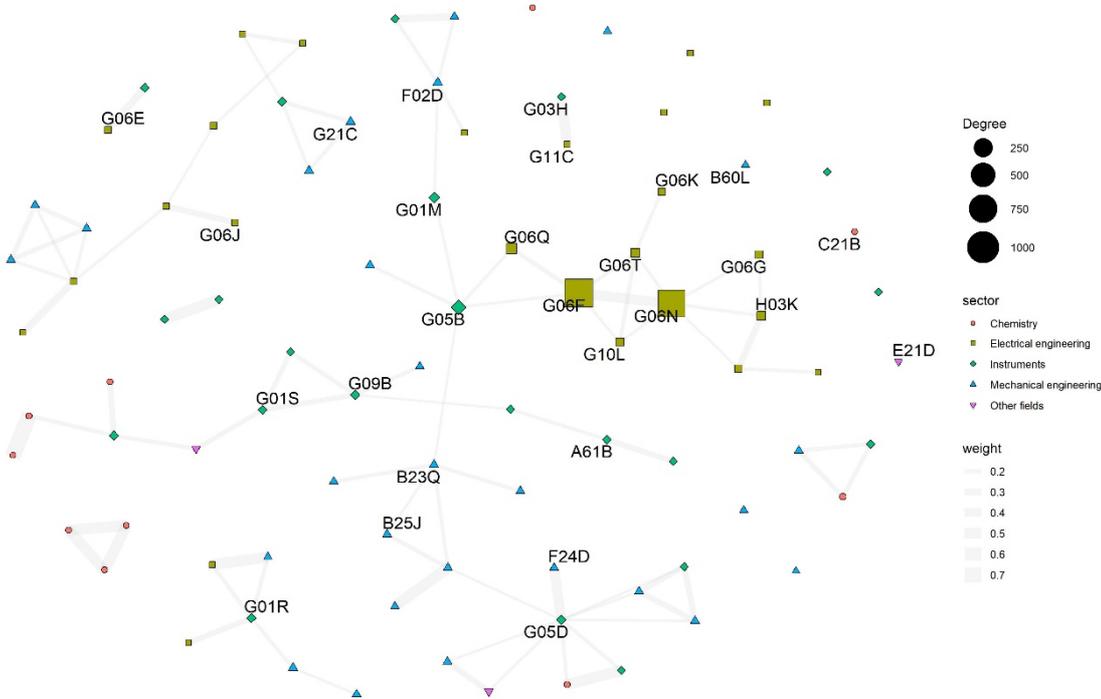
Indicator	Country	Interval		
		Interval 1 (1974-1988)	Interval 2 (1989-2003)	Interval 3 (2004-2018)
a) No. of technological fields in which the Country holds <u>no</u> specialisation all	Japan	12	12	8
	U.S.A.	13	14	18
	South Korea	21	15	14
	China	19	15	14
b) No. of technological fields in which the Country holds only a <u>General specialisation</u>	Japan	7	7	4
	U.S.A.	14	7	4
	South Korea	14	5	1
	China	16	8	7
c) No. of technological fields in which the Country holds only an <u>AI-specific specialisation</u>	Japan	6	7	12
	U.S.A.	5	5	4
	South Korea	0	6	9
	China	0	6	3
d) No. of technological fields in which the Country holds <u>coinciding specialisations</u> (i.e., General plus AI-specific)	Japan	10	9	11
	U.S.A.	3	9	9
	South Korea	0	9	11
	China	0	6	11
e) No. of technological fields in which the Country holds <u>any</u> kind of specialisation in an <u>AI-core technology</u> (i.e., a General OR AI-specific or coinciding)	Japan	6	4	1
	U.S.A.	3	5	5
	South Korea	2	2	3
	China	1	2	4
f) No. of technological fields in which the Country holds any kind of specialisation (i.e., general OR AI-specific OR coinciding)	Japan	23	23	27
	U.S.A.	22	21	17
	South Korea	14	20	21
	China	16	20	21

Appendix D2. Plots of the previous tabular data

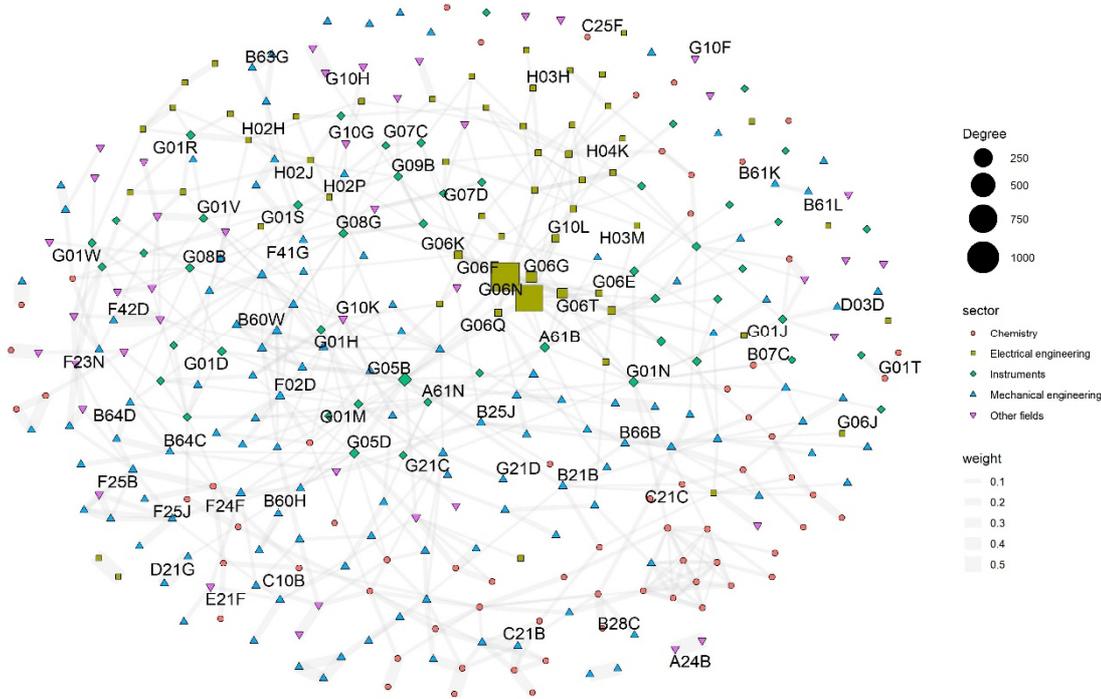


Appendix E: AI relatedness and specialisations over the considered intervals for IPC subclasses.

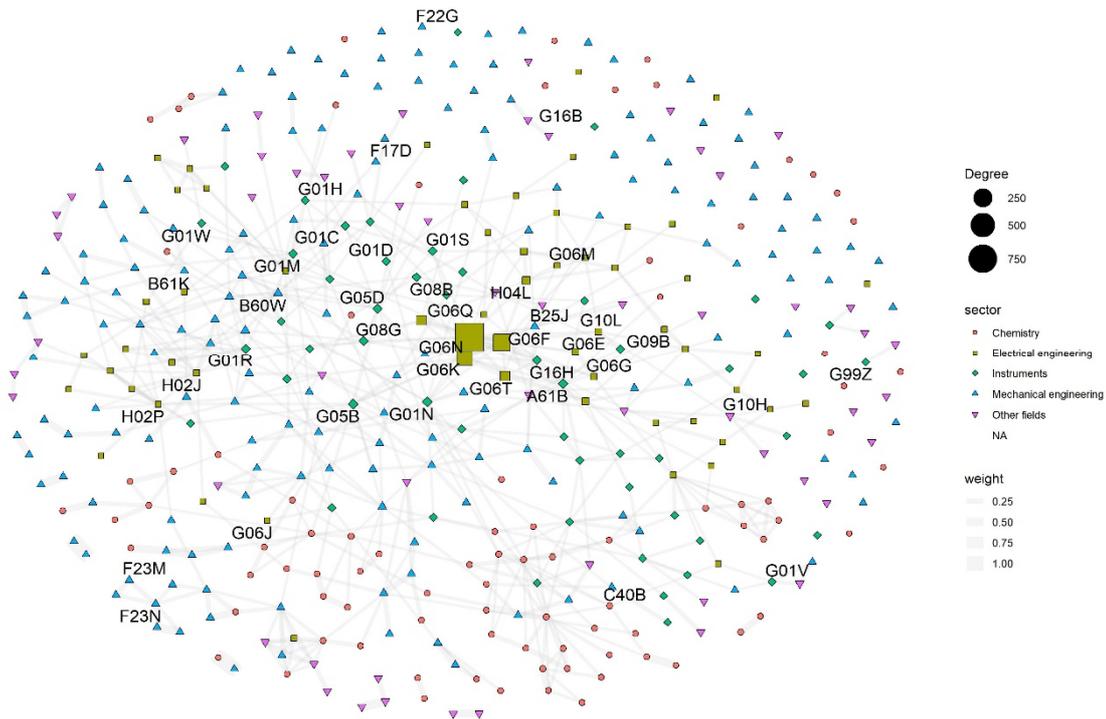
AI-specific technological space: 4-digits IPC level (1974-1988)



AI-specific technological space: 4-digits IPC level (1989-2003)



AI-specific technological space: 4-digits IPC level (2004-2018)



Note that, similarly to Figure 10, AI starts as a narrow network that evolves to a densely and larger connected one centred around some main codes linked to computer technologies. Here, the main central code is G06N, which refers to “Computer systems based on specific computational models”. Note that all ten main subclasses highlighted previously in Figure 15 also appear as central here in the third interval (i.e., with a positive specialisation). Subclasses related to the sectors “Chemistry” and “Other fields” are shown mainly on the periphery of the network, whereas subclasses related to the sectors “Electrical engineering” and “Instruments” occupy again central positions.

Author: Matheus Eduardo Leusin

Abstract: This chapter investigates how the development of AI-related inventions by Multinational Enterprises (MNEs) affects their technological trajectories and innovative performance. I combine a matched-pair analysis with an extension of the Difference-in-Difference method to analyse these effects over a novel panel dataset of MNEs. This dataset links over 30 thousand MNEs to more than 10 million patents that these companies owned directly or indirectly (i.e., through their subsidiaries) in the period from 2011 to 2019. The results indicate that MNEs introducing AI-related inventions increase the relatedness of subsequent inventions by about 10 per cent compared to a control group. These results are robust when accounting for a self-selection bias. AI is thus being used to reinforce the existing technological trajectories, rather than to disrupt them. The results also suggest that the number of subsequent inventions is about 40 per cent higher for MNEs that introduce AI during the observation period compared to the control group, without significant effects on the intensity of R&D expenditures per invention. It is argued that this increase in innovative performance is linked not only to knowledge dynamics created by learning about AI but also by AI's technical potential to be used for learning.

Keywords: Technological trajectory; Relatedness; Artificial Intelligence; Innovative performance;

JEL Classification: D22; O14; O33; L25

Publication

This chapter has been published as a working paper as:

Leusin, M. E., 2022. The Development of AI in Multinational Enterprises – Effects upon Technological Trajectories and Innovation Performance. Bremen Papers on Economics and Innovation No. 2201. University of Bremen, Bremen.

Formal changes, compared to the published version, have been made.

4.1. Introduction

So far, the digital transformation has been manifesting itself primarily by the diffusion of information and communication technologies (ICTs) in business and society (Alcácer, Cantwell, & Piscitello, 2016). The industrial diffusion of technologies related to Artificial Intelligence (AI) represents a qualitatively new development in this current ICT paradigm. AI's transformative qualities arguably offer new opportunities that could significantly change firms' processes of technological development.

First, AI is a digital technology, meaning that it has particular properties linked to convergence and generativity that allow, among other things, its embeddedness into non-physical products and the emergence of combinatorial innovations (Yoo, Boland Jr, Lyytinen, & Majchrzak, 2012). Firms can use AI to transform their existing nondigital products into digital ones, and recombine them to create further innovations with new functionalities. Second, due to its power of classification and prediction, AI is expected to be an Invention of a Method of Invention (IMI). This refers to AI's potential to serve as a tool of "automated discovery" across many knowledge domains, where it could be used to expand the set of problems that firms can feasibly address (Cockburn, Henderson, & Stern, 2018). Third, AI functions as an intelligent autonomous agent and is a modular technology (Nilsson, 2009), which can affect organisations' functions, processes, product functionalities, and competencies (Paschen, Pitt, & Kietzmann, 2020). As a result, AI can both enhance or devalue the existing knowledge and skills held by companies (*ibid*).

These particularities lead to great expectations about AI's potential. The mastering of this technology is expected to give countries advantages over global markets and industries (Cave & ÓhÉigeartaigh, 2018; Cockburn et al., 2018; Klinger, Mateos-Garcia, & Stathoulopoulos, 2018). At the firm level, AI is expected to increase productivity gaps, potentially leading non-adopters to lose large market shares or even to go out of business (European Commission, 2017). AI is also linked to the risk of unemployment due to its automation potential (Goralski & Górniak-Kocikowska, 2020; Iscan, 2021; Su, 2018). But these expectations miss the point that AI will have very distinct impacts across industries. The main argument of this chapter is that AI's true potential depends on how much this technology can change the technological trajectories of a broad set of firms. The underlying premise is that most of AI's impact will not come from the most obvious affected business (e.g., digital platforms such as Facebook, Uber,

Airbnb, or more broadly, ICT-related companies), but rather from how it transforms the larger proportion of existing businesses.

The most recent empirical evidence highlights that AI is linked to an increase in firms' innovative performance in Rammer, Czarnitzki, and Fernández (2021), improvement in wages (Genz, Gregory, Janser, Lehmer, & Matthes, 2021), and that access to digital data contributes significantly to an increase in the number of AI innovations (Beraja, Yang, & Yuchtman, 2020). But there is a knowledge gap regarding how AI affects firms' technological trajectories. After all, is AI a new technological paradigm that allows firms to explore a completely new set of productive problems, or is it an extension of the current ICT paradigm?

Thereby, I focus on multinational enterprises (MNEs), since the current wave of AI adoption is skewed towards large and established firms (WIPO, 2019; Zolas et al., 2021). Firstly, I want to understand how firms' innovations change after AI adoption in comparison to their previous innovations. In particular, this chapter suggests that the introduction of a proprietary AI innovation into an MNE's technological portfolio leads to the emergence of increasingly related innovations. I expect that MNEs will recombine their AI innovation with other technologies from their existing portfolio rather than aiming at new unrelated innovations or technological breakthroughs. Secondly, I want to confirm if AI adoption is linked to a change in firms' innovative performance, and if so, explore what are the mechanisms behind this change. Does AI affect firms' innovative performance due to the learning about this technology, or due to the use of this technology as an IMI? Given current evidence, I expect that the introduction of an AI innovation into MNEs' technological portfolio increases their innovative performance. Moreover, I expect that this result is not only linked to AI's potential as a new piece of knowledge, but also due to its potential as a learning tool.

These two main hypotheses are explored by applying a matching procedure to a novel dataset combining MNE and patent data for the observation period from 2011 to 2019. I match AI adopters to a control group based on industry, company size, age, and innovation output. AI effects are estimated considering the year that it was adopted by each firm using the extension of the Difference-in-Difference (DiD) method proposed in Callaway and Sant'Anna (2018, 2020). The estimated effects indicate that AI adopters develop innovations increasingly related to their knowledge bases, strengthening their existing technological trajectories. This result is consistent even for technological sectors that are unrelated to AI. I also find that AI

adoption increases the performance of MNEs in terms of innovation output. Considering that even unrelated sectors increase their relatedness through AI adoption, I argue that both AI knowledge and its use as a learning tool contribute to the verified increase in firms' innovative performance. I also argue that AI is better linked to a continuous technological change within the current ICT technological paradigm than to a discontinuous change towards a new paradigm.

This chapter is structured as follows: Section 2 provides the relevant theoretical background and specifies the research questions; Section 3 describes the creation of the dataset and the methodology; Section 4 presents the empirical findings; Section 5 discusses the main findings, contributions, and limitations of the chapter.

4.2. Theoretical Background and Research questions

The evolutionary view of technological development highlights that knowledge production is a cumulative, path dependent, and interactive process (Dosi, 1982; Nelson & Winter, 1982). In principle, the adoption of digital technologies, such as AI, offers a technological opportunity that could lead firms to modify the way that they produce new knowledge. This change is understood to be the result of firms' progressing along their technological trajectories (Dosi, 1982).

4.2.1. Digital Technologies and AI

Emerging possibilities created by technological change give firms new windows of technological opportunity, allowing them to change their technological trajectories (Dosi, 1982). These changes can be continuous, if firms stay within their technological paradigms, or discontinuous if they make it possible for firms to enter into a new paradigm⁵⁹ (ibid). One of the most important recent technological developments is the ubiquitous use and diffusion of digital technologies (Alcácer et al., 2016). These technologies spawned novel value creation and value appropriation pathways, besides producing a new breed of innovations (Nambisan, Lyytinen, Majchrzak, & Song, 2017). They give a basis for combinatorial and distributed innovation (Yoo et al., 2012) which may provide firms with a technological opportunity to expand the domains of knowledge that they can explore and absorb efficiently.

⁵⁹ A technological paradigm comprehends the general outlook of the productive problems faced by firms (Dosi, 1982).

The adoption and pervasive use of digital technologies also allow the emergence of affordances linked to generativity and convergence (Yoo et al., 2012), which are created through unique properties of digital technologies, namely reprogrammable functionality and data homogenisation. Generativity refers to the action potential of digital technologies for producing innovations characterised by unprompted change driven by large, varied, and uncoordinated actors. It causes digital technologies to become inherently dynamic and malleable. Convergence, in turn, refers to the action potential of embedding digital technologies in nondigital artefacts. This enables coupling several distinct products or tools into one. Both affordances also favour the emergence of combinatorial innovation, since distinct digital technologies, like software-based modules, can be integrated to produce further innovations (ibid)⁶⁰.

AI is arguably one of the most important advanced digital technologies (Zolas et al., 2021). Most of AI's current power comes from its use as a tool for prediction (Cockburn et al., 2018). Powered by digital inputs, AI can be used to find valuable patterns in large amounts of data, besides also making predictions and decisions accordingly. This refers to AI's potential to be an IMI, which could be used to expand the set of problems that firms can feasibly address (ibid). Through the automation of complex tasks, AI can help companies in finding new ways of creating value and enhance their knowledge and capabilities. But this very same power can also replace the existing knowledge and capabilities of firms (Paschen et al., 2020). AI could also arguably enable firms to discover a completely new set of productive problems, which would make it a bridging technology towards a new technological paradigm.

4.2.2. Relatedness and Innovative Performance

Relatedness was proposed based on the concept of absorptive capacity, which proposes that a firm's ability to explore new knowledge depends on how related it is to the firm's prior knowledge (Cohen & Levinthal, 1990). Extant research documented that large firms are multi-technology corporations that combine a usually greater number of technologies to develop and produce a lower number of products and services (Granstrand, 1998); corporate technological diversification changes only slowly over time (Cantwell & Andersen, 1996); and profiles of technological diversification differ across firms due to history, distinct initial

⁶⁰ The widespread diffusion of Application Programming Interfaces (APIs), which are used to connect distinct software, is an example of how digital innovations can even focus on integrating distinct digital-based systems to allow for the emergence of combinatorial innovation.

conditions, market incentives, distinct institutional settings and other factors (Antonelli, Krafft, & Quatraro, 2010; Ivarsson, Alvstam, & Vahlne, 2015; Le Bas & Sierra, 2002), despite being very similar among large firms producing similar products (Breschi, Lissoni, & Malerba, 2003; Teece, Rumelt, Dosi, & Winter, 1994).

Breschi et al. (2003) propose the concept of relatedness and present evidence that firms diversify their innovative activities mainly by exploring related technologies, with larger innovators being typically more coherent in their technological trajectories than smaller innovators. The importance of relatedness has been linked to knowledge dimensions of proximity, commonality, and complementarity. Knowledge proximity relates to firms' learning processes, which can be unintended or intended. Both are created by firms' focus on learning about technologies that are similar to what they know in terms of knowledge base. Knowledge commonalities imply that firms' innovative activities may span over more than one technology because the same type of knowledge is used in various technologies, whereas knowledge complementarity arises from the need to use distinct technologies together (ibid).

Subsequent studies provide evidence that firms innovating in related areas have higher survival rates (Colombelli, Krafft, & Quatraro, 2013), lower coordination costs (Nesta, 2008), and perform better in knowledge transfer and creation (Weber & Weber, 2010). Relatedness has been shown to moderate positively an inverted U-shaped relationship between technological diversification and technological performance (Kim, Lee, & Cho, 2016; Leten, Belderbos, & Van Looy, 2007). It eases the burden of having a widely diversified portfolio by reducing the costs of learning, but just to a limited extent. Too much diversification – despite offering further opportunities for cross-fertilisation and technology fusion – may hurt innovative performance due to higher coordination and integration costs (Leten et al., 2007). Harmful effects of excessive diversification can be attenuated by firms that develop sufficient competencies around their core technologies (Kim et al., 2016).

The relatedness concept was also extended to explain geographic patterns of innovation (Hidalgo, 2021; Hidalgo, Klingler, Barabási, & Hausmann, 2007). Corresponding research documented that the exploration of related technologies is linked to higher knowledge production (Kogler, Rigby, & Tucker, 2013) and innovative output of regions (Aarstad, Kvitastein, & Jakobsen, 2016; Castaldi, Frenken, & Los, 2015; Delgado, Porter, & Stern, 2014; Solheim, Boschma, & Herstad, 2018; Solheim, Boschma, & Herstad, 2020). The diminishing

returns of related innovations are also recognised (Antonelli et al., 2010; Ejdemo & Örtqvist, 2020). Conversely, the exploration of unrelated knowledge was linked to technological breakthroughs and radical innovations, but at higher costs of learning (Castaldi et al., 2015; Solheim et al., 2018; Solheim et al., 2020).

4.2.3. Hypotheses development

Given the possible influence of digital technologies on firms' absorptive capacity, the effects of introducing such technologies on MNEs' technological portfolios are closely examined. In particular, this chapter analyses how the introduction of an AI innovation⁶¹ – considered a type of advanced digital technology – affects the technological trajectories and innovative performance of MNEs.

These technological trajectories are identified through the concept of relatedness. If MNEs following innovations become technologically different from their previous innovations after the introduction of AI, MNEs' relatedness will decrease. If they are similar or fill close gaps of knowledge that the MNE has, relatedness will remain stable or increase, respectively.

Given the affordances of digital technologies, it is argued that the introduction of an AI innovation in a firm's portfolio will be followed by recombinations of this innovation to create further innovations. On the one hand, given the reduced costs of producing related innovations and the potential of AI for combinatorial innovations, MNEs might recombine their AI innovation with their existing portfolio rather than aiming at new unrelated innovations or technological breakthroughs. The convergence affordance of digital innovations suggests that MNEs can "transform" their nondigital artefacts into digital ones by coupling them with AI, reinforcing their existing technological trajectories. In addition, it is noted that it is very hard to capture the value of digital innovations (Teece, 2018). This is due to difficulties in protecting digital innovations from being copied. Accordingly, it is also hard to reinforce intellectual property rights related to these innovations (ibid). Hence, the generativity aspect of digital innovations offers little direct economic benefits to upstream innovators. Still, these innovators may benefit from technological improvements made to

⁶¹ It is worth highlighting what the "introduction of AI" means in the context of this paper. The term, used throughout the paper, refers to the first moment when an MNE has ownership of an AI patent (directly, or through its subsidiaries). This ownership can happen through the self-development of this patent, or through acquisition (by buying the patent or a company that owns the patent). Hence, in the paper context, MNEs become AI adopters in the first year that they owned (directly or indirectly through one of their subsidiaries) an AI patent. Patents are the proxy for innovation, which is a limitation discussed in the final remarks of the paper.

their protected innovation in downstream sectors. This would allow them (i.e., the innovators who hold property rights over the original digital innovation) to produce additional improved versions of their innovations, increasing their technological relatedness. On the other hand, if relatedness decreases, it would indicate that MNEs use AI as an automated discovery tool to develop innovations in new knowledge domains. Considering that the first case is more likely than the second, I hypothesise:

1. The introduction of an AI innovation into MNEs' technological portfolio leads to the emergence of increasingly related innovations.

It can also happen that the introduction of an AI innovation has opposite effects across MNEs from distinct technological sectors. This would occur if the knowledge dynamics generated through learning about AI vary according to how distinct the new AI knowledge is from the firm's existing knowledge base. For firms from technological sectors that are related to AI, the three dimensions of knowledge (i.e., knowledge proximity, commonalities, and complementarity) are to play out and affect the increase in relatedness. These technologically close sectors are to contain the companies most impacted by AI, e.g., the ICT-related sectors linked to companies like Facebook, Airbnb, and Uber. But for firms from sectors technologically distant from AI, AI knowledge will be a new piece of unrelated knowledge. Hence, only the knowledge complementarity dimension can arise. In this case, the new AI knowledge could only influence relatedness if it would be combined by the firm with another piece of its existing knowledge. This could happen, for example, in the cases where firms use AI to transform their nondigital products into digital ones.

But combinations of unrelated and related knowledge are very difficult and occur more seldomly (Castaldi et al., 2015). In these cases, relatedness would remain unchanged or even decrease at the moment that AI is first introduced. This is because at least one piece of unrelated knowledge will be contained in the firm's following innovations. This piece of knowledge will only become related to the firm's knowledge base by being repeated in further innovations⁶². Due to knowledge dimensions, the introduction of an AI innovation into MNEs'

⁶² To understand how this happens, it is helpful to think about the implementation of the relatedness indicator. Using patents as a proxy for innovations, as it is done in this paper, the indicator compares how similar an innovation is to the existing knowledge base of a firm through the technological codes used to classify innovations (in the case of this paper, the 4-digits IPC codes). Consider a manufacturing company creating an autonomous car. The first innovations of the company may be mixing the existing knowledge (i.e., a technological code linked to "cars") with a piece of knowledge represented by a code linked to software (used for making the

technological portfolio is to have distinct effects across firms from different technological sectors. Hence, I hypothesise:

2. MNEs from sectors that are close to AI's knowledge increase their relatedness, whereas MNEs from sectors more distant to AI's knowledge show no effect or decrease their relatedness after AI introduction.

If no difference is seen across technological sectors, it potentially means that AI effects don't come only from its potential as a new piece of knowledge. The effects could also come from using AI as a learning tool, through its IMI potential for example. To address this possibility, the effects on innovative performance are also analysed. It is suggested that AI innovations affect positively the innovative performance of MNEs due to: i) lower costs linked to the exploration and development of new related knowledge; ii) AI's potential for automating tasks and being a learning tool capable of extending firms' inventive capabilities; and iii) the malleable and dynamic aspects of AI as a digital technology.

The first mechanism refers to possible effects linked to learning about AI (i.e., absorption of AI as a piece of knowledge). Extant research at the regional level links the creation of increasingly related innovations to higher innovative performance (Aarstad et al., 2016; Castaldi et al., 2015; Kim et al., 2016; Leten et al., 2007). Thus I expect a similar relationship at the firm level. The second and third mechanisms refer to particular aspects of AI as a technology. AI has been described as an IMI, which would also lead to an increase in firms' innovative performance. Regarding this potential, Cockburn et al. (2018) highlight that it comes from AI being used as a tool to automate discovery. When applied over large datasets, AI can recognise important patterns and generate accurate predictions which help firms in the development of innovations. The greatest potential of AI as an IMI is that it could be used to learn patterns across very distinct technological sectors. This would allow firms to explore an enormous set of new knowledge (ibid). The literature on digital innovations also suggests that an improvement in innovative performance is expected due to the adoption of this type of technology (e.g., Hanelt, Firk, Hildebrandt, and Kolbe (2021); Huang, Henfridsson, Liu, and Newell (2017); Khin and Ho (2019)). This is linked to the malleability aspect of digital

car autonomous). If the firm never patented software (or did it seldomly), this innovation is only half related to the company knowledge base (thus, relatedness remain unchanged, or is reduced). But then, the company may innovate by improving the software of the car, creating further software-related innovations (which is made possible by the combinatorial aspect of digital innovations). As the company accumulates a considerable number of software-related codes in its knowledge-base, relatedness also starts to increase.

innovations, which allows them to be quickly improved or adapted to new contexts, which also favour the emergence of innovations.

But firms could also reduce their innovative performance due to the introduction of AI. This could happen if firms shift their technological trajectories to develop mainly unrelated innovations, which are linked to lower efficiency since such innovations are typically linked to higher R&D efforts (Solheim et al., 2020). In this case, mechanisms (ii) and (iii) would need to be strong enough to enable firms to increase their patent outputs at a higher rate than the increase in R&D expenditure. Hence, my last hypothesis is:

3. The introduction of an AI innovation into MNEs' technological portfolio increases their innovative performance.

I expect the overall impact of AI introduction on innovative performance to be positive due to the individual or combined potential effects of the three mentioned mechanisms.

4.3. Data & Method

4.3.1. Constructing the Dataset

The dataset combines three sources owned by Bureau van Dijk (BvD), namely Orbis, Orbis IP, and Orbis-Zephyr. All data was downloaded between July 2020 and March 2021. The construction of the company and patent datasets is done separately and merged through the unique standardised identifiers ('BvD ID number'). Details about the creation of the dataset are presented in Appendix A.

The company dataset separates subsidiaries from their Global Ultimate Owners (GUOs) based on a minimum ownership threshold of 25.01%. This results in an ownership structure with up to 21 ownership levels, which is then filtered for corporate companies (see Appendix A, Step I). M&D data is used to recursively extend the 2020 generated ownership structure backward, year by year. To this end, all M&A deals that relate to the GUOs or subsidiaries as of 2020 were downloaded from the Orbis-Zephyr database (Appendix A, Step II). In the following step, a patent dataset is built by downloading from Orbis IP all priority filings⁶³ (granted or not) whose priority dates were within the period from 01/01/2000 to 31/12/2019 (Appendix A, Step III).

⁶³ A priority filing is the first patent application filed to protect an invention. If the same patent is registered in other patent offices, the following registrations are called non-priorities, constituting a patent family linked through the priority filing.

In this step, the AI patents are identified through a specific “AI tag” that allows finding them later. The identification of AI patents adopts the list of keywords proposed in Leusin, Günther, Jindra, and Moehrle (2020), which is based on typical AI techniques. The authors report high accuracy of their proposed search strategy after comparing the quality of their results with alternative approaches (namely the ones proposed in Tseng and Ting (2013) and Fujii and Managi (2018)). All registers containing any of the considered AI-keywords in their titles, abstracts, claims, or description are identified as AI patents. For more details about the keywords adopted, see Appendix A, Table A1.3.

Finally, the BvD ID of patent owners is used to link this patent dataset to the companies’ ownership dataset (Appendix A, Step IV). The patents are linked to their yearly owners, which allows accounting for changes in ownership due to patent or firm acquisitions, for example. As the ownership structure of companies covers only the 2011-2020 period, patents registered from 01/01/2000 to 31/12/2010 are linked to their owners using the ownership structures existent in 2011. These patents are used as the initial stock of companies. The year 2020 is dropped from the analysis due to the possibility of incomplete patent data.

4.3.2. Matching AI Adopters to Non-adopters

Matching methods offer causal inference with fewer assumptions, enabling higher robustness and less sensitivity to modelling assumptions (Ho, Imai, King, & Stuart, 2007). The premise is to construct a dataset where the treated group is as similar as possible to the control group, which means excluding data about units that are too divergent from their possible counterfactuals until a balance is reached. The evaluation of the matching procedure is then straightforward in the sense that the preferred generated dataset is the one that produces the best balance between distance (i.e., the difference between the control and treatment units regarding the considered criteria) and the number of excluded units.

In my case, the treatment is defined as the ownership of at least one AI patent in a given year, whereby both the development of an AI innovation by MNE units as well as the acquisition of an AI innovation (through buying a patent or acquiring a subsidiary that owns an AI patent) are considered. Therefore, the treatment considers all patents that are linked to the identified GUO directly or indirectly (i.e., through its subsidiaries)⁶⁴. GUOs that adopted AI before 2011

⁶⁴ As a robustness test, an alternative measure considering only the patents owned directly by the GUO is also considered when looking at path dependency.

are excluded. The criteria for matching the remaining GUOs are the number of patents owned per year, age, industry (NACE 4-digit level), and size class⁶⁵. Except for the number of patents and age⁶⁶, the remaining criteria are applied through exact matching. Regarding number of patents, all years from the considered period (2011-2019) plus five years of pre-treatment control (i.e., from 2006 to 2010) are considered. This means that treated companies are matched to control companies throughout every year of the period from 2006 to 2019⁶⁷. A caliper⁶⁸ of 0.10 standard deviations is used to define the limits of how “similar” the number of patents has to be. Seldomly used NACE codes are put together under a created label “Less used codes”, which replaces the codes of companies that are found less than five times between AI adopters. Altogether, this label comprehends 20% of the considered NACE codes.

The matching implementation is done using the R package MatchIt (Stuart, King, Imai, & Ho, 2011). The matching method with the most balanced results is based on a genetic algorithm. Once there are considerably more companies that didn’t adopt AI than companies that did it, a ratio of 5 control units is applied in the matching to increase the size of the considered sample (i.e., five non-adopters are matched to one AI adopter for every “matchable” treated GUO).

4.3.3. Estimating the Effects of AI Adoption

I use the extension of the Difference-in-Differences (DiD) model proposed in Callaway and Sant’Anna (2018, 2020) for estimating treatment effects with multiple time periods. It extends the classical model by adding two assumptions: The first is the “Staggered treatment adoption assumption”, which implies that once a unit participates in the treatment it remains treated. The second assumption restricts the DiD parallel trends assumption to consider that there is also no treatment effect heterogeneity across time, across groups receiving treatment in

⁶⁵ According to Orbis classification, companies with 1,000 or more employees are “Very large companies”, companies with less than 1,000 but with 150 or more employees are “Large companies”, companies with less than 150 but with 15 or more employees are “Medium sized companies”, and the rest are considered “Small companies”.

⁶⁶ Assuming that small age differences are more significant for newer companies than the oldest, I consider six periods for the age category: 0-1979, 1980-1999, 2000-2004, 2005-2009, 2010-2019. The last period considers more years due to the very limited number of companies created in the last five years considered.

⁶⁷ This is an important criterion to consider given the structure of my dataset, where the stock of patents of companies “follows” them if they are acquired by other companies (in this way, if the number of patents of a company goes to 0 in a given year, it means the company was sold or stopped operating; given my matching criteria, such company would be matched only if there was another company that was also sold/stopped operating in the very same year).

⁶⁸ A caliper defines the number of standard deviations of the distance measure within which to draw control units.

distinct time periods, or in the treatment sequence⁶⁹. Through this extension, the treatment group can be further divided into smaller groups based on the period on which these distinct groups receive treatment. Thereby a generic treatment group G is divided into $g(i)$ subgroups, with treated ($d=1$) and untreated ($d=0$) units. The units from each subgroup $g(i)$ are treated in a distinct time period “ t_i ” and their group-specific counterfactuals (i.e., the control units where $d=0$) are assumed to follow parallel paths in all post-treatment periods $t \geq g$. The treatment group is divided into 9 subgroups (i.e., $i = 9$), with each subgroup containing the treatment units that adopted AI in a distinct year t , which varies from 2011 to 2019. A public repository with the data and R implementation of the estimations made in this Chapter is available at https://github.com/matheusleusin/AI_and_MNEs.

Callaway and Sant’Anna (2018, 2020) also propose distinct aggregation schemes for highlighting treatment effect heterogeneity and pre-treatment estimates. For the former, I use the aggregation methods of “Dynamic Treatment Effects” and “Simple weighted average”. The pre-treatment estimates, in turn, refer to a “pre-test” to check if the parallel trends assumption holds. The main idea behind these estimates is that they can be used to check if pre-treatment units are similar enough to treated units in the time periods before treatment. It is called a “pre-test” because although the estimates are useful for the pre-treatment period, they can’t say anything about the parallel trends assumption holding during the treatment and post-treatment periods. The parallel trends assumption in these periods, in turn, is fundamentally untestable. The authors point out that this “pre-test” is an important piece of evidence on the credibility of the DiD design in a particular application. Accordingly, I present and discuss the results of this pre-test together with the estimation effects. All effects are estimated using the DID-R package presented in Brantly Callaway (2021).

4.3.4. Dependent Variables Considered

4.3.4.1. Path dependency

Path dependency is calculated following the relatedness density index proposed in Hidalgo et al. (2007) and adapted in Balland (2016) for patents. The implementation of the index is done

⁶⁹ There are two implementation options for this second assumption (i.e., the “Parallel Trends Assumption based on not-yet treated units” and the “Parallel Trends Assumption based on never-treated units”). The “not-yet” option is particularly useful for cases where there is low availability of treatment units, as it allows these units to be used both as treatment i.e., in the time periods before they get treated) and as control (in the time periods in which they are treated. As a limited number of control units is not a particular problem of my dataset, I choose the “never-treated” implementation option.

through the EconGeo R package (Balland, 2017)⁷⁰. Results lay between 0 and 100 percent, with higher values highlighting higher relatedness. The index is calculated as follows:

$$RD_{i,c,t} = \frac{\sum_{j \in c, j \neq i} x_i * \varphi_{ij}}{\sum_{j \neq i} \varphi_{ij}} \times 100$$

Where “i” stands for a technology (in my case, a 4-digits IPC code) found in the company “c” at a time “t”. “ φ_{ij} ” stands for the technological relatedness between two technologies “i” and “j”, which is based on the premise that if two technologies co-occur more often than what would be expected by chance under the assumption of statistical independence, they are more related. These values are normalised following the cosine normalisation procedure (Eck & Waltman, 2009). “ x_i ”, in turn, refers to a binary implementation of the Revealed Technological Advantage (RTA) index⁷¹, taking the value of 1 if the company has a specialisation in technology “i”, and 0 otherwise.

The density equation works by capturing the average relatedness between the technologies in which the firm is specialised. Mathematically, the RTA equation makes it possible to develop specialisations only in a maximum of half of all activities explored by a firm. Considering the exemplary case of an increase in the density index, this may happen for two distinct reasons: The firm may have developed an additional specialisation that is in the average closer to the technological cluster in which the firm is, or the firm may have switched one of its previously existing specialisations for a specialisation in a technology that is closer to its technological cluster. Thus, the emergence of a more coherent technological path, proxied by relatedness, means that firms can be either adding close specialisations or switching to closer specialisations. New specialisations can be added just to a maximum limit of about half of the technologies the firm is patenting in, whereas switching specialisations can occur endlessly.

4.3.4.2. Innovative Performance

Three distinct proxies are considered for measuring innovative performance: i) the number of patents owned by the MNE; ii) the share between the number of patents owned by the MNE

⁷⁰ The average relatedness density, implemented through the “relatedness.density.int.avg” function, which in my case considers as input an incidence matrix with companies depicted in rows and their number of registers in each 4-digits IPC code in columns, and an adjacency symmetric matrix of IPC codes indicating the degree of relatedness between them.

⁷¹ $RTA = \frac{patents_{c,t}(i) / \sum_i patents_{c,t}(i)}{\sum_c patents_{c,t}(i) / \sum_c \sum_i patents_{c,t}(i)}$, rounded to 1 if $RTA \geq 1$, and 0 otherwise.

per unit of turnover (in thousands of USD), and iii) the share between the R&D expenditure (in thousands of USD) per number of patents owned by the MNE. These three proxies offer distinct perspectives over innovative performance: (i) focus directly on innovative output; (ii) considers how this output changes according to the firm’s economic performance; and (iii) is closely linked to the idea of innovative efficiency, which is defined as the ability of firms to translate innovation inputs into outputs (Hollanders & Esser, 2007). These perspectives are complementary: If (i) increases, showing an improvement in innovative output, the shares (ii) and (iii) should decrease (showing a proportionally larger increase in economic performance and a decrease in R&D expending for developing patents, respectively) to corroborate (i).

4.4. Empirical Findings

4.4.1. Overview of Dataset

Overall, the dataset contains data for over 30 thousand GUOs, which are linked to more than 1 million subsidiaries. These companies own together more than 10 million patents, from which roughly 90 thousand are identified as being related to AI. The data also shows the rapid advancement of AI adoption in the considered period: From 5.6% of GUOs having at least one AI patent in 2011 to a share of 10.2% in 2019 (see Appendix B for an overview). It also shows a high concentration of large and old MNEs as the main AI adopters, with a high share of these coming from ICT-related sectors related. These results are very similar to the findings presented in Zolas et al. (2021), which are based on survey data on US companies. Additional information about the main AI adopters according to the number of AI patents owned is presented in Appendix C (which also includes a comparison with the data presented in Fujii and Managi (2018)).

The matching procedure drops most of the top AI adopters since many of them adopted AI before 2011. However, the distribution of firms considering the matching criteria adopted remains similar. AI adopters are primarily from “Computer programming activities” (NACE 6201), 5829 “Other software publishing” (NACE 5829), and “Manufacture of electronic components” (NACE 2611), which together represent about 20% of the industry codes used by adopters. Further information comparing companies before and after applying the matching procedure is presented in Appendix D.

4.4.2. Estimated effects

4.4.2.1. Effects on Relatedness

I use the matched dataset to estimate the aggregated treatment effects of AI adoption (see Table 6). The “Simple weighted average” highlights the weighted average of all considered groups with weights being proportional to the size of each group. One disadvantage of this kind of measure is that early-treated groups are observed for more time periods. The second method, namely “Dynamic Treatment Effects”, avoids this limitation by weighting the average treatment effects according to the different lengths of exposure to the treatment (Callaway & Sant’Anna, 2018, 2020).

In the matched dataset I still have some MNEs classified as holding companies (NACE 6420) and head offices (NACE 7010). Both codes could include firms with quite different industrial profiles, which might be an issue for estimating treatment effects. To account for that, MNEs linked to these two sectors are excluded as a robustness check in (b). Overall, the effects of AI on MNEs’ technological relatedness are significant at the 95% confidence interval for all calculations. The effect becomes stronger when head offices and holdings are dropped (see Table 6, column b). The value of 1.037 means an increase of 10.3% over the average relatedness of 10.095 for non-AI adopters; the value of 1.153, in turn, means an increase of 11.4% over the average relatedness of non-AI adopters⁷².

⁷² As an additional robustness check, similar results hold when considering GUOs’ direct ownership of patents only: The estimated effects are 0.639 under the simple weighted average, which is significant from 0.2672 to 1.0115*, and 0.504 for dynamic treatment effects, which is non-significant for the dynamic treatment effects with a confidence interval from -0.0599 to 1.0679. 871 treated units are included in the sample when only GUOs’ direct ownership is considered.

		a) MNE's relatedness	b) MNEs relatedness without considering holdings and head offices
Simple Weighted Average	Estimated effect	1.037	1.114
	Standard error	0.169	0.179
	Conf. interval	from 0.7054 to 1.3689*	from 0.7625 to 1.4657*
Dynamic Treatment Effects	Estimated effect	1.153	1.253
	Standard error	0.261	0.281
	Conf. interval	from 0.6409 to 1.6643*	from 0.7032 to 1.8036*
Number of treated units considered		1,155	1,054

Table 6: Estimates for the aggregate treatment effects of AI adoption on firms' relatedness.

The "Dynamic Treatment Effects" method can be further disaggregated across subgroups (see Figure 17). Particularly, length of exposure equal to 0 provides the average effect of participating in the treatment across subgroups in the time period when they first participate in the treatment (instantaneous treatment effect). Negative lengths of exposure correspond to time periods before subgroups first participate in the treatment, and lengths of exposure equal to or above 1 correspond to the time periods after initial exposure to the treatment (being 1 the first year after treatment). Grey bars with negative values refer to the estimators from the pre-treatment test proposed in Callaway and Sant'Anna (2018, 2020). This test can detect a broad set of violations against the conditional parallel assumption, which in turn form the basis for all the considered estimation procedures. The interpretation is that for values of 0 within the considered confidence band (in my case, 95%), this assumption is held.

It is seen in Figure 17 that the parallel trends assumption holds for all pre-treatment years, meaning that the considered treated and control companies are similar enough before treatment across groups (given the considered criteria)⁷³. The corresponding effect of AI

⁷³ These pre-treatment values go until -12 to pick the maximum number of years that one subgroup stayed in the sample without being treated (in this case, the Subgroup 2019, which wasn't treated during the 4 years

adoption increases for the first 5 years after adoption, before it starts to decline albeit remaining positive. The exclusion of MNEs classified as holding companies and head offices does not change this trend (see Figure 17b).

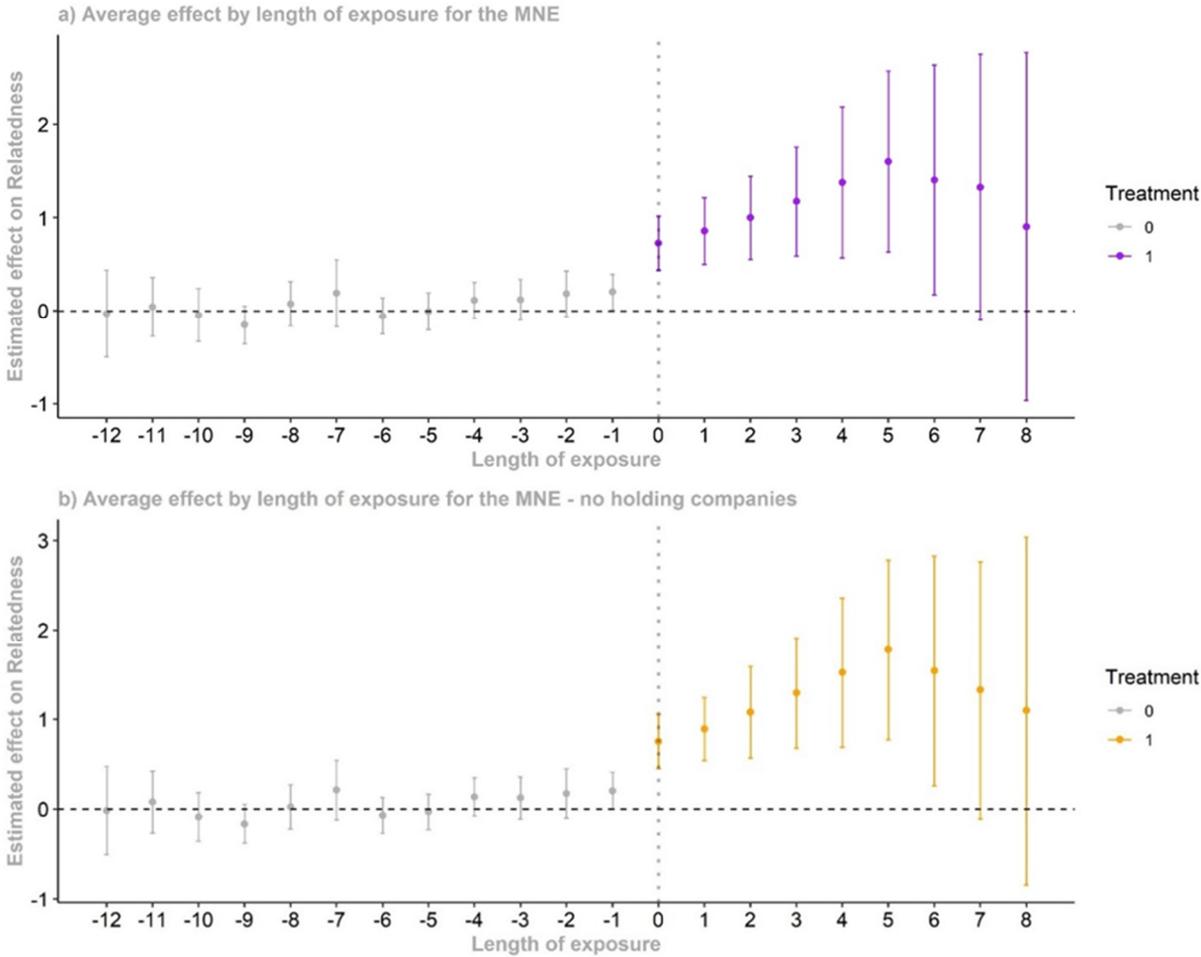


Figure 17: Estimated Dynamic Treatment Effects of AI adoption on the relatedness of all MNEs (17a) and when holdings and head offices are excluded (17b).

4.4.2.2. Relatedness effects across sectors and the self-selection bias

Using a nonrandomised sample for estimating the effects of AI may raise concerns regarding a self-selection bias. Although the “pre-test” may pick differences across MNEs before they receive treatment for the considered criteria, a self-selection problem may still exist if AI adopters are inherently different from non-adopters in a relevant but unaccounted way. I foresee two main situations where a self-selection would occur in the considered context. It may happen i) if AI-adopters and non-adopters differ in regards to their innovation strategy,

following the initial year of analysis, which is also reflected for all groups [i.e., considering the 4 years following 2006 until 2010] plus 8 years [from 2011 to 2018] in which this particular Subgroup was not treated whereas all the other subgroups were). Once the data goes just until 2019, the effects of the 9th group (i.e., companies that adopted AI firstly in 2019) are considered just for the immediate effect (i.e., at the length of exposure 0).

or ii) if adopting AI requires such a low effort that firms do it for trivial reasons (e.g., as a way to signalise something to competitors, investors, to the company itself, etc.). Although the dataset does not allow accounting for innovation strategies and other potentially relevant parameters, it is argued that the effects on firms' relatedness would have significant differences according to firms' knowledge distance to AI.

In the first case, in which AI adoption is driven by the effects of different "perceptions" towards AI, AI adopters are to increase their number of AI innovations. This could happen because of the favourable view towards AI, or to create a head start over competitors⁷⁴. In this situation, the knowledge distance to AI would lead companies from far away sectors to rapidly decrease their relatedness. Or, at the very least, companies from these sectors don't substantially change their innovation patterns (e.g., despite the AI strategy, innovative efforts are minimal), and relatedness doesn't change significantly due to AI adoption. The reason is straightforward: The technological "movement" towards AI for companies with knowledge bases close to AI means creating innovations that are relatively similar to the ones that these companies already had (increase in relatedness), whereas for companies distant to AI this would lead to the creation of innovations that are unrelated from what the companies had so far (decrease in relatedness). The effect would be similar for the second possible self-selection bias, albeit for a distinct reason. Some AI adopters could enter in the treatment just because the effort is minimal, but the costs would be greater the larger the distance is to AI (i.e., unrelated knowledge is more expensive to learn). Accordingly, firms with knowledge bases that are distant to AI aren't likely to create several costly AI innovations for trivial reasons, and no effects should be seen on relatedness.

Hence, I disentangle the matched pairs of treated and non-treated MNEs into three subgroups with distinct knowledge distances to AI. In doing so, I also verify the validity of the second proposed hypothesis. I apply two options to calculate distances: First, a very simplistic one in which all companies that belong to ICT-related sectors (NACE 5829, 6201, 6311, 6312, 6391, and 6399) are put in the closest to AI category, MNEs previously classified into the "Less used codes" (which comprehends the 20% NACE codes less used by AI adopters) are put in the category with the greatest distance to AI, and all remaining MNEs are put in a median group.

⁷⁴ Considering that patents offer the possibility of protecting the use of an invention, AI adopters could create several AI-related patents in a defensive way to protect knowledge seen as potentially having future economic value, for example.

In the second option, the average distance of every NACE code to the “AI cluster” is measured, and quartiles are used to create three distinct categories (see for a detailed description in Appendix E)⁷⁵.

The results of the new estimations considering the three distance groups across the two distinct distance measures are presented in Table 7. Figures for the disaggregation across groups under the “Dynamic Treatment Effects” are presented separately in Appendix F. Overall, the positive results are maintained for all groups, with the “Median group” being the only exception with non-significant results for the estimations based on knowledge distance and Dynamic effects. It is noted, however, that the magnitude of the effect changes across the considered groups. Particularly, closer subgroups present effects that are up to 7 times stronger. Regarding Appendix F, the results are also very similar to the trend seen previously: The effects on relatedness are increasingly positive in the first years after adoption, and then decline albeit staying positive. A particular difference is that the increasing trend remains longer for closer groups in comparison to the other two groups. Considering the assumption that a self-selection bias would be reflected in companies from sectors distant to AI showing no effects on relatedness, I assume that the consistency of significant results is indicative of no self-selection bias.

⁷⁵ The proposed measure seems to perform well in putting “digital MNEs” in the top quartile. The reference for this assessment is based on the Technical Annex entitled “The Top 100 Digital MNEs” made by (available at https://unctad.org/system/files/official-document/wir2017ch4_Annex_en.pdf), which identifies companies related to internet platforms, IT and Telecom companies, and other sectors as the main adopters of digital technologies.

			Closer group (1)	Median group (2)	Farthest away group (3)
Check 1: Based on codes' usage in the treated units	Simple Weigh. Avg.	Estim. effect	2.426	0.761	0.622
		Stand. error	0.465	0.213	0.310
		Conf. interval	from 1.5152 to 3.3368*	from 0.3439 to 1.1772*	from 0.0137 to 1.2299*
	Dyn. Treat. Effects	Estim. effect	2.888	0.807	0.571
		Stand. error	0.590	0.327	0.466
		Conf. interval	from 1.731 to 4.0449*	from 0.1665 to 1.4483*	from -0.3413 to 1.484
	N. of treated units considered		204	655	302
Check 2: Based on technological distance to the AI cluster	Simple Weigh. Avg.	Estim. effect	1.990	0.249	0.777
		Stand. error	0.297	0.302	0.243
		Conf. interval	from 1.4078 to 2.5713*	from -0.3429 to 0.8417	from 0.2997 to 1.2535*
	Dyn. Treat. Effects	Estim. effect	2.413	-0.016	0.856
		Stand. error	0.446	0.439	0.378
		Conf. interval	from 1.5389 to 3.2863*	from -0.8775 to 0.8451	From 0.1152 to 1.5977*
	Number of treated units considered		400	368	387

Table 7: Results estimated for the self-selection bias across the considered two distinct knowledge distances to AI.

4.4.2.3. Effects on Innovative Performance

Orbis provides company data for an up to 10-years span in its online version. This limitation affects two of the three considered proxies for innovative performance, namely Turnover and R&D expenditures. In addition, there is very often missing data for these indicators for some years⁷⁶. Therefore, I extended the data backward and then forward to fill the missing values of individual companies for these two indicators. All MNEs with available data for 2 or fewer

⁷⁶ 46.7% of the Turnover data is missing for at least one year; this share increases to 62.7% when R&D expenses data is considered.

years (2011-2019) are excluded. Furthermore, I excluded the group treated in 2011 when considering these two variables due to the lack of data for 2010, required for pre-testing before treatment. As R&D expenditure is more frequently missing than Turnover data, the final samples used for each proxy are slightly different (see Appendix G). Particularly, very large MNEs are overrepresented due to their more complete availability of data. To understand how much this potentially affects the results, I estimate the effects on the number of patents owned by the entire MNE for the sample of companies from the Turnover dataset (i.e., complete matched dataset minus the 38.5% companies lost in the Turnover sample). Table 8 summarises the estimation effects for this additional estimation, together with the other three considered proxies used to measure innovative performance. Estimations with holdings and head offices excluded from the samples are shown within square brackets.

		No. of patents owned by the MNE	No. of patents owned by the MNEs with Turnover data	No. of patents owned by the MNE/ Turnover	R&D expenses/No. of patents owned by the MNE
Simple Weigh. Avg.	Estimat. effect	25.01	32.134 [31.54]	-0.007 [-0.013]	-29,865.73 [-30,110.03]
	Stand. error	4.99	6.58 [7.79]	0.0144 [0.014]	18,008.61 [17,407.32]
	Conf. interval	from 15.23 to 34.80*	from 19.24 to 45.02* [from 16.27 to 46.82*]	from -0.035 to 0.021 [from -0.041 to 0.015]	from -65,161.96 5,430.50 [from -64,227.76 to 4,007.69]
Dyn. Treat. Effects	Estimat. effect	35.48	43.95 [44.77]	-0.013 [-0.013]	-27,061.3 [-27,347.35]
	Stand. error	7.05	9.66 [11.55]	0.014 [0.012]	18,619.97 [17,481.74]
	Conf. interval	from 21.66 to 49.30*	from 25.01 to 62.89* [from 22.13 to 67.42*]	from -0.041 to 0.016 [from -0.037 to 0.010]	from -63,555.77 9,433.17 [from -61,610.93 to 6,916.23]
N. of treated units considered		1155	690 [620]	690 [669]	444 [442]

Table 8: Estimated effects for the three distinct dependent variables related to innovative performance.

The results indicate that AI adoption significantly increases the innovative output of firms regarding the number of patents owned (see first two columns of Table 8). This effect is stronger for the sample with larger companies. The estimated effect of 25.017 for the simple weighted method in the complete dataset represents an increase of 40.0% over the average number of patents owned by non-treated MNEs (average of non-treated: 62.48; treated: 108.52); the estimated effect of 32.134 in the Turnover-related sample, which has the larger companies, represents an increase of 35.7% over the average number of patents owned by non-treated units (average of non-treated: 89.90, average of treated: 162.58). The effects, therefore, are relatively smaller for the sample with larger companies (despite being larger in absolute values). The dynamic effects for the number of patents across these two distinct samples are shown in Figure 18. It is seen that, compared to relatedness, innovative output increases more steeply and consistently. There is a drop in the last three years, but it is very little and the results remain significantly positive except for the last considered year⁷⁷.

⁷⁷ Please note that in Figure 18, treated units don't pass the pre-test up to 4 years before receiving treatment (for the complete sample) and 3 years for the larger-companies sample (18b and 18c), meaning they may be different from non-treated units and, accordingly, results should be considered with caution. Overall, the trend is very similar for the three considered samples.

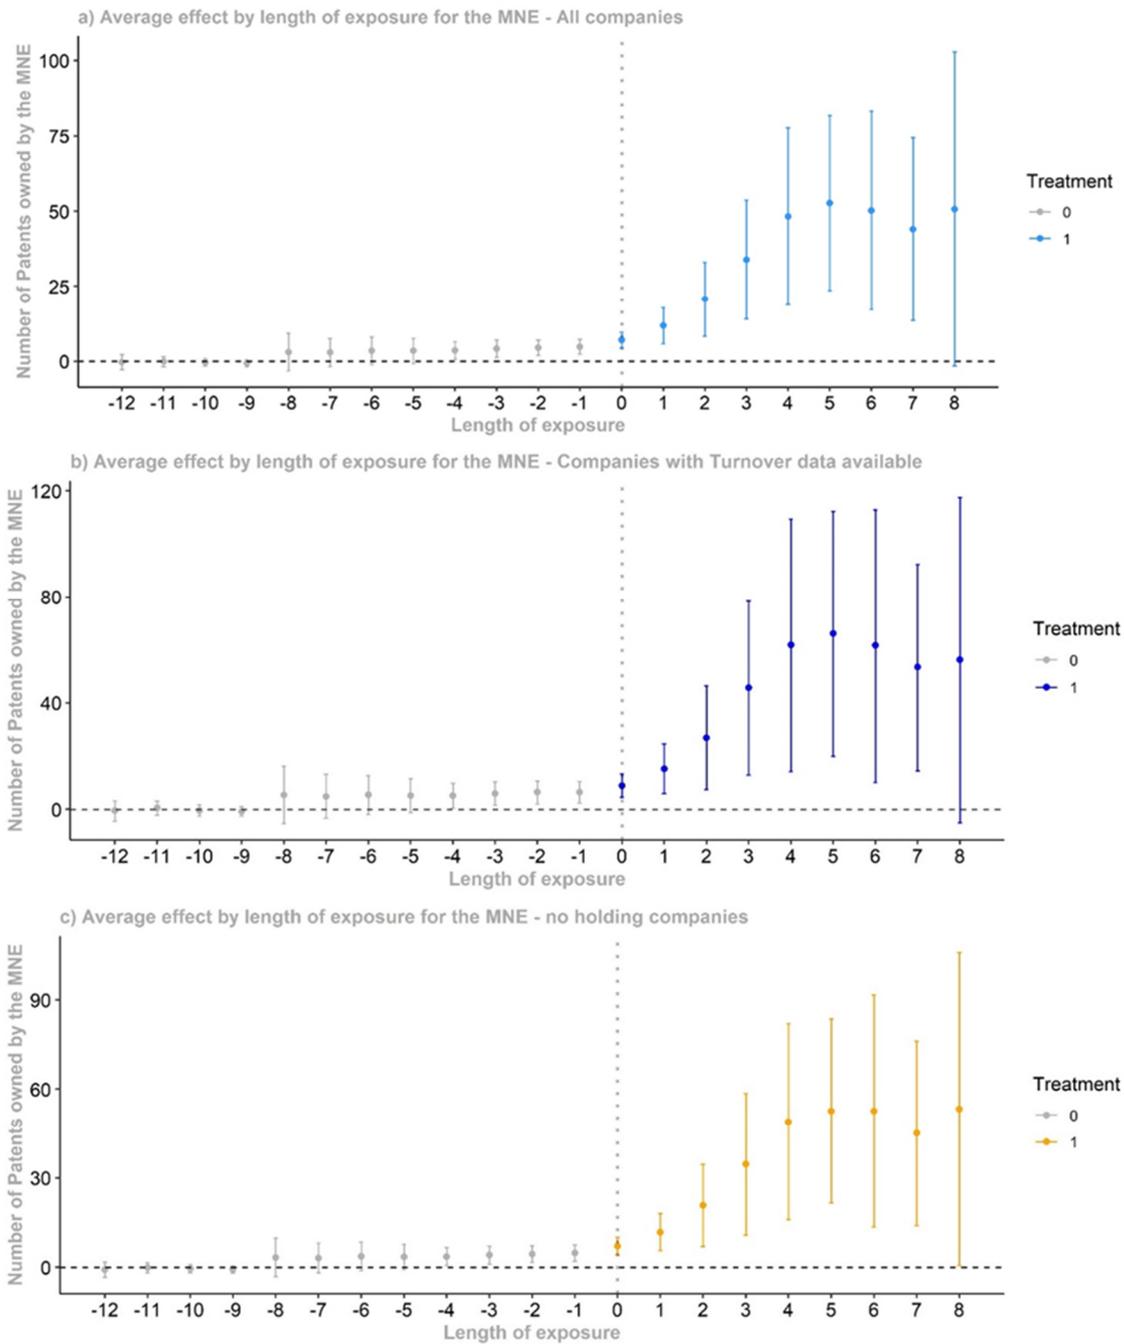


Figure 18: Estimated effects for the number of patents owned by the entire MNE for the entire (18a) and for the smaller samples with larger companies (18b and 18c).

The results for the other two considered proxies show that AI adoption decreases insignificantly both the number of patents owned by the MNE per unit of turnover and the share of R&D expenses per patent (see Table 8). The estimated negative effect of 0.007 for the simple weighted method represents a decrease of 83.42% over the average number of patents per unit of Turnover (which is 0.0084 for non-treated units, and 0.0098 for treated). The estimated negative effect of 29,865.73 for the same method represents a decrease of 80.25% over the average share between R&D expenditures and the number of patents owned

(which is 37,215.73 for non-treated units, and 112,120.46 for treated). The dynamic effects for these indicators⁷⁸ are shown in Figure 19.

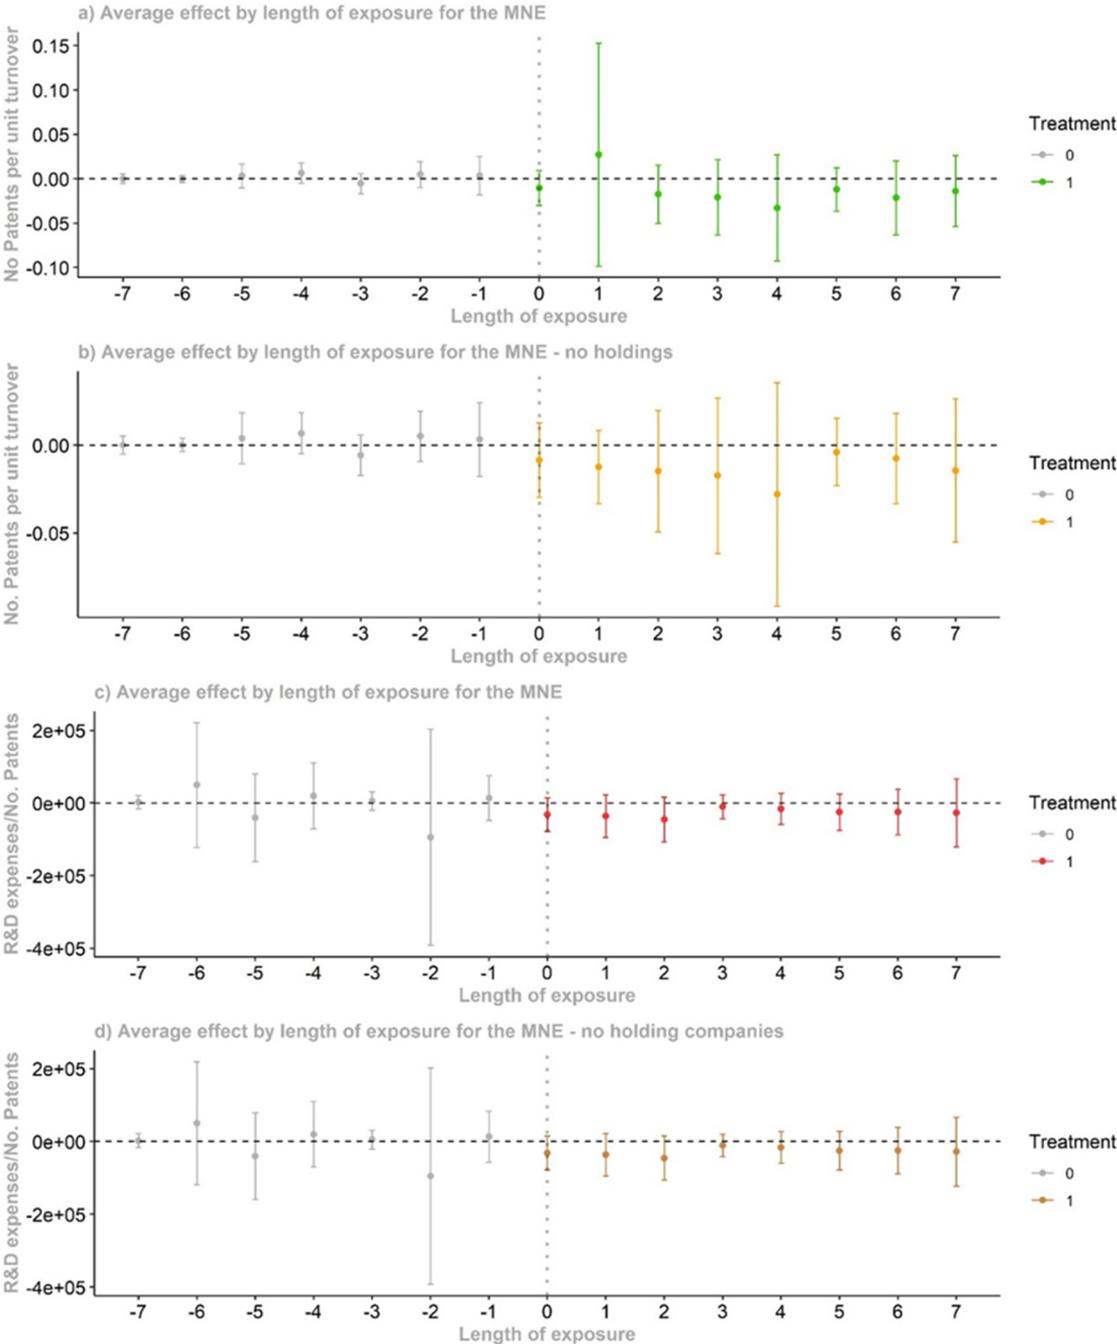


Figure 19: Estimated effects for Number of patents per unit of turnover (19a and 19b) and R&D expenses per patent (19c and 19d).

⁷⁸ Please note that there is data available only for the last 10 years for both of these parameters, the pre-test for the parallel trends assumption goes only until -7 (instead of -12 as in the previous figures). Also, as the 2011 group is dropped from both samples, the length of exposure goes only until 7 years instead of 8. Particularly regarding the exclusion of holdings and head offices, it is seen that the exposure for the first proxy seems more consistent after the exclusion (19b), whereas for the second proxy it seems to make no difference. Moreover, in comparison with the previous proxies, the negative effects seen here are relatively stable over time rather than increasing over time for some years and just then stabilising or diminishing.

4.5. Discussions and Conclusions

4.5.1. Discussion

Summing up, the results indicate that MNEs introducing AI-related innovations increase the relatedness of their subsequent innovations by about 10 to 11 per cent. The corresponding effect increases over the first 5 years after the introduction, before declining albeit remaining positive. The result seems robust when accounting for a possible self-selection bias. It is also seen that AI effects on relatedness remain positive across sectors with knowledge bases far away from AI knowledge. This means that even companies from sectors unrelated to AI increase the relatedness of their following innovations after its introduction. The evidence also indicates that MNEs increase by about 40 per cent the number of innovations they own after the introduction of an AI innovation. This effect grows steeply and considerably stronger over time in comparison to the effect seen for relatedness. This increase is also shown to occur without significantly changing the intensity of R&D expenditures per innovation. Thus, the empirical findings are in line with hypotheses (i) and (iii).

Particularly, the first hypothesis refers to AI's potential to increase firms' relatedness. It is confirmed that the introduction of an AI innovation is associated with such an effect, increasing firms' technological coherence. Thus, rather than being used as a tool to automate knowledge discovery across distant technological domains (Cockburn et al., 2018), AI reinforces the learning of related knowledge. This finding corroborates the anecdotal evidence presented in Brock and Von Wangenheim (2019) that AI is being deployed mostly to support companies' existing businesses. Taking AI as an example of advanced digital technology allows better understanding how the most recent wave of technological developments is potentially affecting firms' innovative activities. The evidence indicates the increasing role of absorptive capacity in the current information age. This finding is somewhat puzzling. Some aspects of digital innovations could arguably lead firms to become less path dependent as they become digitalised, e.g., digital innovations are reprogrammable (which makes them "malleable" to changes as highlighted in Yoo et al. (2012)), and digital resources allow companies to reduce their sunk costs related to physical goods. But these very same aspects are also to reinforce path dependency of technological sectors unrelated to digital technologies. It is very unlikely that firms from the mining industry, manufacture, or linked to agriculture – i.e., sectors that are shown in Appendix E to be far away from the AI knowledge

cluster – would change their technological trajectories due to digitalisation. Digital technologies will be used by these firms to innovate within their technological trajectories.

Additional evidence to confirm this assumption comes from the findings linked to the second hypothesis. This hypothesis suggests that the introduction of an AI innovation may have distinct effects across MNEs, depending on their knowledge distance to AI. I expected that MNEs from sectors that are close to AI's knowledge would increase their relatedness, whereas MNEs from sectors more distant to AI's knowledge would show no effect or decrease their relatedness after AI introduction. However, the results indicate that the effects are consistently positive across all technological sectors, the only difference being the intensity of the effect. But hypothesis (ii) considers solely the possible effects of AI as being introduced as a new piece of knowledge. The dimensions of knowledge can explain why the effects on relatedness are larger for firms that are closer to AI (i.e., the three dimensions are to play out, whereas for unrelated firms the only effects are to come from AI being used as a knowledge complementarity), but cannot explain how the introduction of unrelated knowledge leads to an immediate and significant increase in the production of related knowledge. However, the introduction of an AI innovation does not comprehend only the knowledge aspect. AI is also a technology that can be used to discover patterns in large amounts of digital data. This technological use is at the core of AI's expected IMI potential (Cockburn et al., 2018).

The fact that AI adoption is linked to an increase in firms' innovative performance⁷⁹ is additional evidence of AI's specific technological potential. This finding is associated with the third hypothesis presented. This hypothesis considers that AI can increase firms' innovative performance by being a valuable new piece of knowledge and due to its technological particularities. As a new piece of knowledge, AI increases firms' relatedness, which is associated in the geographic view of relatedness with higher innovative performance (Aarstad et al., 2016; Castaldi et al., 2015; Kim et al., 2016; Leten et al., 2007). AI's power as a technology, in turn, is created by its digital affordances and IMI potential. When used as a technology, AI can help firms to learn from digital data and to create "malleable" adaptable

⁷⁹ The gains in innovative performance linked to AI corroborate the findings presented in Rammer et al. (2021), which were based on survey data. Considering that AI is a type of digital technology, these gains in innovative performance contradict the findings presented Usai et al. (2021). The authors combine survey data about ICT usage in European firms and their innovation performance. One possible reason for these distinct results is that Usai et al. (2021) focus on the use of digital innovations, whereas I focus on the ownership of these innovations via intellectual property rights.

innovations suited to recombination to other contexts. But the ability to recognise which data and contexts are valuable to explore is dependent on firms' existing knowledge bases. The data available to a firm is also to be specific to the firm's current operations, which also determines what kind of innovations can be deployed⁸⁰. This specificity of data and the path dependence of firms in judging the value of opportunities explain why the introduction of AI is linked to an increase in knowledge relatedness. AI innovations are thus path dependent on existing knowledge bases: Firms innovate according to the data and accumulated knowledge they have.

Thereby, AI's technological potential to be a learning tool can create additional knowledge spillovers of related knowledge. Spillovers generated through learning processes are typically linked to the dimension of knowledge proximity (Breschi et al., 2003). But in the AI's case, these spillovers are generated from using AI as a technology, rather than due to the learning process generated through its absorption as a new piece of knowledge. This use of AI as a learning technology explains why the results on relatedness are consistent even across sectors that are unrelated to AI. This effect cannot be explained by other technological particularities of AI, as its digital aspects of generativity and convergence⁸¹.

This finding allows understanding the concept of absorptive capacity as an ability that can be changed by firms' technological choices. Despite the recognition in the literature that some technologies give firms specific benefits or advantages (as explored here in the case of digital technologies following the arguments presented in Yoo et al. (2012)), this possibility was not yet explored. MNEs that introduced AI extended their innovative activities also as a result of particular technological potentialities that AI offers. This is in addition to changes linked to learning about AI as a new piece of knowledge. Breschi et al. (2003) highlight that firms extend their innovative activities in knowledge-related domains as a consequence of their learning processes and due to specific features of knowledge and its links. This chapter adds to that

⁸⁰ Beraja et al. (2020) give an example of how deterministic data is to the emergence of AI innovations. The authors show that the availability of surveillance camera network data has a causal effect on the emergence of AI software linked to facial recognition.

⁸¹ Generativity stresses the role of unprompted change driven by large, varied, and uncoordinated actors. Uncoordinated actors are not likely to produce knowledge that fill the gaps from technologically distant firms. Convergence, in turn, refers to the action potential of embedding digital technologies in nondigital artefacts. This affordance is to have a negative effect on relatedness particularly for firms from sectors technologically distant from digital technologies. This technological distance would lead relatedness to be reduced at least in the first moments after AI introduction, when the firm is transforming its related nondigital artefacts into unrelated digital innovations. The dynamic effects for relatedness across distinct sectors show that this is not the case (See Appendix F, Figures F1 and F2).

view by showing that the adoption of some specific technologies may also impact firms' ability to explore knowledge-related domains. This finding is particularly interesting because it shows that even unrelated technologies can be used to create related knowledge.

Therefore, digitalisation seems to reinforce the aspect of knowledge production as being a cumulative, path dependent, and interactive process (Dosi, 1982; Nelson & Winter, 1982). In this context, AI seems associated with continuous technological change (Dosi, 1982) within the existing ICT technological paradigm, rather than a discontinuous change towards a new paradigm. The fact that AI reinforces existing technological trajectories also highlights that this technology is not being applied yet to its full potential as an IMI. AI is not being used to develop unrelated knowledge. This corroborates similar findings from Bianchini, Müller, and Pelletier (2020) regarding the potential use of AI as an IMI in science. The authors find that scientists use AI to explore scientific areas that are related to their knowledge domains. In the way that AI is being deployed (i.e., to strengthen technological trajectories), it also seems that firms are using this technology to enhance their existing competencies instead of replacing them (Paschen et al., 2020). Nevertheless, AI's potential to be a competence-destroyer technology can play out through competition between firms as AI enhances competences within firms. AI adopters may be using AI to deploy functions that replace what other firms offer.

4.5.2. Contributions

This chapter analysed the effects that affordances provided by digital technologies may have on firms' technological trajectories. For AI as an example of an advanced digital technology, I show that AI adoption is linked to a significant increase in the technological coherence and innovative performance of MNEs. Given the consistency of AI effects on relatedness across distinct sectors, I argue that sectors technologically distant from AI manage to increase their technological relatedness by using AI as a tool for learning rather than by just using AI as a new piece of knowledge. This has implications for the concept of absorptive capacity, linking it to a changing ability that can be influenced by firms' technological decisions. As a result, not only knowledge dimensions, as suggested by Breschi et al. (2003), but also technologies hold the potential to influence the emergence of relatedness.

4.5.3. Limitations and future research

There are several limitations I need to acknowledge: First, I consider inventions as a direct proxy of innovations. However, only a fraction of AI innovations is patented. Innovators may not patent their innovations for several reasons, e.g., to avoid disclosing technical aspects of their inventions and/or avoid copying, due to difficulties in proving the invention's novelty, due to prices related to patenting, etc. In addition, patent applications do not reflect activities by MNEs in open source software development, which might be a relevant aspect of AI development by firms. Second, I use patent applications rather than granted patents, which might introduce a quality bias. Third, there are missing values for two out of the three indicators considered for innovative performance in the dataset, which may compromise the analysis related to these particular indicators.

Future work is particularly needed – with a focus on alternative technologies – to understand whether the effects measured for AI are part of a broader aspect of how firms' technological trajectories change when they learn about a new technology, or not. In this sense, a detailed analysis of how specific technologies impact sectors from distinct knowledge distances would be particularly interesting. Future research could also focus on alternative data sources to measure innovation and on the possible effects of AI's adoption in the competition between firms (i.e., to understand AI's competence-destroyer potential).

References

- Aarstad, J., Kvitastein, O. A., & Jakobsen, S.-E. (2016). Related and unrelated variety as regional drivers of enterprise productivity and innovation: A multilevel study. *Research policy*, 45(4), 844-856. doi: <https://doi.org/10.1016/j.respol.2016.01.013>
- Alcácer, J., Cantwell, J., & Piscitello, L. (2016). Internationalization in the information age: A new era for places, firms, and international business networks? In: Springer.
- Antonelli, C., Krafft, J., & Quatraro, F. (2010). Recombinant knowledge and growth: The case of ICTs. *Structural Change and Economic Dynamics*, 21(1), 50-69. doi: <https://doi.org/10.1016/j.strueco.2009.12.001>
- Balland, P. A. (2016). Relatedness and the geography of innovation. In *Handbook on the geographies of innovation*: Edward Elgar Publishing.
- Balland, P. A. (2017). Economic Geography in R: Introduction to the EconGeo package. Available at SSRN 2962146. doi: <http://dx.doi.org/10.2139/ssrn.2962146>
- Beraja, M., Yang, D. Y., & Yuchtman, N. (2020). *Data-intensive innovation and the state: evidence from AI firms in China*.
- Bianchini, S., Müller, M., & Pelletier, P. (2020). Deep learning in science. *arXiv preprint arXiv:2009.01575*.
- Brantly Callaway, P. H. C. S. A. (2021). R Package Difference-in-Differences (update from 2021-05-07). Retrieved from <https://github.com/bcallaway11/did>
- Breschi, S., Lissoni, F., & Malerba, F. (2003). Knowledge-relatedness in firm technological diversification. *Research policy*, 32(1), 69-87. doi: [https://doi.org/10.1016/S0048-7333\(02\)00004-5](https://doi.org/10.1016/S0048-7333(02)00004-5)
- Brock, J. K.-U., & Von Wangenheim, F. (2019). Demystifying AI: What digital transformation leaders can teach you about realistic artificial intelligence. *California Management Review*, 61(4), 110-134.
- Callaway, B., & Sant'Anna, P. H. (2018). Difference-in-differences with multiple time periods and an application on the minimum wage and employment. *arXiv preprint arXiv:1803.09015*, 1-47.
- Callaway, B., & Sant'Anna, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of econometrics*.
- Cantwell, J., & Andersen, B. (1996). A statistical analysis of corporate technological leadership historically. *Economics of innovation and new technology*, 4(3), 211-234.
- Castaldi, C., Frenken, K., & Los, B. (2015). Related variety, unrelated variety and technological breakthroughs: an analysis of US state-level patenting. *Regional studies*, 49(5), 767-781. doi: <https://doi.org/10.1080/00343404.2014.940305>
- Cave, S., & ÓhÉigeartaigh, S. S. (2018). *An AI race for strategic advantage: rhetoric and risks*. Paper presented at the Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society.
- Cockburn, I. M., Henderson, R., & Stern, S. (2018). *The Impact of Artificial Intelligence on Innovation*.
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative science quarterly*, 128-152. doi: <https://doi.org/10.2307/2393553>
- Colombelli, A., Krafft, J., & Quatraro, F. (2013). Properties of knowledge base and firm survival: Evidence from a sample of French manufacturing firms. *Technological Forecasting and Social Change*, 80(8), 1469-1483.
- Delgado, M., Porter, M. E., & Stern, S. (2014). Clusters, convergence, and economic performance. *Research policy*, 43(10), 1785-1799.
- Dosi, G. (1982). Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change. *Research policy*, 11(3), 147-162. doi: [https://doi.org/10.1016/0048-7333\(82\)90016-6](https://doi.org/10.1016/0048-7333(82)90016-6)
- Eck, N. J. v., & Waltman, L. (2009). How to normalize cooccurrence data? An analysis of some well-known similarity measures. *Journal of the American society for information science and technology*, 60(8), 1635-1651. doi: <https://doi.org/10.1002/asi.21075>

- Ejdemo, T., & Örtqvist, D. (2020). Related variety as a driver of regional innovation and entrepreneurship: A moderated and mediated model with non-linear effects. *Research policy*, 49(7), 104073. doi: <https://doi.org/10.1016/j.respol.2020.104073>
- European Commission. (2017). *Harnessing the economic benefits of artificial intelligence*.
- Fujii, H., & Managi, S. (2018). Trends and priority shifts in artificial intelligence technology invention: A global patent analysis. *Economic Analysis and Policy*, 58, 60-69. doi: <https://doi.org/10.1016/j.eap.2017.12.006>
- Genz, S., Gregory, T., Janser, M., Lehmer, F., & Matthes, B. (2021). How do workers adjust when firms adopt new technologies? *ZEW-Centre for European Economic Research Discussion Paper* (21-073).
- Goralski, M. A., & Górnica-Kocikowska, K. (2020). Handling resultant unemployment from artificial intelligence. *J. Mark Munoz and Al Naqvi (eds.), Handbook of Artificial Intelligence and Robotic Process Automation: Policy and Government Applications*, 67-76.
- Granstrand, O. (1998). Towards a theory of the technology-based firm. *Research policy*, 27(5), 465-489.
- Hanelt, A., Firk, S., Hildebrandt, B., & Kolbe, L. M. (2021). Digital M&A, digital innovation, and firm performance: an empirical investigation. *European Journal of Information Systems*, 30(1), 3-26.
- Hidalgo, C. A. (2021). Economic complexity theory and applications. *Nature Reviews Physics*, 3(2), 92-113.
- Hidalgo, C. A., Klinger, B., Barabási, A.-L., & Hausmann, R. (2007). The product space conditions the development of nations. *science*, 317(5837), 482-487. doi: <https://doi.org/10.1126/science.1144581>
- Ho, D. E., Imai, K., King, G., & Stuart, E. A. (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political analysis*, 15(3), 199-236.
- Hollanders, H., & Esser, F. C. (2007). *Measuring innovation efficiency*: Citeseer.
- Huang, J., Henfridsson, O., Liu, M. J., & Newell, S. (2017). Growing on steroids: Rapidly scaling the user base of digital ventures through digital innovation. *Mis Quarterly*, 41(1).
- Iscan, E. (2021). An Old Problem in the New Era: Effects of Artificial Intelligence to Unemployment on the Way to Industry 5.0. *Journal of Yaşar University*, 16(61), 77-94.
- Ivarsson, I., Alvstam, C., & Vahlne, J.-E. (2015). Global technology development by collocating R&D and manufacturing: the case of Swedish manufacturing MNEs. *Industrial and corporate change*, dtw018.
- Khin, S., & Ho, T. C. (2019). Digital technology, digital capability and organisational performance: A mediating role of digital innovation. *International Journal of Innovation Science*.
- Kim, J., Lee, C.-Y., & Cho, Y. (2016). Technological diversification, core-technology competence, and firm growth. *Research policy*, 45(1), 113-124.
- Klinger, J., Mateos-Garcia, J., & Stathoulopoulos, K. (2018). Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology. doi: <https://arxiv.org/pdf/1808.06355.pdf>
- Kogler, D. F., Rigby, D. L., & Tucker, I. (2013). Mapping knowledge space and technological relatedness in US cities. *European Planning Studies*, 21(9), 1374-1391.
- Le Bas, C., & Sierra, C. (2002). 'Location versus home country advantages' in R&D activities: some further results on multinationals' locational strategies. *Research policy*, 31(4), 589-609.
- Leten, B., Belderbos, R., & Van Looy, B. (2007). Technological diversification, coherence, and performance of firms. *Journal of Product Innovation Management*, 24(6), 567-579.
- Leusin, M. E., Günther, J., Jindra, B., & Moehle, M. G. (2020). Patenting patterns in Artificial Intelligence: Identifying national and international breeding grounds. *World Patent Information*, 62, 101988. doi: <https://doi.org/10.1016/j.wpi.2020.101988>
- Nambisan, S., Lyytinen, K., Majchrzak, A., & Song, M. (2017). Digital Innovation Management: Reinventing innovation management research in a digital world. *Mis Quarterly*, 41(1).
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Cambridge, Massachusetts, and London, England: The Belknap Press of Harvard University Press.

- Nesta, L. (2008). Knowledge and productivity in the world's largest manufacturing corporations. *Journal of Economic Behavior & Organisation*, 67(3-4), 886-902.
- Nilsson, N. J. (2009). *The quest for artificial intelligence*: Cambridge University Press.
- Paschen, U., Pitt, C., & Kietzmann, J. (2020). Artificial intelligence: Building blocks and an innovation typology. *Business Horizons*, 63(2), 147-155.
- Rammer, C., Czarnitzki, D., & Fernández, G. P. (2021). Artificial intelligence and industrial innovation: Evidence from firm-level data. *ZEW-Centre for European Economic Research Discussion Paper* (21-036).
- Solheim, M. C., Boschma, R., & Herstad, S. (2018). *Related variety, unrelated variety and the novelty content of firm innovation in urban and non-urban locations*. Retrieved from Papers in Evolutionary Economic Geography (PEEG) 1836.
- Solheim, M. C., Boschma, R., & Herstad, S. J. (2020). Collected worker experiences and the novelty content of innovation. *Research policy*, 49(1), 103856.
- Stuart, E. A., King, G., Imai, K., & Ho, D. (2011). MatchIt: nonparametric preprocessing for parametric causal inference. *Journal of Statistical Software*.
- Su, G. (2018). Unemployment in the AI Age. *AI Matters*, 3(4), 35-43.
- Teece, D. J. (2018). Profiting from innovation in the digital economy: Enabling technologies, standards, and licensing models in the wireless world. *Research policy*, 47(8), 1367-1387. doi: <https://doi.org/10.1016/j.respol.2017.01.015>
- Teece, D. J., Rumelt, R., Dosi, G., & Winter, S. (1994). Understanding corporate coherence: Theory and evidence. *Journal of Economic Behavior & Organisation*, 23(1), 1-30.
- Tseng, C.-Y., & Ting, P.-H. (2013). Patent analysis for technology development of artificial intelligence: A country-level comparative study. *Innovation*, 15(4), 463-475.
- Weber, C., & Weber, B. (2010). Social capital and knowledge relatedness as promoters of organisational performance: An explorative study of corporate venture capital activity. *International Studies of Management & Organisation*, 40(3), 23-49.
- WIPO. (2019). *WIPO Technology Trends 2019: Artificial Intelligence*. Retrieved from https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf
- Yoo, Y., Boland Jr, R. J., Lyytinen, K., & Majchrzak, A. (2012). Organizing for innovation in the digitized world. *Organisation science*, 23(5), 1398-1408. doi: <https://doi.org/10.1287/orsc.1120.0771>
- Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Dinlersoz, E. (2021). *Advanced Technologies Adoption and Use by US Firms: Evidence from the Annual Business Survey*.

Appendix A: Constructing a dataset combining Patents and Multinational Enterprises data (2011-2019).

The creation of the dataset presented in this Chapter is a joint work between Björn Jindra, Felix Lüders, and Matheus E. Leusin.

The creation and matching of the firm and patent data consist of four main steps, which can be summarised as follows:

- Step I: Construct a dataset of corporate Global Ultimate Owners and their subsidiaries as of 2020;
- Step II: Recreate ownership structures of the past 10 years by adding M&A data and incorporation dates;
- Step III: Create a patent dataset containing the historical ownership of patents;
- Step IV: Matching of firm- and patent data and creation of additional firm-level variables;

The construction of the complete dataset was developed in R. Each of the steps carried in the dataset development is detailed in the subsections below.

Step I: Creating a dataset of corporate Global Ultimate Owners and their subsidiaries as of 2020

This step combines a bottom-up identification of Multinational Enterprises (MNEs) on Orbis, followed by a top-down approach used to download data to complement the coverage of the companies considered. Both approaches are based on data downloaded from Orbis between 17/07/2020 and 25/01/2021. In the bottom-up approach, foreign subsidiaries are linked to their immediate parents and, ultimately, their respective global ultimate owners (GUOs). All subsidiaries with foreign shareholders were extracted from Orbis with an ownership criterion of 25.01%. The files were downloaded in the following format:

Set 1: Files "GUO" - 2,803,941 companies: Subsidiaries and related ownership data

Company name	Latin	MNE-BvD Number	ID	GUO_BvD Number	ID	ISH_BvD ID Number (Immediate parent)
Siemens GmbH		DE10001				
Flender GmbH		DE10012		DE10001		LU10002
Siemens Mechanical Drives GmbH		LU10002		DE10001		DE10011

Building on this data structure, a vlookup function⁸² was used to extend the ownership levels upwards (look for subsidiaries that are parents of other subsidiaries) and downwards (look for parents that are subsidiaries of other subsidiaries). This resulted in an ownership network structure for each of the subsidiaries consisting of a maximum of 21 ownership levels. The search for more subsidiaries/parents stopped when there were no more subsidiaries of a company and when the parent equalled the GUO. This leads to a first outcome file with the following structure:

⁸² This vlookup function is applied through the expss package. It looks for a vector of looked up values in a reference column of a data frame, returning the corresponding values from this data frame (<https://cran.r-project.org/web/packages/expss/expss.pdf>). In the current implementation, the vector values are a list of individual BvD ID Numbers.

Outcome 1: "Interim1" / 2,609,242 companies: Ownership Network of MNEs (bottom-up)

GUO	Own 1	Own 2	Own 3	Own 4	Own 5	Own 6	Own7
DE10001		DE10011	LU10002	DE10012			
DE10001			DE10011	LU10002	DE10012		
DE10002		DK10003	DE10021	DE10022			
DE10002			DK10003	DE10021	DE10022	DE10023	

As this method only allows the connection of subsidiaries across borders and not all subsidiaries might be captured, this initial dataset was then complemented with a top-down search for shareholders with foreign subsidiaries. Thus, the shareholders and all their subsidiaries were downloaded in the following format:

Set 2: Files "Shareholders" - 3,152,693 companies: Shareholders and related subsidiary data

Company name	Latin alphabet	Subsidiary-BvD Number	ID	GUO_BvD ID Number	Shareholder_BvD Number	ID
Siemens GmbH		DE10001				
Flender GmbH		DE10012		DE10001	LU10002	
Siemens Mechanical Drives GmbH		LU10002		DE10001	DE10011	

By treating the shareholder as the parent company, a vlookup function was then used to establish the ownership levels of each company, leading to a second Outcome file with the following structure:

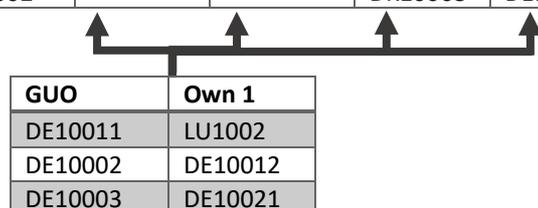
Outcome 2: "Interim2" / 1,752,876 companies: Ownership Network of MNEs (top-down)

GUO	Own 1	Own 2	Own 3	Own 4	Own 5	Own 6	Own7
DE10001		DE10011	LU10002	DE10012			
DE10001			DE10011	LU10002	DE10012		
DE10002		DK10003	DE10021	DE10022			
DE10002			DK10003	DE10021	DE10022	DE10023	

As both data structures are incompatible to match due to differences in the final ownership structures, the bottom-up data frame is used as a basis, which is then complemented with the shareholder level data. Therefore, a data string of all shareholders and subsidiary relations across all ownership level was established and merged to each ownership level of Interim 1. An overview of how this Structure-based output is created is presented below.

Match 1: "Structure" / 3,212,537 companies: Ownership Network of multinational enterprises

GUO	Own 1	Own 2	Own 3	Own 4	Own 5	Own 6	Own7
DE10001		DE10011	LU10002	DE10012			
DE10001			DE10011	LU10002	DE10012		
DE10002		DK10003	DE10021	DE10022			
DE10002			DK10003	DE10021	DE10022	DE10023	



After removing duplicates and empty columns, this new dataset contains data from all shareholders and all of their subsidiaries – domestic and foreign –, which add up to 21 ownership levels for the longer firm's ownership structure. As the focus is on identifying corporate GUOs only, the original bottom-up Orbis search was modified by filtering for different types of GUOs. This set (Set 3: Files "Diverse" - 2.803.941 companies) had the same structure as the first set but was bundled by the types of GUOS including Banks, Insurance, Private Equity Funds, Hedge Funds, Mutual and Private Equity

Funds, Research Institutes, Public Agencies, Individuals, Employees, and Corporations. By loading these datasets into R and matching them with the ownership structure dataset, it was possible to create a “type” variable for each company. This led to a dataset with all multinational enterprises, their global structure and a column indicating non-corporate GUOs, with the structure presented below.

Outcome 3: “Type” / 2,609,982 companies: Ownership Network of MNEs with Type

GUO	Own 1	Own 2	Own 3	Own 4	Own 5	Own 6	Type
DE10001		DE10011	LU10002	DE10012			
DE10001			DE10011	LU10002	DE10012		
DE10002		DK10003	DE10021	DE10022			
DE10002			DK10003	DE10021	DE10022	DE10023	

To avoid non-corporate GUOs in the dataset, ownership structures were moved down if the GUO was not a corporate company. This means that states or holdings having an ownership share in several corporations were excluded and corporations were made the highest entity in the corporate network. This resulted in a dataset of 342,787 corporate GUOs representing the top node of the ownership structure as of 2020.

Step II: Recreating ownership structures of the past 10 years by adding M&As and dates of incorporation

After establishing the corporate network ownership structure, two data frames were constructed to allow extending the 2020 ownership structure backwards with additional M&A historical data. First, one data frame consisting of all GUOs in one column and all subsidiaries in the second column (as well as one duplicate of the GUO) was created. Based on this structure, GUOs without subsidiaries or subsidiaries located only in the home country were removed, leaving 280,206 unique GUOs and 1,815,997 subsidiaries. This data serves as the basis for matching all data to all companies in the dataset while still allowing the aggregation to the (corporate) GUO level. To do so, the dataset was split into many GUO-ownership-level data frames, recombined, and cleaned for duplicates. Second, a two-column string of all company subsidiary relations was established by splitting the corporate network data across all ownership levels. This helps defining the subsidiaries of every single nod, which is important if a company (and its subsidiaries) is acquired by another firm. The structure of these two files is presented below.

Outcome 4: “Own” (2,096,203) & “Subs” (3,770,613): Company strings

GUO	Subsidiaries			Company	Subsidiaries
DE10011	LU1002			DE10011	LU1002
DE10002	DE10012			LU1002	DE10012
DE10003	DE10021			DE10003	DE10021

To establish past ownership structure, all M&A deals that relate to the GUOs or subsidiaries as of 2020 were downloaded from the Orbis-Zephyr database. These files included deal and company characteristics, particularly the target, vendor and acquiror information, and were downloaded in the following format:

Set 4: Files "DealsX" - 486,770 transactions: Mergers and acquisitions (since 2010) + company data

Deal Number	Target BvDID	Acquiror BvDID	Vendor BvDID	Deal Type	Deal date	Deal value
351998	DE10015	DE10001	LU10002	Acquisition 50%	2011	n.a.
352048	DE10015	DE10011	UK10001	Acquisition 50%	2015	n.a.
352061	LU10004	LU10002	LU10003	Minority Stake unknown%	2014	n.a.

After excluding deals that are not relevant to ownership changes (unknown % or acquisitions below 25%), 32.334 deals were left for the period between 2010 and 2020. To calculate the ownership structure backwards, subsidiaries that were sold in year t were added to the ownership structure in year t-1. This was done to identify companies that only belonged to the ownership structure in previous years and to capture all subsidiaries that belonged to the GUO during the examination period. This structure was treated as a basis of 2010 to adjust the ownership structure through forward calculation in three main ways:

1. Merge by target: If the subsidiary as of 2010 was target of a deal in 2011, a vlookup function is used to Check if the acquiror of the target may still be in the ownership group according to the 2020 dataset. If the acquiror's GUO was not the same as the target GUO (indicating an inter-company acquisition), the company-subsiary string was matched to detect all companies related to the acquired company. After re-checking if the acquiror belonged to the corporate group, all targets and their related subsidiaries were deleted from the 2010 structure to arrive to a new the 2011 structure.
2. Merge by acquiror: If a company within the GUOs ownership structure of 2010 acquired a new subsidiary in 2011, the target company was added to the ownership structure in 2011, Therefore, the match by acquiror displayed all targets that were added by the company in 2011. Then, the company-subsiary string was added to identify all subsidiaries of the acquired company. Then, if the acquiror of the targets belonged to the same ownership group, all targets and their subsidiaries were added to the new corporate group of the acquiror.
3. Merge by vendor: If a company within the GUOs ownership structure of 2011 (resulting string of the previous two steps) sold a subsidiary in 2012, the target company and its subsidiaries were added post hoc to the ownership structure in 2011.

In this way, the ownership structures were calculated from 2011 to 2019 and the company-subsiary data was also updated synchronously to generate a complete picture of ownership structures based on their acquisition deals. The validity of the process is also demonstrated by the consistency of the original Orbis structure and the resulting 2020 ownership structure as there is only a minor difference of 40.000 companies, which can be explained with lacking updates of the Orbis data and deletion of companies with a 25.01% ownership change. The following illustration shows the process of removing and adding new subsidiaries.

Match 2: "Own" + "Deals": Matching ownership structures with M&A data

BvDID	Subsidiaries
DE10001	DE10001
DE10001	DE10011
DE10001	DE10012
DE10001	DE10013
LU10001	LU10001
LU10001	LU10002

Structure 2010
(Basis for match)

Company	Subsidiaries
DE10001	DE10001
DE10001	DE10011
DE10001	DE10012
DE10001	DE10013
DE10013	DE10011
LU10001	LU10001
LU10001	LU10002
LU10002	LU10003

Match by target

BvDID	Subsidiaries	Acquiror	GUO1	Subsidiaries	Acquiror	GUO2	Delete
DE10001	DE10001						
DE10001	DE10011	DE10012	DE10001				
DE10001	DE10012				LU10001	LU10001	1
DE10001	DE10013	LU10001	LU10001				1
LU10001	LU10001						
LU10001	LU10002	DE10012	DE10001	LU10003			1
LU10001	LU10003				DE10012	DE10001	1

Match by acquiror

BvDID	Subsidiaries	Target	Subsidiaries	Acquiror GUO	Add	Delete
DE10001	DE10001					
DE10001	DE10011			DE10001		
DE10001	DE10012	DE10011		DE10001		
DE10001	DE10012	LU10002	LU10003	DE10001	1	
DE10001	DE10013	LU10001	LU10001			1
LU10001	LU10001	DE10013	DE10011		1	
LU10001	LU10002					1
LU10001	LU10003					1

Match by vendor in 2020

BvDID	Subsidiaries
DE10001	DE10001
DE10001	DE10012
DE10001	DE10013
DE10001	LU10002
DE10001	LU10003
LU10001	LU10001
LU10001	DE10011

BvDID	Subsidiaries	Target
DE10001	DE10001	DE10014
DE10001	DE10012	
DE10001	DE10013	
DE10001	LU10002	
DE10001	LU10003	
LU10001	LU10001	
LU10001	DE10011	AD10011

BvDID	Subsidiaries
DE10001	DE10001
DE10001	DE10012
DE10001	DE10013
DE10001	LU10002
DE10001	DE10014
DE10001	LU10001
LU10001	LU10001
LU10001	DE10013
LU10001	AD10011

Structure 2011

This process resulted in nine separate ownership files (one for each year considered between 2011-2019), which were further refined through the inclusion of the date of incorporation. The data of incorporation allows differentiating between the foundation of new subsidiaries (greenfield investment) from acquisitions (brownfield investment). This additional data was downloaded in the following format:

Set 5: Files “Date_data” - 2.463.340 (605,532 NAs) companies: Date of incorporation data

Company name Latin alphabet	BvD ID Number	Date of incorporation
Siemens GmbH	DE10015	08/12/2014
Flender GmbH	DE10015	23/06/2005
Siemens Mechanical Drives GmbH	LU10004	12/03/1956

This data was matched to each yearly dataset and companies that were founded after the last considered year were excluded. Plus, companies that were founded in the specific year of interest were indicated with a dummy variable. The resulting dataset has the following structure:

Match 3: “Full20xx” + “Date_data”: Matching companies and their dates of incorporation

BvDID	Subsidiaries		BvDID	Subsidiaries	Greenfield
DE10001	DE10001	→ Match by BvDID	DE10001	08/12/2020	0
	DE10011		DE10011	23/06/2011	1
	DE10012		DE10012	n.a.	0

The changes of the forward ownership calculation and date match are summarised in Tables A1.1 and A1.2.

Table A1.1: Forward Calculation of ownership structures (Original 2020 value 2,096,203)

Year	Companies	By vendor	Year	Companies	By target	By acquiror	By vendor
'20	2,096,203	/	'10	2,096,242			
'19	2,096,252	663 (2,999)	'11	2,089,584	1,228 (2,995)	2,272 (3,597)	1,327 (4)
'18	2,096,254	995 (3,297)	'12	2,086,754	1,299 (2,838)	2,290 (3,248)	1,566 (18)
'17	2,096,262	972 (3,857)	'13	2,084,039	1,363 (1,229)	2,331 (2,897)	1,701 (91)
'16	2,096,258	987 (4,154)	'14	2,080,873	1,764 (2,958)	2,921 (4,018)	3,480 (138)
'15	2,096,265	915 (3,542)	'15	2,075,449	2,170 (3,089)	3,281 (6,593)	1,563 (218)
'14	2,096,225	1,034 (3,783)	'16	2,071,648	2,035 (2,722)	2,939 (9,206)	1,624 (158)
'13	2,096,229	881 (3,313)	'17	2,066,043	2,138 (3,740)	3,025 (5,263)	1,657 (48)
'12	2,096,239	746 (2,517)	'18	2,060,179	2,133 (4,124)	2,947 (7,848)	2,004 (73)
'11	2,096,239	707 (2,277)	'19	2,055,048	1,937 (4,409)	2,698 (8,543)	1,309 (134)
'10	2,096,242	704 (2,253)	'20	2,055,977	1,223 (3,013)	1,694(3,773)	/

Table A1.2: Forward Calculation of ownership structures (New 2020 value: 2.055.977)

Year	'11	'12	'13	'14	'15	'16	'17	'18	'19
Companies	1,649,863	1,696,713	1,747,034	1,798,762	1,852,586	1,907,127	1,958,278	2,006,336	2,044,415
Difference	+46,850	+50,321	+51,728	+53,824	+54,541	+51,151	+48,058	+38,080	+51,728

Step III: Creating a dataset of patent applications with historical ownership data

One additional dataset with patent data applications was collected from Orbis IP. The development of the patent dataset is divided into five steps, which refer to the creation of a general patent dataset (step 1), creation of a complementary dataset containing additional information about patents with changes in ownership (step 2), creation of an additional AI tag to identify AI patents (step 3), filtering and merging of these three datasets (step 4), and the slicing of the final patent dataset into yearly data (step 5). The details of each step are presented below.

Creating a complete dataset of priority patents

First a dataset based on all patent filings (granted or not) whose priority dates were within the period between 01/01/2009 and 31/12/2019 is created. Particularly, the year 2020 was excluded due to the

existing lag in the application process of patents. The initial 45,183,514 patents attending this criterion were further filtered for priorities. Orbis IP allows filtering priorities through the “Families” option, where the filter “Most representative family member only” can be applied. This filter reduces the original dataset of patents from 45,183,514 registers to 27,415,475 priorities registered at the date of the search (11/11/2020). It is still worth highlighting that the option ““Complete” family (DOCDB definition - extended approach)” is the one chosen (conversely to the option ““Simple” family (invention based)”. These 27,415,475 priorities were downloaded in the following format:

Set 6: Files “Priorities” - 27,415,475 priority filings: Patents applications with priority dates (2009-2019)

Publication number	Priority date	Current owner(s) BvD ID	direct BvD ID	Applicant(s) BvD ID Number(s)	IPC code (main)	IPC code (others)	Granted
AP2016094D0	02/09/2015	HK0022459176		HK0022459176	n.a.	n.a.	
AM170U	05/06/2009				n.a.	n.a.	Yes
AT12187U3	19/04/2009	DE2190544304		DE2190544304	E04H12/04	...	Yes
BG111715A	05/03/2014				D01H4/02	...	

Creating a complete dataset of priority patents with change in ownership

In the second step, additional patent information for the patents that had a change in ownership after the first considered year of 2009 is downloaded. There are 773,721 patents that attend to those criteria, and they were downloaded in the following format:

Set 7: Files “OwnershipChange” - 773,721 Priority filings: All patent applications with a change in ownership (2009-2019)

Publication number	Current owner(s) BvD ID	Previous owner(s) BvD ID	Vendor BvD ID Number	Acquiror BvD ID Number	Transaction date
AT12626U1	CHCHE10592			US133668641	16/10/2018
				CHCHE1059295	16/10/2018
AT12628U1	AT91108193	AT91105013	AT91105013	AT9110819336	23/10/2014
AU201124213	FI10942595	AU096768704	FI15380325	SE5565878054	13/07/2020
		SE5565878054	FI15380325	FI120464	13/08/2020
		FI120464	AU0967687	FI15380325	07/02/2019

Due to some issues with ambiguous and duplicated data in Orbis IP, additional filterers and checks were employed to separate irrelevant data from the relevant transactions useful for the analysis. In short, four filters were applied for cleaning the data: i) one for identifying changes where the current owner⁸³ is also the acquirer of the patent or the applicant (if the vendor information is not available); ii) one for counting the number of changes that actually happened, based on the dates (i.e., excluding the possibility of several changes occurring in the same day); iii) one for selecting which of the repeated data should be selected (based on the acquiror/vendor and owner/applicant information); and iv) one for identifying different owners and applicants, so that data about any of these companies is not missed. After applying these filters and dropping irrelevant data, a total of 694,859 patents are identified as having at least one ownership change in the considered period. These are further filtered for priorities and expired patents by merging this dataset to the priority’s dataset, created in the last step (unmatched registers are dropped). The resulting dataset contains information about ownership changes for 429,227 priority patents.

⁸³ It is worth mentioning that Orbis IP allows for the possibility of a patent having more than one owner, which is also considered in the creation of the patent dataset presented here.

Creating an “AI” tag

As the current chapter has a particular focus on AI technologies, these are additionally tagged in the patent dataset. The identification of these patents follows the keyword-based search proposed in Leusin et al. (2020). The list of keywords considered is also presented in Table A1.3.

Table A1.3: Keywords proposed in Leusin et al. (2020) for identifying AI patents

AI Technique Keyword	Additional Wikipedia Synonym Keyword
%neural network%	No additional synonym
%machine learn%	No additional synonym
%artificial intelligence%	%machine intelligen%
%expert system%	No additional synonym
%support vector machin%	%support vector network%
%fuzzy logic%	No additional synonym
%graphical model%	%structured probabilistic model%
%pervised learn%	No additional synonym
%deep learn%	%deep structured learn% and %hierarchical learn%
%classification tree% OR %regression tree%	%decision tree learn%
%reinforced learn%	%reinforcement learn%
%logic programming%	No additional synonym
%rule learn%	%rule induction%
%probabilistic reason%	%probability logic% and %probabilistic logic%
%task learn%	No additional synonym
%logical learn% OR %relational learn%	No additional synonym
%latent represent%	No additional synonym
%bio-inspired approach%	%bio-inspired comput% and %biologically inspired comput%
%instance-based learn%	%memory-based learn%
%ontology engineer%	No additional synonym
%description logistic%	Keyword not found

Conversely to Leusin et al. (2020), which identifies as AI patents all registers containing any of the mentioned AI-keywords in the titles or abstracts, also the claims and description of the patents are considered in this chapter. These criteria resulted in 717,860 identified AI patents, which are further filtered for priorities, resulting in a final “AI tag” added to 440,698 priorities.

Filtering and merging patent datasets

Next, the patents with a change in ownership are extracted from the general comprehensive priorities’ dataset. Priorities without any information about current owners are also excluded from the main dataset, remaining 17,581,241 priorities (and 429,227 priorities in the unchanged “Dataset with changes in ownership”). Both datasets are further filtered for withdrawn or abandoned patents by dropping the ones containing the status of “Expired” in Orbis IP. Considering the usual expiration time of 20 years and the time frame considered between 2009 and 2019, no expired patents should be found in the considered period, meaning that applications with this status were very likely withdrawn or abandoned by the applying companies. Therefore, they don’t represent any significant innovation results and are excluded from the patent dataset.

To allow for the direct matching of the patent data with the MNEs’ ownership structure later on, a structure with the yearly ownership data information of patents throughout each year of the considered period is created. For patents without any change in ownership, the owner information is

repeated for each year and each owner for the number of times that the patent has existed in the considered period. For the dataset with changes in ownership, distinct criteria for selecting the owners of the patents in distinct periods of time are applied, based on the dates in which the changes in ownership occurred. More precisely, until the first change in ownership, the owners of the patents are assumed to be the applicants of the patent. For the period after the last change in ownership of a patent, the owners of the patent are defined to be the current owners of the patents. For all changes of ownership occurring between these two periods, the owners of the patents are defined based on the "Acquiror BvD ID Number" information of each ownership change transaction. In addition, if there is missing information about applicants, owners or acquirors in the years that this information is used for defining the owners, the patents are not dropped but just stay without an owner until the next transaction without missing data happens.

Finally, these two datasets were merged with the "AI tag" using the Publication number information. In this way, AI patents gain additional identification information to separate them from the rest of the patents. The final dataset has the following structure:

Outcome 6: "Patents" / 1,527,441 companies: Ownership Structures with firm-and country variables

Publication Number	Curr. Owner BvDID	Prior Year	Diff 1	Owner	Cur. Year	IPC main	IPC other	Grant	No of family mem.	AI tag
CN2082366U	CN9387053	2018	2	CN9387053	2019	NA	NA	NA	NA	NA
CN2082366U	CN9387053	2018	2	CN9387053	2018	NA	NA	NA	NA	NA
CN2082367U	CN31995PC	2018	2	CN31995PC	2019	NA	NA	NA	NA	NA
CN2082367U	CN31995PC	2018	2	CN31995PC	2018	NA	NA	NA	NA	NA
CN2082375U	CN9371220	2017	3	CN9371220	2019	NA	NA	NA	NA	NA

"Diff1" refers to how many years the patent existed in the period 2009-2019. "CurrYear" refers to each year of "existence" of the patent in this period. The column "Owner" refers to the owner of the patent in a given year. Finally, for AI patents, the last column, namely "AI", indicates if the register is an AI patents ("Yes") or not (NA). Descriptive information about the generated patent ownership data is presented in Table A1.4.

Table A1.4: Characteristics of the patent dataset

Description	Full Dataset	
Total number of priority patents for the period 2009-2019 including withdraw/abandoned patents	18,010,468	
Detailed results:	Considering the period 2011-2019:	
Total number of non-expired priority patents	15,735,583	100%
Total number of non-expired priority patents with no changes in ownership	15,438,678	98.1%
Total number of non-expired priority patents with no changes in ownership and with just one owner	14,582,343	92.7%
Total number of non-expired priority patents with no changes in ownership and with more than one owner	856,335	5.4%
Total number of non-expired priority patents with changes in ownership	296,905	1.9%

Finally, to complement the patent stock, additional patent data is downloaded. But this time, the data is downloaded focusing on the relevant BvD IDs identified. This means that the ownership structure of GUOs as of 2010 is used as input to collect additional patents that were owned by these companies before the year 2011. This additional patent stock data was downloaded in 02/03/2021. In total, 13,545,086 additional priorities were identified. All these patents (plus the patents from 2009 to 2010) were aggregated to the 2010 ownership structure, which served as the basis for the patent stock variables.

Slicing the Patent dataset into yearly data and creation of additional variables

The “Patent dataset” created in the previous step is further separated by year into 10 slices, which covers a stock file (with data from 2000 to 2010) plus yearly ownership information from 2011 to 2019. The created variable “Owner” is used to create the following additional variables: Sum of patents (NoPat), sum of new patents per year (App), sum of AI patents (AINoPat), sum of new AI patents per year (AInew) These variables are aggregated to the owner and include all MNEs and non-MNEs that patented in the examination period. The yearly dataset has the following structure:

Outcome 7: “Full11” / 1.527.441 companies: Patent dataset for 2011

BvDID Owner	Year	No Pat	App	AINoPat	AIApp	AInew
GB000414	2011	254	0	0	0	0
UG800100	2011	0	0	0	0	0
GB054441	2011	0	0	0	0	0

Step IV: Matching corporate and Patent data, and download of additional variables

Finally, the patent data was matched with the GUO-subsidary data based on the owners BvD IDs, so that the patent-related variables per company were attached to the different entities of the MNEs. As the main interest is on the top node of the MNE hierarchy, the patent variables were aggregated up to the corporate GUO. Then, the list of GUOs considered was uploaded to Orbis as a search criterion, so that additional specific firm-level variables could be downloaded for each MNE. The following additional variables were downloaded from Orbis: “NACE Rev. 2, core code (4 digits)”, “Country ISO code”, “Size classification” (which follows Orbis classification), “Operating revenue (Turnover) th USD” (one column of data for each year within the period 2011-2019), “R&D expenses/Operating revenue (%)” (also one additional column of data for each year within the period 2011-2019), in addition to the

variables “age” and “Type of entity”, which were already downloaded in the previous steps. This additional data was downloaded on 26/03/2021.

To further reduce the sample, exclusion criteria on the number of total patents (<1), foreign subsidiaries (<1), employees (<50.000) and revenue (<50.000.000) were applied. These criteria comply with the definition of large companies according to the European Commission. This step also ensures the comparability of firms in the dataset. General information about the complete dataset is presented on Appendix B.

Appendix B: General descriptives of the complete dataset.

Indicator	Year								
	2011	2012	2013	2014	2015	2016	2017	2018	2019
General information about number of GUOs and their subsidiaries									
Total number of GUOs considered	30,228	31,551	32,756	33,806	34,992	36,210	37,279	38,103	38,168
Total number of subsidiaries	785,931	817,887	848,821	880,528	911,233	943,850	974,776	1,003,895	1,016,299
Total number of subsidiaries with at least one patent (and share in relation to total number of subsidiaries)	59,067 (7.5%)	62,653 (7.7%)	66,040 (7.8%)	69,078 (7.8%)	72,121 (7.9%)	75,146 (8.0%)	77,964 (8.0%)	80,325 (8.0%)	80,857 (8.0%)
General information about patents									
Total number of patents identified	6,026,034	6,590,477	7,155,959	7,739,903	8,370,595	9,001,192	9,571,261	10,168,201	10,303,302
Total number of patents owned by GUOs	3,737,247	4,018,636	4,285,994	4,554,733	4,820,442	5,092,258	5,302,449	5,522,523	5,479,364
Total number of patents owned by subsidiaries (and share in relation to total)	2,288,787 (38%)	2,571,841 (39%)	2,869,965 (40%)	3,185,170 (41%)	3,550,153 (42%)	3,908,934 (43%)	4,268,812 (45%)	4,645,678 (46%)	4,823,938 (47%)
Total number of GUOs with direct ownership of patents	17,165 (57%)	17,870 (57%)	18,557 (57%)	19,154 (57%)	19,775 (57%)	20,436 (56%)	21,010 (56%)	21,401 (56%)	21,399 (56%)
Total number of GUOs with direct or indirect ownership of patents	18,760 (62%)	19,692 (62%)	20,513 (63%)	21,253 (63%)	22,075 (63%)	22,903 (63%)	23,660 (63%)	24,242 (64%)	24,335 (64%)
Information about AI patenting									
Total number of GUOs without AI patents	28,523	29,656	30,675	31,520	32,472	33,405	34,101	34,370	34,271
Total number of GUOs with at least one AI patent (and share related to the total number of GUOs with direct or indirect ownership of patents)	1,705 (5.6%)	1,895 (6%)	2,081 (6.4%)	2,286 (6.8%)	2,520 (7.2%)	2,805 (7.7%)	3,178 (8.5%)	3,733 (9.8%)	3,897 (10.2%)

Total number of AI patents identified owned by the considered GUOs	15,265	17,793	21,068	25,509	32,224	41,625	56,131	79,469	91,742
Total number of AI patents owned at the GUO level	9,247	10,448	11,912	13,910	16,511	20,603	26,766	35,823	39,193
Total number of AI patents owned at the subsidiary level (and share in relation to total)	6,018 (39%)	7,345 (41%)	9,156 (43%)	11,599 (45%)	15,713 (49%)	21,022 (51%)	29,365 (52%)	43,646 (55%)	52,549 (57%)

Note: Both total number of AI patents accumulated and total number of AI patents identified consider an initial stock from 2000 to 2010. In particular, the role of subsidiaries in patenting seems to grow over time, although the number of subsidiaries that create at least one patent changes just marginally. The number of patents owned by subsidiaries related to the considered GUOs increases from 38% in 2011 to 47% in 2019, whereas the total number of subsidiaries with at least one patent changes from 7.5% to 8.0%. When AI patents are considered separately, the increase in subsidiaries' relevance is even more expressive: Their ownership share grow from owning 39% of the AI patents related to the GUOs in 2011 to a share of 57% in 2019. Note the difference between the three indicators presented in the table: The indicator named "Total number of GUOs considered" considers the accumulated number of Global Ultimate Owners (GUOs) that appear in my dataset. It comprehends both companies that still exist and companies that are no longer active (due to mergers or other reasons). The two indicators "Total number of GUOs with direct ownership of patents" and "Total number of GUOs with direct or indirect ownership of patents", in turn, allow differentiating the number of active companies in a given year. This differentiation is possible through the structure of my dataset: As patents "follow" their previous owners once they are acquired by other companies, a GUO that has no ownership over any patent is not active or was merged to another GUO. Finally, the higher values seen for "Total number of GUOs with direct or indirect ownership of patents" over "Total number of GUOs with direct ownership of patents" mean that in some cases the GUOs didn't register any patent in the considered period, although subsidiaries related to them did.

Appendix C: Data about 30 larger AI adopters in 2016 (Table C1) and 2019 (Table C2).

No.	Name of the GUO	No. of patents owned	No. of AI patents owned	Country
1	Microsoft Corporation	49,912	1,754	US
2	Samsung Electronics Co., Ltd.	65,999	1,343	KR
3	International Business Machines Corp	85,963	1,183	US
4	Alphabet Inc.	31,602	970	US
5	Softbank Group Corp	15,608	884	JP
6	Siemens Ag	52,783	845	DE
7	Intel Corp	50,365	677	US
8	Sony Corporation	92,139	555	JP
9	Qualcomm Inc	43,850	534	US
10	Koninklijke Philips N.V.	33,878	493	NL
11	General Electric Company	47,308	488	US
12	Hitachi Ltd	144,544	455	JP
13	Fujitsu Limited	78,317	440	JP
14	Nec Corporation	69,580	389	JP
15	Nokia Oyj	45,775	385	FI
16	Naspers Limited	10,680	385	ZA
17	Toshiba Corporation	139,962	344	JP
18	Halliburton Co	12,568	344	US
19	Panasonic Corporation	218,691	314	JP
20	Mathworks Inc	1,004	305	US
21	Mitsubishi Electric Corporation	105,359	297	JP
22	Canon Incorporated	165,658	287	JP
23	Ford Motor Co	19,176	269	US
24	Facebook, Inc.	5,642	260	US
25	Ericsson (Telefonaktiebolaget NI) Ab	33,658	247	SE
26	Zte Corporation	37,475	241	CN
27	Hp Inc.	30,430	238	US
28	Robert Bosch Industrietreuhand Kommanditgesellschaft	40,219	212	DE
29	Boeing Company (The)	14,949	187	US
30	General Motors Company	25,382	173	US

Table C1: Main AI adopters identified in 2016. In comparison to the results presented in Fujii and Managi (2018), Table 3, one can see significantly higher numbers here. All top 10 AI adopters presented in Fujii and Managi (2018) are also seen here, although in slightly distinct positions (these top 10 are highlighted in the above table). Please note that Alphabet is the parent company of Google. When the top 15 are considered, Yahoo and D-Wave don't appear here. When the top 20 are considered, SAP and Xerox don't appear.

No.	Name of the GUO	No. of patents owned	No. of AI patents owned	Country
1	Softbank Group Corp	29,727	3,137	JP
2	Naspers Limited	20,529	2,752	ZA
3	Microsoft Corporation	51,903	2,719	US
4	Samsung Electronics Co., Ltd.	82,699	2,638	KR
5	Siemens Ag	59,342	1,747	DE
6	Alphabet Inc.	33,012	1,736	US
7	Intel Corp	52,849	1,524	US
8	International Business Machines Corp	77,273	1,513	US
9	Koninklijke Philips N.V.	35,334	983	NL
10	Sony Corporation	93,803	972	JP
11	Lg Corp.	102,404	889	KR
12	Fujitsu Limited	83,578	833	JP
13	Hitachi Ltd	157,540	813	JP
14	General Electric Company	48,639	778	US
15	Nokia Oyj	48,228	750	FI
16	Gree Electric Appliances, Inc. Of Zhuhai	51,911	742	CN
17	Qualcomm Inc	47,882	724	US
18	Robert Bosch Industrietreuhand Kommanditgesellschaft	51,863	722	DE
19	Nec Corporation	71,228	717	JP
20	Mitsubishi Electric Corporation	122,082	644	JP
21	Fujifilm Holdings Corporation	119,637	600	JP
22	Canon Incorporated	180,316	571	JP
23	Jd. Com Incorporated	5,797	566	KY
24	Panasonic Corporation	226,173	542	JP
25	Ericsson (Telefonaktiebolaget NI) Ab	38,329	538	SE
26	Apple Inc.	25,351	457	US
27	Xiaomi Corporation	12,302	452	KY
28	Hp Inc.	31,585	444	US
29	Ford Motor Co	24,256	438	US
30	Halliburton Co	14,166	421	US

Table C2: Main AI adopters of 2019. In comparison to the previous table, the overall number of AI patents increases considerably, with 8 players owning now more than 1,000 patents, in comparison to 3 seen in 2016. Particularly some smartphone manufacturers come into the picture, seen with the first appearance of Apple, LG, and Xiaomi.

Appendix D: Data about AI adopters before and after matching.

The three most used Nace codes are the same before and after the matching, although the code 6201 in particular has a considerably higher share for the matched sample. Size classes are also very similar for both samples (D2), as well as Age (D3) and Year of first adoption (D5). For countries, the three leaders hold some distance from the rest of the countries for both samples, whilst Japanese companies overcome Chinese ones in the matched sample. Overall, the considered AI adopters are overwhelmingly very large companies, usually from the sectors 6201 (Computer programming activities), 5829 (Other software publishing) and 2611 (Manufacture of electronic components), which concentrate 22.5% of the Nace sectors. AI adopters are also mostly companies created before 2001 (63%), situated in the United States, China or Japan (51%), and 41% of them adopted AI in 2017 or 2018.

Data pre-matching (n = 1798)		Data post-matching (n = 871)	
Nace code	Representativeness (%)	Nace code	Representativeness (%)
6201	6.9%	6201	11.4%
5829	5.8%	5829	6.3%
2611	5.3%	2611	4.8%
6420	5.3%	2120	3.8%
6209	3.3%	6209	3.3%
2120	2.6%	2899	3.1%
2899	2.6%	2651	2.8%
2630	2.4%	6420	2.5%
2651	2.4%	7112	2.1%
2620	2.1%	2630	2.0%

Table D1: Nace codes' representativeness before and after the matching procedure.

Size class	Representativeness Pre-Matching (%)	Representativeness Post-Matching (%)
Very large company	72%	73%
Large company	13%	11%
Medium sized company	10%	8%
Small company	6%	8%

Table D2: Size class' representativeness before and after the matching procedure.

Age category	Representativeness Pre-Matching (%)	Representativeness Post-Matching (%)
(0,1980]	23%	25%
(1980,2000]	32%	38%
(2000,2005]	15%	12%
(2005,2010]	13%	12%
(2010,2021]	16%	13%

Table D3: Age categories' representativeness before and after the matching procedure.

Data pre-matching (n = 1798)		Data post-matching (n = 871)	
Country	Representativeness (%)	Country	Representativeness (%)
USA	24%	USA	28%
China	15%	Japan	13%
Japan	11%	China	11%
Germany	5%	Germany	6%
United Kingdom	4%	Taiwan	4%

Table D4: Main countries' representativeness before and after the matching procedure.

Year first AI adoption	Representativeness Pre-Matching (%)	Representativeness Post-Matching (%)
2011	7%	7%
2012	9%	9%
2013	7%	6%
2014	9%	8%
2015	9%	8%
2016	12%	14%
2017	16%	16%
2018	24%	25%
2019	7%	8%

Table D5: Distribution groups (based on the year of AI adoption) before and after the matching procedure.

Appendix E: Average distance of every Nace sector to the AI technological cluster.

Below it is presented the estimated distance for all Nace codes available, from the technologically closest ones (subgroup 1) to the most distant ones (subgroup 3). The calculation of these distances is done in the following way: i) patents are separated into their subclasses (4-digits IPC codes) and aggregated in the Nace sectors of the companies that own them; ii) the relatedness density of all AI patents considered is calculated; iii) the relatedness density between the considered Nace sectors to the AI density cluster is calculating, generating a matrix in which each Nace sector presents values between 0 and 100 for every subclasses of technologies existent; iv) the Mean Absolute Error (MAE) for every sector from their density matrix to the AI cluster is calculated and used as final measure for knowledge distance. The implementation of step 2 and 3 is done using the EconGeo package, through the formulae “relatedness.density.int” and “relatedness.density.ext”, respectively.

Sector	Avg. Dist.	Subgroup number						
6209	11.38	1	8219	14.19	1	7110	15.06	1
4743	11.79	1	5221	14.19	1	6110	15.06	1
4761	11.79	1	3513	14.24	1	6900	15.06	1
6190	11.84	1	6203	14.24	1	6611	15.06	1
5829	11.93	1	6491	14.25	1	5814	15.07	1
6020	12.34	1	8559	14.25	1	8424	15.11	1
6201	12.36	1	4312	14.32	1	8421	15.15	1
6311	12.59	1	9523	14.33	1	8622	15.15	1
5320	12.60	1	4774	14.36	1	3240	15.16	1
5121	12.63	1	9002	14.37	1	1391	15.17	1
6202	12.64	1	8413	14.40	1	6622	15.20	1
5819	12.75	1	6130	14.41	1	5913	15.25	1
8020	12.82	1	6312	14.52	1	6629	15.25	1
8010	12.92	1	9603	14.53	1	7220	15.27	1
8623	13.00	1	4742	14.55	1	9512	15.30	1
2630	13.26	1	8790	14.60	1	8230	15.30	1
6411	13.26	1	2620	14.61	1	1412	15.30	1
4791	13.28	1	4939	14.61	1	9601	15.32	1
2740	13.30	1	9312	14.62	1	2823	15.35	1
6399	13.34	1	4520	14.63	1	4762	15.36	1
7711	13.45	1	4222	14.66	1	4531	15.38	1
7722	13.50	1	4212	14.69	1	0162	15.38	1
8510	13.60	1	9602	14.70	1	4631	15.40	1
8621	13.61	1	9329	14.70	1	6530	15.40	1
9411	13.63	1	5223	14.70	1	4540	15.41	1
4648	13.65	1	4730	14.73	1	1813	15.42	1
8610	13.65	1	9200	14.73	1	3320	15.42	1
4332	13.72	1	4637	14.73	1	4510	15.43	1
4651	13.75	1	9313	14.77	1	7311	15.45	1
7420	13.77	1	7820	14.78	1	2822	15.46	1
8220	13.82	1	7733	14.79	1	4665	15.46	1
8291	13.84	1	4931	14.83	1	4615	15.47	1
7312	13.87	1	5914	14.84	1	1390	15.48	1
7911	13.87	1	9003	14.86	1	4310	15.49	1
2600	13.89	1	5813	14.86	1	7731	15.49	1
6391	13.89	1	8560	14.86	1	0142	15.51	1
6120	13.89	1	2100	14.89	1	4772	15.51	1
4532	13.90	1	6010	14.89	1	8292	15.52	1
6200	13.93	1	5912	14.90	1	3821	15.52	1
8422	13.94	1	2811	14.97	1	7200	15.53	1
5010	14.00	1	7021	14.98	1	0123	15.55	1
2640	14.08	1	9319	15.00	1	6610	15.55	1
			8430	15.04	1	5629	15.56	1
			7912	15.04	1	9604	15.57	1

2700	15.58	1
6800	15.59	1
4777	15.60	1
8899	15.61	1
4775	15.61	1
2540	15.61	1
9311	15.61	1
4932	15.62	1
2731	15.63	1
7020	15.64	1
3220	15.65	1
4619	15.67	1
4770	15.69	1
4940	15.69	1
2500	15.69	1
7400	15.71	1
3040	15.71	1
4765	15.72	1
4331	15.73	1
0990	15.73	1
8730	15.75	1
0620	15.76	1
2651	15.76	1
5920	15.76	1
4799	15.77	1
7800	15.78	1
3812	15.78	1
1700	15.79	2
5630	15.80	2
582	15.81	2
9104	15.81	2
4910	15.82	2
4750	15.83	2
3530	15.87	2
2572	15.87	2
4613	15.88	2
3314	15.88	2
7000	15.88	2
8425	15.88	2
2312	15.89	2
0150	15.89	2
9420	15.89	2
2349	15.89	2
4213	15.90	2
0240	15.90	2
4776	15.90	2
0129	15.90	2
2360	15.91	2
2440	15.91	2
4636	15.91	2
2840	15.91	2
4622	15.91	2
5000	15.91	2
4100	15.91	2
4511	15.92	2
7712	15.92	2
4313	15.93	2
0160	15.93	2
1410	15.94	2

4700	15.94	2
0125	15.94	2
4664	15.95	2
2310	15.95	2
3020	15.95	2
5030	15.95	2
4612	15.96	2
16	15.96	2
3523	15.96	2
0721	15.96	2
3311	15.97	2
2443	15.97	2
4290	15.97	2
5040	15.97	2
9810	15.97	2
3315	15.97	2
0124	15.97	2
0146	15.97	2
3800	15.97	2
691	15.97	2
7310	15.97	2
7735	15.97	2
7830	15.97	2
8290	15.97	2
9524	15.97	2
4610	15.97	2
8122	15.97	2
7729	15.98	2
0164	15.99	2
4110	15.99	2
803	15.99	2
4670	15.99	2
2344	15.99	2
3250	15.99	2
9600	16.00	2
4291	16.01	2
2510	16.01	2
7430	16.01	2
4643	16.02	2
2050	16.03	2
4753	16.03	2
2365	16.04	2
0147	16.04	2
4663	16.05	2
0220	16.05	2
1720	16.06	2
8200	16.06	2
2710	16.06	2
9511	16.06	2
2750	16.07	2
3100	16.10	2
3102	16.10	2
3820	16.11	2
4652	16.11	2
4600	16.12	2
4725	16.12	2
0149	16.13	2
2010	16.14	2
4200	16.14	2

2450	16.15	2
2363	16.15	2
4530	16.16	2
2120	16.16	2
6490	16.16	2
1624	16.18	2
4311	16.18	2
7734	16.19	2
2931	16.20	2
3213	16.20	2
2433	16.21	2
3510	16.21	2
5590	16.21	2
0130	16.22	2
2611	16.23	2
0112	16.23	2
0163	16.23	2
1030	16.23	2
1106	16.23	2
4616	16.23	2
3313	16.23	2
8110	16.23	2
1300	16.24	2
8720	16.24	2
1040	16.24	2
3010	16.25	2
3514	16.26	2
8129	16.27	2
0140	16.27	2
4724	16.27	2
1104	16.27	2
4781	16.27	2
0322	16.28	2
2053	16.30	2
2364	16.32	2
1000	16.33	2
2720	16.33	2
2451	16.34	2
2660	16.35	2
3103	16.35	2
2530	16.38	2
3832	16.38	2
4391	16.39	2
9491	16.41	2
2670	16.41	2
6512	16.41	2
0520	16.42	2
2352	16.42	2
4644	16.42	2
2040	16.43	2
4630	16.43	2
3092	16.44	2
6600	16.44	2
2790	16.44	2
1411	16.45	2
4778	16.45	2
9321	16.46	2
4660	16.47	2
2591	16.47	2

2512	16.54	2
2815	16.54	2
3822	16.54	2
2810	16.54	2
0114	16.56	2
1092	16.57	2
4662	16.58	2
7810	16.58	2
0119	16.60	2
1071	16.60	2
4611	16.60	2
2733	16.61	2
0128	16.63	2
5020	16.64	2
2370	16.66	2
1052	16.66	2
2752	16.67	2
1032	16.69	2
2825	16.70	2
1072	16.71	2
2220	16.72	2
6400	16.72	2
7990	16.74	2
3291	16.75	2
1520	16.75	2
2042	16.75	2
7490	16.76	2
3101	16.76	2
4647	16.76	2
2680	16.77	2
2812	16.79	2
2890	16.80	2
4764	16.81	2
2594	16.84	2
8121	16.88	2
8211	16.90	2
2222	16.90	2
7120	16.90	2
1086	16.91	2
7320	16.93	2
1623	16.98	2
4754	16.98	2
2652	16.99	2
7500	17.00	2
2590	17.04	2
9609	17.04	2
0122	17.05	2
3312	17.06	2
1431	17.08	2
0311	17.10	2
0141	17.10	2
3319	17.10	2
2824	17.11	2
2751	17.13	2
3012	17.14	2
2841	17.15	2
5222	17.15	2
8690	17.17	2
3900	17.17	2

1394	17.18	2
4299	17.18	2
1051	17.19	2
2431	17.20	2
2894	17.21	2
7111	17.21	2
1811	17.23	2
2000	17.25	2
9499	17.26	2
1814	17.26	2
4677	17.28	2
4639	17.28	2
3521	17.29	2
1820	17.30	2
4339	17.31	2
2320	17.35	2
3512	17.36	2
1039	17.37	2
1089	17.38	2
1073	17.40	2
3299	17.40	2
2332	17.41	2
3230	17.41	2
2342	17.42	2
1812	17.44	2
4741	17.49	2
6419	17.50	2
1723	17.50	2
4632	17.51	2
1396	17.51	2
1622	17.51	2
4771	17.53	2
5530	17.54	2
1621	17.54	2
1511	17.56	2
1419	17.58	2
1085	17.58	2
4661	17.59	2
3091	17.60	2
1413	17.61	2
1320	17.61	2
1393	17.63	2
4646	17.63	2
4729	17.63	2
4333	17.64	2
2893	17.64	2
4666	17.66	2
4321	17.66	2
1721	17.69	2
4623	17.69	2
1083	17.71	2
1200	17.74	2
6511	17.77	2
1084	17.77	2
1061	17.79	2
2592	17.79	2
2017	17.79	2
1610	17.84	2
6832	17.84	2

2830	17.85	2
2369	17.86	2
2820	17.87	2
1062	17.88	2
1414	17.88	2
2571	17.88	2
3511	17.89	2
8299	17.90	2
2052	17.93	2
2211	17.94	2
0321	17.95	2
2612	17.96	2
2910	17.96	2
5811	17.97	2
2331	17.97	2
2849	17.98	2
2361	18.03	2
8130	18.04	2
5224	18.05	2
4399	18.06	2
1042	18.06	2
1102	18.07	2
2712	18.07	2
5210	18.07	2
1439	18.08	2
1729	18.12	2
7732	18.13	2
0610	18.17	2
6619	18.19	2
4645	18.24	2
0161	18.25	2
0113	18.26	2
2932	18.26	2
1711	18.29	2
2896	18.29	2
2410	18.30	2
2821	18.30	2
4211	18.32	2
2030	18.34	2
7211	18.34	2
2051	18.35	2
2441	18.35	2
2711	18.36	2
0812	18.38	2
4773	18.39	2
7739	18.39	2
1392	18.41	3
2444	18.41	3
2110	18.42	3
4719	18.42	3
2452	18.43	3
1101	18.45	3
4334	18.45	3
7022	18.45	3
3212	18.47	3
4638	18.51	3
8542	18.52	3
0710	18.53	3
6612	18.53	3

4617	18.54	3
2362	18.55	3
1012	18.55	3
2895	18.56	3
4621	18.56	3
4676	18.56	3
2434	18.57	3
2573	18.58	3
3109	18.58	3
3522	18.58	3
0116	18.60	3
4322	18.62	3
1512	18.64	3
1724	18.65	3
3030	18.65	3
3099	18.67	3
2012	18.67	3
4120	18.70	3
4633	18.70	3
0910	18.73	3
1722	18.76	3
1395	18.79	3
1082	18.79	3
1105	18.81	3
4642	18.81	3
6920	18.85	3
9412	18.87	3
2314	18.87	3
2453	18.88	3
2059	18.89	3
5911	18.91	3
4759	18.91	3
2445	18.96	3
4641	18.97	3
3700	18.98	3
2015	18.99	3
1020	18.99	3
2041	19.00	3
0510	19.02	3
2319	19.03	3
1013	19.03	3
2562	19.07	3
0111	19.09	3
2014	19.09	3
2529	19.10	3
2891	19.10	3
4752	19.12	3
2892	19.14	3
2920	19.14	3

4690	19.14	3
2521	19.16	3
8541	19.17	3
1081	19.19	3
2341	19.21	3
1330	19.25	3
2813	19.26	3
2013	19.30	3
2561	19.31	3
0893	19.33	3
1920	19.33	3
5110	19.36	3
7210	19.36	3
2229	19.37	3
2219	19.39	3
2016	19.40	3
2432	19.41	3
1712	19.42	3
1310	19.42	3
2313	19.42	3
4711	19.44	3
1031	19.46	3
8412	19.48	3
4672	19.50	3
6831	19.53	3
4519	19.53	3
0811	19.54	3
4614	19.56	3
1041	19.57	3
1107	19.57	3
2343	19.57	3
2311	19.59	3
0210	19.61	3
4674	19.61	3
5229	19.61	3
7740	19.64	3
2454	19.65	3
8411	19.65	3
1091	19.68	3
6910	19.70	3
5610	19.70	3
2442	19.71	3
2060	19.73	3
8552	19.75	3
2223	19.78	3
5510	19.79	3
2221	19.80	3
6500	19.81	3
0729	19.82	3

2420	19.84	3
1011	19.85	3
2732	19.86	3
4329	19.89	3
4675	19.93	3
1420	19.97	3
2550	19.98	3
3811	19.98	3
3600	20.04	3
4634	20.04	3
1910	20.07	3
4950	20.07	3
2351	20.10	3
2814	20.13	3
2391	20.17	3
2011	20.17	3
4669	20.20	3
6492	20.21	3
6810	20.25	3
4673	20.28	3
2800	20.40	3
1399	20.41	3
2020	20.42	3
2593	20.47	3
2399	20.49	3
2511	20.50	3
6630	20.51	3
0891	20.55	3
3011	20.58	3
6430	20.61	3
4221	20.84	3
2899	20.88	3
6820	21.10	3
1629	21.11	3
6420	21.17	3
2829	21.24	3
4941	21.26	3
4618	21.47	3
7410	21.47	3
2599	21.81	3
0899	22.02	3
4671	22.04	3
7112	22.23	3
7010	22.31	3
6499	22.31	3
4649	22.55	3
7219	22.98	3

Appendix F: Dynamic effects for the self-selection analysis for the first (Figure F1) and second (Figure F2) classification methods used.

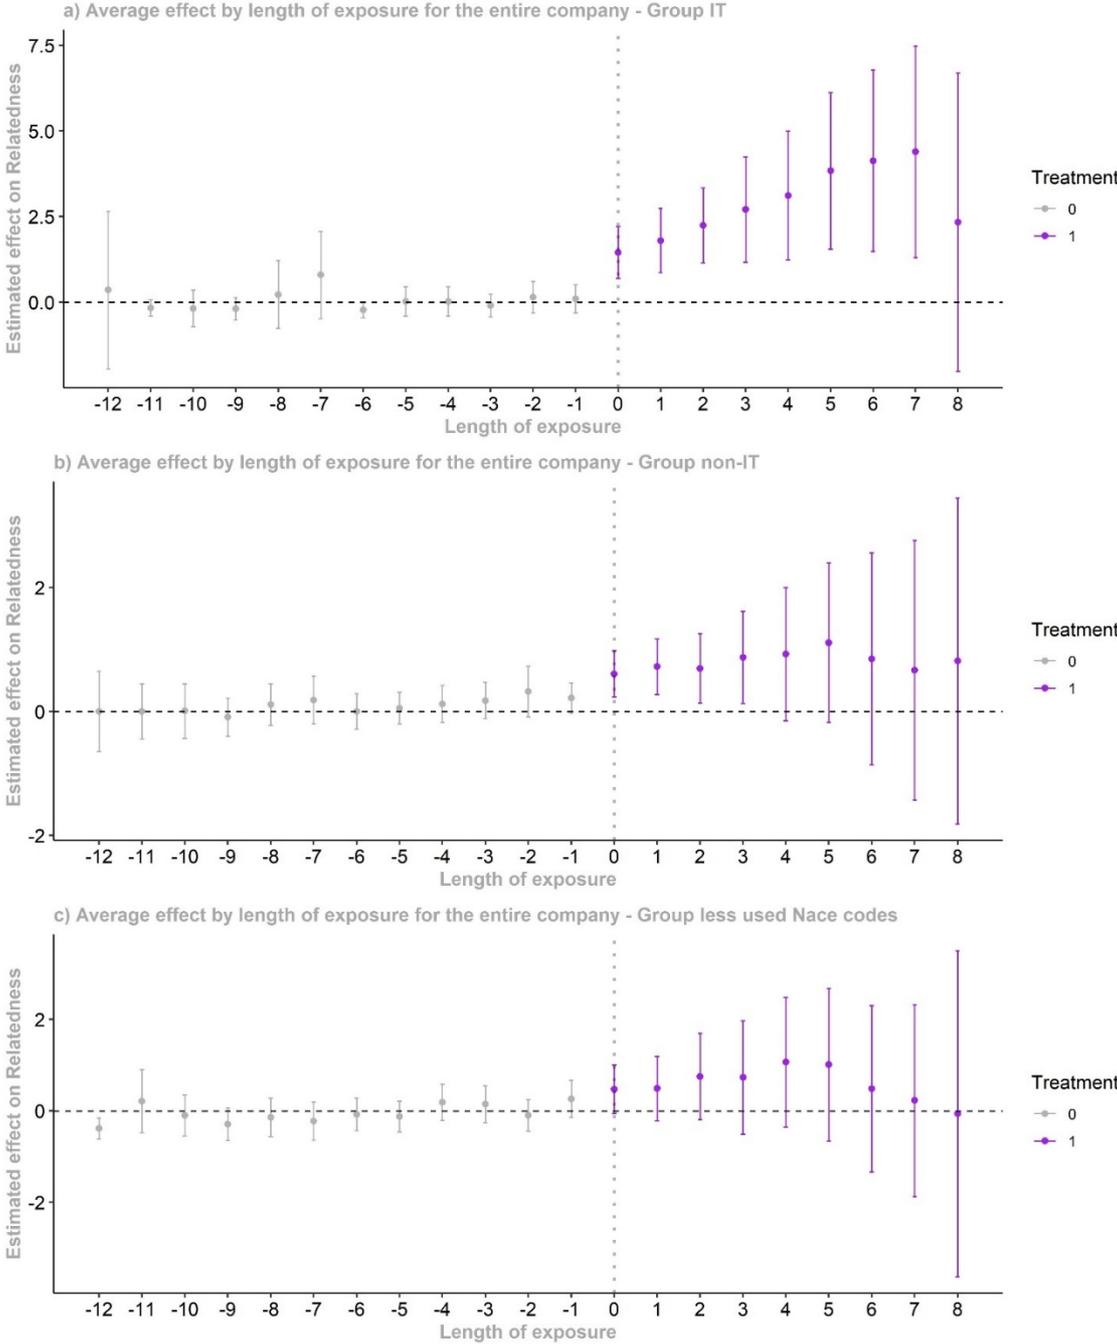


Figure F1: Estimated effects on firms' relatedness considering Nace-based classification.

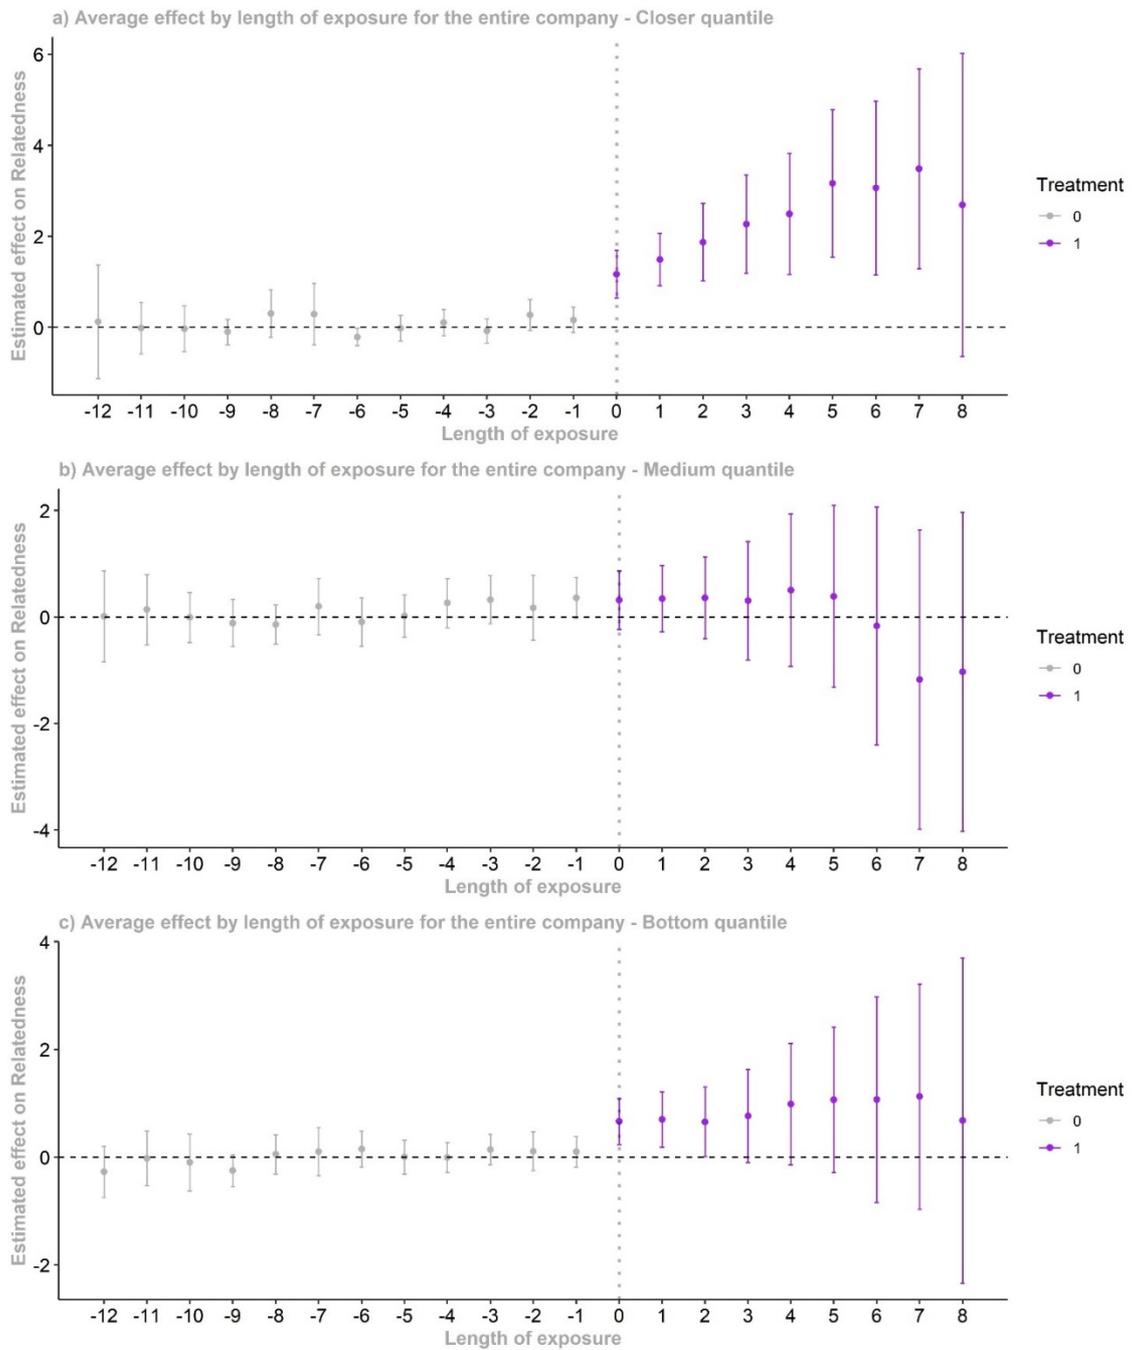


Figure F2: Estimated effects on firms' relatedness considering knowledge distance-based classification.

Appendix G: Distribution of distinct considered samples regarding the size of the companies.

Size	Distribution all matched companies	Distribution companies with Turnover data	Distribution of companies with R&D expenses data
Very large company	73%	87.6%	99.5%
Large company	11%	7.6%	0.5%
Medium sized company	8%	3.8%	0.0%
Small company	8%	1.0%	0.0%
Number of MNEs in the sample	6,930	4,259	2,649
Units lost in comparison to the total		38.5%	61.9%

PERSONAL CONTRIBUTIONS TO THE PAPERS OF THE CUMULATIVE DISSERTATION

Patenting Patterns in Artificial Intelligence: Identifying National and International Breeding Grounds (Chapter 2).

This paper is a joint work with Prof. Dr. Jutta Günther, Prof. Dr. Björn Jindra, and Prof. Dr. Martin G. Möhrle. It is largely based on my own work, supervised by the three co-authors. Besides the supervision, Jutta Günther provided specific resources for the collection of data (i.e., through PATSTAT), Björn Jindra provided additional assistance with the identification of relevant literature regarding economic indicators and resources that allowed the use of language editing services, and Martin G. Möhrle contributed to the conceptualisation and writings on the draft version of the paper. This information is supported by the “CRediT authorship contribution statement” published in the paper (Pg. 11).

Break-through or Break-in? How AI Becomes a Part of National Technological Trajectories (Chapter 3).

This paper is a joint work with Prof. Dr. Björn Jindra and Dr. Daniel S. Hain. It is largely based on my own work supervised by the two co-authors. This includes my participation in the complete coding, data collection, empirical analyses, conceptualisation, and writing of the submitted version of the paper. Besides supervision, both co-authors also helped in the conceptualisation. Particularly, Björn Jindra also contributed with writings on the draft version of the paper and provided resources that allowed the use of language editing services.

The Development of AI in Multinational Enterprises – Effects upon Technological Trajectories and Innovation Performance (Chapter 4).

This paper was written by me as the sole author. However, special thanks and recognition should be addressed to Dr. Felix Lüders and to Prof. Dr. Björn Jindra regarding the development of the dataset used in this Chapter. Felix and I developed this dataset together (meaning joint efforts in downloading, writing code, and analysing each other’s results) in a joint collaboration that lasted several months. We further developed the dataset for our specific needs (which were linked to two separate sole author papers, one by each of us) by ourselves. Björn closely supervised the complete endeavour and development of the dataset.

ACKNOWLEDGEMENTS

I watched lots of movies in which AI is responsible for a catastrophic future, with machines subduing humans. The more realistic near future, without any AI killer robot, seems worse than most of these movies (if not all). Every day our planet gets closer to the tipping point of human-induced irreversible climate change, and yet, we manage to create additional catastrophes in the meantime. Instead of uniting efforts to fight global challenges, we divide ourselves. The current war in Ukraine is painful and terrible evidence of how inept we are for cooperating on a global scale. It also shows how susceptible the lives of a whole population of innocent people are to the wishes of autocratic regimes. A couple of men exercise political power in one country, and a despairing large number of families have their lives interrupted and taken away in another country. People are attacked and killed in their homes, hospitals, and even in humanitarian corridors. Amid this, the risk of nuclear power and mutually assured destruction between countries arises once again, menacing also those who think are safe by refusing to act.

But at the same time, there is an increasing number of people doing something positive to help. The large majority are willing to risk their comfort in the endeavour, with the caveat that some may lack the courage. Between the courageous people, I met some exceptional ones that made me believe that there is always something that can be done to make a positive difference. These people give me hope that things can get better. One of the exceptional people I'm grateful to have met in the last three and a half years was my supervisor, Prof. Jutta Günther. I wasn't looking for a role model when I started my PhD, but I found one in her. Jutta showed me that with enough courage, it is possible to be unbelievably loyal to one's own beliefs. I had no idea that seeing closely someone doing what is right every day – even (and especially) when these actions are against the actor's own comfort – could be so transformative and inspiring. Jutta changed my understanding of who I want to be as a person. The other wonderful surprise of this journey was Prof. Björn Jindra, my favourite co-author. I would think that it would take almost an entire lifetime to accumulate the wisdom and patience that Björn has, but he is not much older than me. His persistent support and diligent care not only with our scientific work but also with my personal development took many hours of his life to make mine easier and more enjoyable. Björn went out of his way countless times to help me, and I'm very thankful to him for that.

If I mention going out of one's own way, I have to mention who did that literally too. I would like to thank Monize Heimann, who has fully embarked on this journey with me. Monize has way more faith in me than I do, and I haven't figured out why yet. But I feel extremely lucky for that, as well as grateful. If she reads this one day, I hope she is reminded that I'll never stop admiring her.

Back to the people I met during the PhD, I'd like to thank my working colleagues. First, the ones that shared an office with me: Jarina Kühn, Katharina Friz, and Kehinde Medase. Thanks for the great and stimulating conversations, companionship, and friendship! I'd also like to thank the veterans of the group that were always ready to help: Ann Hipp, Mariia Shkolnykova, Marcel Lange, and Tobias Wendler. Marcel connects the whole group: there is no technical problem he is not capable of solving, and his gentleness and competence while doing it has repeatedly impressed me. He was happy to help me solve personal problems as well, like translating and making sense of the last German letter that seems very important but I don't understand any of the content. The other veterans helped me directly or indirectly by making invariably great suggestions and comments. I will always remember Ann for her combination of kindness and extreme honesty, Tobi for his friendly attitude and charisma, and Masha for her commitment and unmeasurable strength. I'm also thankful for meeting colleagues from other working groups. Prof. Michael Rochlitz and his working group, in particular, were always close

valuable friends. I'm especially grateful for having known and shared good times with Michael, as well as David Karpa and Olga Masyutina.

Last but not least, I'd like to thank my other friends and people I love. I'd like to thank my parents, who always provided me with unconditional love. They never discouraged me from following what I want, even when that meant moving to a faraway country and not seeing them for a painfully long time. I'm also grateful to my uncles and aunts, André, Carolina, Ademir, and Tânia. You were my home when I needed, and I always carry you in my heart. I'd like to thank friends that shared the PhD struggle with me, Philipp Tobias Hohn and Thiago Carrano de Albuquerque Bernardes. It was great having your company while complaining about the many things that went wrong! And also, of course, just to chat: sometimes while going to Mensa with Tobias, sometimes via 12 minutes-long WhatsApp voice messages with Thiago. I'd also like to thank Christoph Klose for his friendship, support, and kind advice. It was always a pleasure to have you around and to see that my poor German skills never demotivated you from trying to teach me. I'm very grateful to my former masters too. Prof. Mauricio Uriona Maldonado and Prof. Enzo Morosini Frazzon, in particular, introduced me to academic research during my Master's. Mauricio was always generous with his time and gave me all the support I needed. Enzo always looked to my future and incentivised me to look for life-changing experiences. He forwarded me an e-mail out of the blue long ago, which happened to be the one announcing the PhD position that I took. That meant the world to me and gave me the opportunity to have the wonderful experience that I had in the last years in Bremen.

I'd also like to thank my oldest best friends, Joabel Luis Kasper and Felipe Gaspareto Seminotti. I can tell Joabel anything at any time. It is comforting to know we can interrupt contact for ages and still restart it from the same point later, using the same jokes. The strength to carry a gigantic head above the shoulders is proof of the impressive willpower this man has, which is quite uplifting. To Felipe, I'm thankful for the old memories, which I always remember with fondness. They become painful only when I realize I'm getting old, while he couldn't.

If the world ends soon, I wish everyone had the chance to meet at least half as many wonderful people as I did.