When one teaches, two learn:
The bidirectional learning process supervisor-PhD student

Pierre Boutros
Université Côte d'Azur, CNRS, GREDEG, France
pierre.boutros@etu.univ-cotedazur.fr

Michele Pezzoni
Université Côte d'Azur, CNRS, GREDEG, France
Observatoire des Sciences et Techniques, HCERES, Paris, France
michele.pezzoni@univ-cotedazur.fr

Sotaro Shibayama
Institute for Future Initiatives,
The University of Tokyo, Tokyo, Japan
shibayama@ifi.u-tokyo.ac.jp

Fabiana Visentin
UNU-MERIT, Maastricht University, the Netherlands
f.visentin@maastrichtuniversity.nl

Abstract: In contexts involving teachers and students, knowledge transfer is commonly assumed from the former to the latter. However, what if teachers learn from students? This paper investigates the bidirectional knowledge transfer between PhD students and their supervisors. We consider 40,852 PhD students who graduated in the STEM fields in France between 2010 and 2018. Focusing on Artificial Intelligence (AI) knowledge transfer, we find evidence that a student supervised by a supervisor with AI knowledge is 10 percentage points more likely to write a thesis in AI than a student with a supervisor with no AI knowledge, denoting an AI knowledge transfer from supervisors to students. We also find that a supervisor with no AI knowledge is 19 percentage points more likely to publish an article with AI content if exposed to a student with AI knowledge than a supervisor not exposed to a student with AI knowledge, denoting an AI knowledge transfer from students to supervisors. Those results confirm the bidirectionality of the learning process.

Keywords: PhD training, Artificial Intelligence, Knowledge transfer, STEM in France

JEL codes: J24, O30

Acknowledgments: We thank participants of the seminar series at Scuola Superiore Sant'Anna, Pisa, Italy; the KID summer school, Nice, France; and the summer school organized by REGIS, Madrid, Spain. We also thank our colleagues at the University of Bordeaux, Bordeaux, France, and Maastricht University, Maastricht, Netherlands, for their valuable comments and suggestions. We are also grateful to the French government for its financial support of this work through the grant ANR-24-CE26-1154.

Competing interest: The authors disclose no relevant financial interests related to the research in this paper.

CRediT authorship contribution statement: *Pierre Boutros*: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. *Michele Pezzoni*: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. *Sotaro Shibayama*: conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. *Fabiana Visentin*: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization.

1. Introduction

Academic training is a period during which students and supervisors exchange scientific knowledge (Kyvik and Smeby, 1994; Colbeck, 1998; Bozeman and Corley, 2004; Laudel and Glaser, 2008). Once trained, students disseminate the acquired knowledge within the academic community and to society (Dasgupta and David, 1994; Sumell et al., 2009; Stephan, 2001, 2012). Critical to academic training is the relationship between the student and the supervisor. In this relationship, the learning process has been traditionally viewed as the supervisor unidirectionally transferring knowledge to the student (Kam, 1997; Delamont and Atkinson, 2001; Campbell, 2003). However, unidirectionality might not always be the case. Students may acquire specific knowledge before their supervisors (Zuckerman, 1988; Stephan and Levin, 1992) and transfer that knowledge to them. This is especially true in fast-changing technological landscapes (Fleming and Sorenson, 2003), in which supervisors often learn from their students the use of new technologies. One such technology is Artificial Intelligence (AI). The use of AI in science has grown fast in the last decade, transforming scientific research across research fields (Flanagan et al., 2023). Given its recent upsurge, students might have better and more up-to-date AI knowledge than their supervisors, becoming a source of AI knowledge.

Our paper aims to empirically demonstrate the bidirectional transfer of AI knowledge between supervisors and students. Specifically, we ask: Are supervisors with AI knowledge transferring this knowledge to their students? Are students with AI knowledge transferring this knowledge to their supervisors?

From the extant literature on higher education, we know that the role of the supervisor in student-supervisor relationships is to provide the student with knowledge, expertise, and guidance (Cox, 1988; Kam, 1997; Delamont and Atkinson, 2001; Hilmer and Hilmer, 2011; Shibayama, 2019; Patsali et al., 2024). We also know that the supervisor's training styles have different effectiveness in determining the successful accomplishment of the student's training (Cox, 1988; Kam, 1997; Delamont and Atkinson, 2001; Gurr, 2001). When looking at the role played by the students in student-supervisor relationships, the extant literature on academic organizations shows that students are crucial actors in the knowledge production process for their supervisors and scientific teams (Conti and Liu, 2015; Igami et al., 2015; Shibayama, 2019; Wang and Shibayama, 2022). While these studies advanced our knowledge of the role played by supervisors and students in science, the theoretical discussion about academic training has always assumed a unidirectional knowledge transfer process from supervisors to students. However, the unidirectionality of knowledge transfer can be questioned by drawing

from results in the organizational learning literature. Organizational learning literature suggests that organization members are crucial in reinforcing and updating the organization's knowledge basis, showing that hiring individuals might be the source of new fresh knowledge for organizations (Song et al., 2003; Amara et al., 2008; Tzabbar, 2009; Slavova et al., 2016). Most contributions to the organizational learning literature considered relatively large corporate organizations as a unit of analysis (Song et al., 2003; Tzabbar, 2009), neglecting the study of individuals during the learning process. Also, they considered cases in which experienced workers changed employers (Slavova et al., 2016; Franzoni et al., 2018), neglecting the case of the recruitment of inexperienced trainees. Applying organizational learning theory to the academic context, we argue that seemingly inexperienced students can become a source of exploratory learning for the organization's senior members, i.e., their supervisors, generating a bidirectional knowledge transfer in the context of academic training. Finally, regarding the specific context of our analysis, we know that AI knowledge is becoming a highly valuable asset as shown by an increasing demand and reward for AI professionals in the labor market (Acemoglu et al., 2022; Manuel, 2024) and by its increased use in scientific and technological research (Cockburn et al., 2018; Bianchini et al., 2022).

To empirically document the AI knowledge transfer process, we draw on a unique dataset covering the population of 40,852 French PhD students in STEM who graduated between 2010 and 2018. We identify AI documents by selecting those documents that include at least one AI-specific keyword in their abstract or title. Then, we identify AI students and AI supervisors as those using AI keywords in their theses and published articles, respectively. To illustrate the bidirectional AI knowledge transfer, we conduct two analyses at the student-supervisor pair level. First, we investigate the *supervisor-to-student* knowledge transfer by looking at the probability of writing a PhD thesis with AI content when the supervisor has published articles with AI content before the student's enrollment. Second, for those supervisors who had no AI knowledge before the student's enrollment, we investigate the *student-to-supervisor* knowledge transfer by looking at the supervisor's probability of publishing an article with AI content after mentoring an AI student.

Our empirical analysis shows that students in student-supervisor pairs characterized by supervisors with AI knowledge before the student's enrollment are 10 percentage points more likely to write an AI thesis than students in pairs in which supervisors have no AI knowledge before the student's enrollment. Restricting to student-supervisor pairs in which the supervisor has no AI knowledge before the student enrollment, we find that pairs characterized by students

writing a thesis in AI are 19 percentage points more likely to have a supervisor publishing an AI article three years after the student's graduation. These two findings show evidence of bidirectional learning between supervisors and students.

This paper provides four main contributions to the literature. First, it contributes to the debate on academic training (e.g., Shibayama, 2019; Patsali et al., 2024) by illustrating the bidirectionality of knowledge transfer between supervisors and students. In other words, it proposes and empirically tests the idea that supervisors and students can benefit and learn from each other. Second, this study speaks to the literature on organizational learning in the context of academic organizations (Conti and Liu, 2015; Igami et al., 2015; Shibayama, 2019; Wang and Shibayama, 2022) by showing that students can be a channel for exploring and updating the knowledge basis of academic organizations. Third, this study contributes to the literature on higher education by focusing on the semantic content of the training, i.e., AI knowledge, complementing the previous studies focusing mainly on the styles and skills of training. Fourth, our study represents the first attempt to study the supply of AI talents in the French AI labor market.

The remainder of this paper is structured as follows. Section 2 outlines the previous studies on PhD training and knowledge transfer. Sections 3 and 4 illustrate our empirical context and data, respectively. Section 5 illustrates our methodology, which implements a propensity score matching approach to analyze bidirectional learning. Section 6 presents the main results and further empirical analyses. Finally, Section 7 concludes and discusses the implications of our findings.

2. The bidirectional learning process in academic training

2.1 knowledge transfer from supervisors to students

The literature on academic organizations has explored the role of organizational factors in scientific research, demonstrating that academic organizations, such as laboratories, can achieve their research goals by arranging a portfolio of research projects, combining knowledge from different sources, and allocating tasks and responsibilities among their members (Latour and Woolgar, 1979; Knorr-Cetina, 1999; Owen-Smith, 2001; Carayol and Matt, 2006; Heinze et al., 2009; Conti and Liu, 2015; Shibayama et al., 2015). However, academic organizations are meant not only to produce research output but also to serve as academic training, transforming early career researchers into independent researchers (Cox, 1988; Hacket, 1990;

Gurr, 2001; Campbell, 2003; Bozeman and Corley, 2004; Laudel and Glaser, 2008; Shibayama, 2019).

Academic training is typically provided through the so-called apprenticeship model, which occurs when early career researchers, i.e., students, are in postgraduate degree training programs (Stephan, 2012: Ch.7; Kahn and Ginther, 2017). In this model, while students offer labor to supervisors, supervisors provide coaching, transfer knowledge to students, and guide them to accumulate human capital. This training model has a collaborative nature and mainly relies on a learning-by-doing process (Kyvik and Smeby, 1994; Colbeck, 1998; Bozeman and Corley, 2004; Laudel and Glaser, 2008). The learning-by-doing process consists of a student engaging in a research project and assigned tasks under the mentorship of a supervisor who transfers knowledge to the student. The supervisor transfers two main types of knowledge to the student: unspecific and specific knowledge.

Regarding the transfer of *unspecific knowledge*, the students gain psycho-social and career-related knowledge from supervisors (Marsh et al., 2002; Paglis et al., 2006). Studies on academic organizations empirically tested the transfer of unspecific knowledge from supervisors to students. For example, Wang and Shibayama (2022) study if unspecific creativity skills are transferred during doctoral training. They find a correlation between the novelty of supervisors' research output and their PhD students' work after they earned degrees and attained tenured positions. Yoshioka-Kobayashi and Shibayama (2021) also find that PhD candidates and postdocs exposed to certain research practices during training tend to replicate the same practices throughout their careers.

Regarding the transfer of *specific knowledge*, due to the nature of apprenticeship, students are often assigned to research topics related to their supervisors' research interests (Kam, 1997; Delamont and Atkinson, 2001; Patsali et al., 2024). By working on topics of interest for their supervisors, students benefit from their knowledge, resources, and attention (Delamont et al., 1997). These conditions create a favorable environment for the students to learn specific knowledge, such as AI techniques. Transferring specific knowledge from supervisors to students can have a long-lasting impact, influencing the students' research topic selection in later career stages (Austin, 2002; Laudel and Glaser, 2008). Indeed, researchers tend to choose coherent research topics after graduation, often continuing to collaborate with their supervisors and citing their work (Shibayama, 2019).

In this paper, we aim to add to the scarce literature on the transfer of specific knowledge from supervisors to students by measuring students' learning of AI knowledge from their supervisors. Moreover, as described in section 2.2, we aim to complement the extant literature by tackling the almost unexplored hypothesis that academic training allows supervisors to learn specific knowledge from their students.

2.2 Knowledge transfer from students to supervisors

Exploration is at the core of all creative activities, including academic research (McFadyen and Cannella 2004; Foster et al. 2015). As with any other type of organization, academic organizations must find ways to update their knowledge to explore new research avenues (Gupta et al., 2006; Amara et al., 2008). The literature on organizational learning has suggested that organizations can use three main approaches to obtain the new knowledge needed for exploration.

The first approach is to rely on *external collaborators*. In academic organizations, researchers often obtain the knowledge needed to accomplish a project by forming teams of collaborators who collectively have the knowledge needed (Ayoubi et al., 2017; Katz and Martin, 1997; Wuchty et al., 2007). The limitation of relying on external collaborators is that the knowledge brought to the organization from external collaborators may remain external and not be internalized. Different from this latter approach, i.e., the use of external collaborators, the two following approaches allow organizations to internalize the needed knowledge.

The second approach organizations use to explore new research avenues is to ask their members to invest in updating their knowledge (Amara et al., 2008). In the academic context, researchers codify the knowledge produced in research articles that become publicly accessible sources of knowledge. Thus, organizations' members may be able to retrieve new knowledge from public sources, test it in-house, and internalize it. Among the organization members, students might be encouraged by senior members, i.e., their supervisors, to acquire the knowledge needed. In this case, the knowledge acquisition process is expected to be monitored by the supervisor and undertaken in practice by the student during the training period, resulting in student-supervisor *co-learning*. Leveraging students to acquire new knowledge, rather than senior researchers, has some advantages. One advantage is that students are a cheap scientific labor force. Moreover, students may have greater flexibility in the choice of research topics or approaches compared to senior researchers, which makes them an ideal vehicle for acquiring new knowledge. Finally, students' lack of experience may be an advantage, removing any possible bias or constraint shaped by prior experiences and making them a potential source of serendipity and innovation

(Zuckerman, 1988; Stephan and Levin, 1992). A few works have investigated the beneficial effect of students' learning on academic organizations. Shibayama et al. (2015), examining the division of labor between PhD students and their supervisors in life science research, argued that supervisors learn from students and build on students' learning. Drawing on survey data about research teams in the U.S. and Japan, Igami et al. (2015) suggested that postdoctoral fellows are fundamental actors in fast-moving and competitive scientific research areas. Based on complete personnel data from a biology university department, Conti and Liu (2015) emphasized the role played by graduate students and postdocs in achieving breakthrough scientific discoveries. Finally, Zanchetta et al. (2017) found that mentors and mentees co-develop and share new perspectives in an undergraduate-level nursing program context.

The third approach organizations use to explore new research avenues is to obtain new knowledge by relying on a *learning-by-hiring* process. Learning-by-hiring consists of employing individuals who already have the target knowledge needed. Updating the knowledge basis of organizations through learning-by-hiring has been broadly studied (Song et al., 2003; Tzabbar, 2009). For instance, it has been shown that R&D-intensive organizations use new hires to transform their technological portfolio and access new scientific knowledge (Zucker and Darby, 1997; Jain, 2016; Buenstorf and Heinisch, 2020). In the academic context, only a few studies refer to learning-by-hiring. Slavova et al. (2016) examined the effects of researchers' inbound mobility on incumbent researchers' performance in U.S. chemical engineering departments. Similarly, Franzoni et al. (2018) surveyed scientist teams in 16 countries and investigated the effect of internationally mobile (incoming) researchers on the teams' performance. These studies also identified contextual factors that make learning-by-hiring more effective. The identified factors include close interaction between the incumbent and new hires (Slavova et al., 2016) and commitment to creativity (Franzoni et al., 2018). These studies examined experienced researchers, and recruiting students has not been considered a potential way of learning-by-hiring. Nonetheless, hiring students can be a compelling opportunity for organizational learning. For instance, students in master's degree programs or pre-doctoral training programs (e.g., internships) may learn advanced AI techniques, which might appeal to academic organizations looking for PhD candidates with AI knowledge.

In this paper, we argue that these two latter approaches, i.e., co-learning and learning-by-hiring, can be used in academic training to acquire the knowledge needed for exploration. In other words, an academic organization may leverage students to explore new research areas by encouraging them to learn specific knowledge and bring it within the organization's perimeter

(*co-learning*) or recruiting students who have already learned the necessary knowledge from their previous experiences (*learning-by-hiring*).

3. Empirical context

To explore the bidirectional learning between supervisors and students, we focus on the AI knowledge exchanged between supervisors and students during the PhD training in French universities in STEM fields.

We focus on AI as the area of knowledge transferred because it is a growing technology transforming scientific research, affecting scientists' productivity (Flanagan et al., 2023). Moreover, AI is a general-purpose technology with applications across all scientific fields, fostering interdisciplinary approaches and driving innovation (Cockburn et al., 2018; Bianchini et al., 2022). Finally, countries worldwide are investing heavily in AI research and education. France, for instance, launched its National Strategy for AI[1] in 2017, allocating more than 3.7 billion euros for the 2018-2025 period. The French AI Strategy includes supporting the production of AI knowledge in universities and training AI talents. However, despite the substantial use of public resources, large-scale quantitative studies are lacking to assess how effective these investments are. This makes France our ideal empirical setting.

We limit our empirical analysis to STEM fields for three main reasons. First, STEM fields assure a reliable use of peer-reviewed publications as an encompassing proxy for researchers' scientific outcomes. Research outcomes in STEM fields are more likely to be published in peer-reviewed journals than in Social Sciences, in which books and book chapters are alternative valuable outcomes for researchers. The coverage of research outcomes other than peer-reviewed publications is unclear in bibliographic datasets, creating potential problems when analyzing Social Sciences. Moreover, the lack of peer review for outcomes other than journal publications does not assure a minimum quality of those outcomes. A second reason for considering only STEM fields is that the delay between the research results and their publication is likely to be shorter than in Social Sciences, allowing for a timely measure of the scientific activity of the researchers using publication data. Finally, in our analyses, we use keywords to classify documents (publications and theses) with AI content. The precision in the use of the language in STEM fields assures the accuracy of our document classification. Indeed, documents in STEM fields are expected to use more precise language than Social Sciences, in which broad terms and synonyms are often used. Moreover, the wide adoption of the English

[1] https://www.economie.gouv.fr/strategie-nationale-intelligence-artificielle#

language in STEM fields copes well with our list of English keywords used to classify documents with AI content.

4. Data

4.1 Data sources

In our empirical analysis, we rely on four datasets. The first dataset is the French national repository of the PhD theses defended. It includes information about the thesis documents, students, and supervisors. Among the information relevant to our analysis, we gathered the thesis title and abstract, defense year, field of study, and the names of both the students and supervisors. The second dataset is *OpenAlex*, from which we collected bibliographic data on supervisors (Priem et al., 2022). We assigned journal articles collected in *OpenAlex* to supervisors by matching their given names and surnames[2]. For each identified journal article, we retrieved its title and abstract, publication year, author list, and the *OpenAlex* classification of the content (i.e., concepts[3]). Third, to identify the genders of students and supervisors, we use the official French gender-name dataset[4], matching given names to those listed in the dataset. Finally, we gather information on national research funding using the complete list of individual grants awarded by the *Agence Nationale de la Recherche* (ANR)[5], the French national funding agency.

After merging the four datasets, we ended up with a study sample of 40,852 theses defended between 2010 and 2018. Our sample includes 17,287 distinct supervisors, who, on average, mentor 2.36 students each. Each of the 40,852 theses identifies a student-supervisor pair for which we assess the bidirectional knowledge transfer.

4.2 AI theses and AI publications

To identify theses with AI content, we create a reliable list of AI keywords by merging two widely recognized lists. The first list of 189 AI keywords is obtained from an OECD report aiming to develop a methodology to identify publications and patents with AI content (Baruffaldi et al., 2020; Calvino et al., 2023). The second list of 39 AI keywords is obtained from a pioneering study aiming to identify publications and patents in AI (Cockburn et al., 2018). After merging the two lists, including plural versions of the keywords reported only in their singular form, and removing the keywords too generic to identify AI documents, we ended

[2] More details about the matching process can be found in Appendix B
[3] https://docs.openalex.org/api-entities/concepts
[4] https://www.data.gouv.fr/fr/datasets/liste-de-prenoms/
[5] https://anr.fr/

up with a list of 272 AI keywords.[6] A Student's thesis is classified as an *AI thesis* if it includes at least one AI keyword in its abstract or title.[7] In our study sample of 40,852 theses, we identified 3,012 theses with AI content, 7.37% of the total. We apply a similar approach to identifying the supervisors' publications with AI content, defining *AI publications* as those with at least one AI-specific keyword in their abstracts or titles. Among the 1,315,638 publications attributed to the 17,487 distinct supervisors in our study sample along their careers (up to 2024), we identified 35,027 publications with AI content, 2.66% of the total.

Once theses and publications that include AI keywords are identified, we evaluate the bidirectional transfer of AI knowledge between students and their supervisors. To do so, we define AI students and AI supervisors as individuals with AI knowledge content codified in their theses and publications, respectively. Specifically, we calculate three variables to identify an *AI student*, an *AI supervisor after* the student's defense, and an *AI supervisor before* the student's enrollment in the PhD program.[8] All three variables are calculated by adopting the student-supervisor pair as the unit of analysis. *AI student* is a dummy variable that takes value one if the student writes a PhD thesis with AI content. *AI supervisor after* is a dummy variable that takes value one if the supervisor has at least one AI-related publication three years after the student's defense year (*t*), i.e., from *t+1* to *t+3*. Finally, *AI supervisor before* is a dummy variable that takes value one if the supervisor has at least one publication in AI at any time before the student's enrollment, i.e., in *t-4* or before.[9]

We study the supervisor-to-student knowledge transfer on a sample of 40,852 student-supervisor pairs. 10,333 of these pairs are characterized by a supervisor with AI knowledge before the supervision period (*AI supervisor before* = 1), while 30,519 have a supervisor lacking AI knowledge (*AI supervisor before* = 0). Table 1 shows that 18.60% of the pairs in which the supervisor has AI knowledge also show a student writing a thesis with AI content (*AI student* =

[6] A detailed description of the procedure used to identify AI keywords is reported in Appendix A. In Appendix E, we rerun our main analyses identifying AI documents using separately the two AI keyword lists by Cockburn et al. (2018) and by OECD (Baruffaldi et al., 2020; Calvino et al., 2023). We find that the results reported in our main analyses are qualitatively the same as the ones obtained using the two lists separately.

[7] In Appendix J, we use the full text of a subset of 12,382 theses to validate the use of titles and abstracts in identifying AI knowledge. We find that theses classified as AI theses using the title and abstract show intensive use of AI keywords throughout the full text, while theses classified as non-AI using the title and abstract are unlikely to show AI keywords in the full text. This result reassures us on the reliability of the method used to identify AI documents.

[8] The information about the precise date of student enrollment is not available. Therefore, coherently with the official statistics (Pommier et al., 2022), we assume that the student's enrollment year in the PhD program is *t-3*, where *t* is the thesis defense year.

[9] We have calculated the share of AI papers among those supervisors with the variable *AI supervisor before* equal to one. Our findings reveal that, on average, an *AI supervisor* has approximately 25 percent of her publications in the field of AI *before* the student's enrollment.

1), while only 3.57% of the pairs in which the supervisor has no AI knowledge show a student writing a thesis with AI content (*AI student* = 1). In other words, these descriptive statistics show that students are more likely to write a thesis with AI content if their supervisor has AI knowledge.

Table 1: Supervisor-to-student knowledge transfer descriptive statistics.

	AI student = 1	AI student = 0	Total
AI supervisor before = 1	1,922 (18.60%)	8,411 (81.40%)	10,333 (100%)
AI supervisor before = 0	1,090 (3.57%)	29,429 (96.43%)	30,519 (100%)
			40,852 student-supervisor pairs

Concerning the opposite direction of knowledge transfer, we study the student-to-supervisor knowledge transfer process on a sample of 30,519 student-supervisor pairs, all of which have supervisors initially lacking AI knowledge (*AI supervisor before* = 0). Within this group, 1,090 pairs are characterized by a student writing a thesis in AI (*AI student* = 1), and the remaining 29,429 pairs by a student writing a thesis not in AI (*AI student* = 0). Table 2 shows that, among the pairs characterized by students with AI knowledge, 29.36% have a supervisor publishing an AI article within three years following the student's defense date (*AI supervisor after* = 1), while among the pairs characterized by students with no AI knowledge, only 9.58% of the supervisors publishing an AI article. These descriptive statistics show that supervisors without AI knowledge are more likely to start publishing in AI after mentoring a student writing an AI thesis than supervisors mentoring students not writing an AI thesis. Table 2 also suggests that a non-negligible part of the learning process between the supervisor and the student concerns the knowledge acquired by the supervisors from students, supporting the idea that it exists a bidirectional student-supervisor learning.

Table 2: Student-to-supervisor knowledge transfer descriptive statistics.

	AI supervisor after = 1	AI supervisor after = 0	Total
AI student = 1	320 (29.36%)	770 (70.64%)	1,090 (100%)
AI student = 0	2,819 (9.58%)	26,610 (90.42%)	29,429 (100%)
			30,519 student-supervisor pairs*

NOTE: *The total number of observations reduces from 40,852 to 30,519 student-supervisor pairs because in this table we consider only pairs in which the supervisor has no AI knowledge before the student enrollment.

5. Methodology

5.1 Model

In studying the bidirectionality of learning between supervisors and students, we aim to assess the relationship between the variables *AI supervisor before* and *AI student*, i.e., *supervisor-to-student* knowledge transfer, and the relationship between *AI student* and *AI supervisor after*, i.e., *student-to-supervisor* knowledge transfer. One possible problem in identifying these relationships is that supervisors with AI knowledge might have different characteristics than supervisors without AI knowledge. Similarly, supervisors who mentor AI students might have different characteristics than those who do not. To address this issue, we employ a *Propensity Score Matching* (PSM) approach. PSM allows matching student-supervisor pairs on a range of observable characteristics to create comparable groups.

To analyze the *supervisor-to-student* knowledge transfer, we first identify the student-supervisor pairs in which the supervisor has AI knowledge before the student's enrollment (*AI supervisor before* = 1). For each of these pairs, we find a similar pair having a supervisor with no AI knowledge before the student's enrollment (*AI supervisor before* = 0). Given that the supervisor's characteristics solely determine her experience in AI before the student's enrollment, we consider the vector of observable student characteristics (X) as irrelevant for constructing our control sample. Equation 1 predicts the propensity score used to construct the control sample for the *supervisor-to-student* knowledge transfer analysis. We estimate Equation 1 using a logit model. Based on the predictions of Equation 1, each student-supervisor pair in which the supervisor has AI knowledge is matched with a pair in which the supervisor has no AI knowledge but with the closest possible propensity score value.[10]

$$\text{AI supervisor before} = \gamma_0 + \gamma_1 Z + \gamma_2 L + \gamma_3 D + \gamma_4 F + \varepsilon$$

Equation 1: Supervisor-to-student knowledge transfer PSM equation

To analyze the *student-to-supervisor* knowledge transfer, we identify the student-supervisor pairs in which students develop AI knowledge during their thesis (*AI students* = 1). Then, for each of these pairs, we find a similar pair characterized by a student with no knowledge in AI (*AI students* = 0). Contrary to the previous analysis aiming to assess the *supervisor-to-student*

[10] To construct our control samples, we use the *Nearest Neighbor* approach *without replacement*. This technique ensures that each control unit is matched exclusively to one treated unit based on the closest propensity score. Moreover, this method offers the clear advantage of reducing the bias associated with repeated use of the same control units (Smith, 1997). While Smith and Todd (2005) argue that matching without replacement might lead to increased variance in the estimator due to a potentially limited selection of matches, our dataset is sufficiently large, allowing for accurate matches without the need to reuse control units (Caliendo and Kopeinig, 2008).

knowledge transfer, in this analysis, we consider both the student (*X*) and supervisor (*Z*) observable characteristics as relevant characteristics to construct our control sample. Indeed, the acquisition of AI knowledge during the thesis can be influenced by factors related to both the student and the supervisor. Equation 2 predicts the propensity score for the *student-to-supervisor* knowledge transfer analysis. We estimate Equation 2 using a logit model. We apply the predictions of Equation 2 to match each student-supervisor pair in which the student is doing a thesis in AI with a pair in which the student is not doing a thesis in AI but has a similar propensity score value.

$$\text{AI student}_i = \delta_0 + \delta_1 X + \delta_2 Z + \delta_3 L + \delta_4 D + \delta_5 F + \varepsilon$$

Equation 2: Student-to-supervisor knowledge transfer PSM equation

5.2 Student and supervisor characteristics

In Equations 1 and 2, we include two vectors of variables, i.e., *X* and *Z,* to control for the student and supervisor characteristics. Moreover, we include controls for the location of the student and the supervisor (*L*), for the defense year (*D*), and for the research field (*F*). All the student and supervisor characteristics are calculated by adopting the student-supervisor pair as the unit of analysis.

As student characteristics (*X*), we consider the student's gender and the number of peers. We calculate the variable *Female student* that takes value one in the case of a female student and zero otherwise. We also define the variable *N. peers* as the number of students in the training period at the same time as the focal student and supervised by the focal student's supervisor.

As supervisor characteristics (*Z*), we include a dummy variable, *Female supervisor,* equal to one if the supervisor is a female scholar and zero otherwise. We also include the supervisor's *Seniority*, computed as the years elapsed from the supervisor's first publication to the year the student began the PhD program. We also calculate the supervisor's *Mentorship experience* by counting the number of students she has successfully mentored since 1990 up to the enrollment in the PhD program of the focal student. We also define the variable *N. of publications* as the number of supervisor's publications in peer-reviewed journals before the student's enrollment into the PhD program. Moreover, we explore the supervisor's research interests. We determine the supervisor's multidisciplinary nature by calculating the number of fields she published before the enrollment of the PhD student. We use *OpenAlex*'s concept classification system to determine the publication fields. This system assigns tags to each publication, identifying the specific academic fields it covers. We define the variable *N. of fields* as the number of distinct

OpenAlex's concepts across all publications attributed to each supervisor. Finally, as a proxy for the supervisor's capability in securing funding, we calculate the dummy variable *Grant*, which equals one if the supervisor serves as the principal investigator for an ANR grant before the enrollment of the PhD student in the program, and zero otherwise.

As geographical location (*L*), we consider the location of the university where the student and supervisor work. Specifically, we group universities in major French cities into distinct categories. For instance, all educational institutions in Paris are combined into a single *Paris* dummy variable. We calculated ten dummy variables, one for each of the ten cities with the largest number of PhD students in France. All other universities in France that are not located in these specific cities are classified under the residual dummy variable *Other*. In the econometric models, we consider *Paris* as the reference dummy variable. As defense years (*D*), we introduce a set of dummy variables corresponding to the students' *Defense year*. In the econometric models, we consider *2010* to be the reference dummy variable. As fields of study (*F*), we introduce eight dummy variables representing major research fields in which the supervisor mentors the student: *Mathematics, Engineering, Computer Science, Physics, Geology, Medicine, Biology, and Chemistry.* In the econometric models, we consider *Computer Science* to be the reference dummy variable.

5.3 Finding a control sample

This section has two parts. The first part regards the *supervisor-to-student* knowledge transfer analysis. It reports the matching exercise conducted to find, for each student-supervisor pair with an AI supervisor before the student's enrollment, a similar control pair without an AI supervisor. The second part regards the *student-to-supervisor* knowledge transfer analysis. It reports the matching exercise conducted to find, for each student-supervisor pair characterized by an AI student, a similar control pair with a non-AI student.

5.3.1 *Supervisor-to-student* knowledge transfer

Table 3 reports the estimates of Equation 1 used to predict the propensity score values corresponding to the probability of observing a student-supervisor pair with an AI supervisor before the student's enrollment (*AI supervisor before*=1). Female and senior supervisors are less likely to be AI supervisors. On the contrary, experience in mentoring students, conducting interdisciplinary research, and principal investigators of research grants are more likely to be AI supervisors. Not surprisingly, when analyzing the research fields, we find that supervisors in *Computer Science* are more likely to be AI supervisors than supervisors in other fields. We also find that supervisors mentoring students in recent years are more likely to be AI

supervisors. Finally, our estimates show that supervisors affiliated with universities in Rennes have the highest probability of being AI supervisors than those in other French cities, contrary to supervisors in Montpellier, who have the lowest probability of being AI supervisors.

Table 3: *Supervisor-to-student* knowledge transfer. Regression predicting the probability of observing a student-supervisor pair in which the supervisor has AI knowledge before the student's enrollment (*AI supervisor before*). Logit estimates.

Dependent variable: AI supervisor before	Estimate	Std.error	t-statistic	P-value
(Intercept)	-2.021	0.071	-28.482	0.000
Supervisor characteristics				
Female	-0.052	0.029	-1.764	0.078
Missing gender[11]	-0.023	0.060	-0.390	0.697
N. publications	0.000	0.000	0.819	0.413
Seniority	-0.009	0.001	-5.987	0.000
Grant	0.144	0.033	4.376	0.000
Experience	0.025	0.006	3.974	0.000
Interdisciplinarity	0.219	0.005	48.101	0.000
Filed				
Biology	-1.512	0.047	-32.033	0.000
Chemistry	-3.117	0.070	-44.362	0.000
Engineering	-1.239	0.042	-29.527	0.000
Geology	-2.065	0.068	-30.209	0.000
Mathematics	-1.341	0.059	-22.786	0.000
Medicine	-1.416	0.054	-26.272	0.000
Physics	-1.997	0.050	-39.641	0.000
Computer science	ref	ref	ref	ref
Location				
Bordeaux	0.248	0.071	3.498	0.000
Lille	0.106	0.070	1.507	0.132
Lyon	-0.038	0.053	-0.703	0.482
Marseille	0.015	0.072	0.213	0.831
Montpellier	-0.286	0.072	-3.975	0.000
Nantes	-0.142	0.089	-1.586	0.113
Others	-0.032	0.033	-0.965	0.335
Rennes	0.526	0.071	7.420	0.000
Strasbourg	-0.214	0.078	-2.749	0.006
Toulouse	0.106	0.053	2.007	0.045
Paris	ref	ref	ref	ref
Year				
2010	ref	ref	ref	ref
2011	0.020	0.063	0.322	0.748
2012	0.028	0.061	0.456	0.648
2013	0.112	0.060	1.880	0.060
2014	0.092	0.060	1.533	0.125
2015	0.140	0.059	2.356	0.018
2016	0.160	0.059	2.700	0.007
2017	0.233	0.059	3.933	0.000
2018	0.244	0.061	4.022	0.000
Pseudo R2	0.161			
Student-supervisor pairs	40,852			

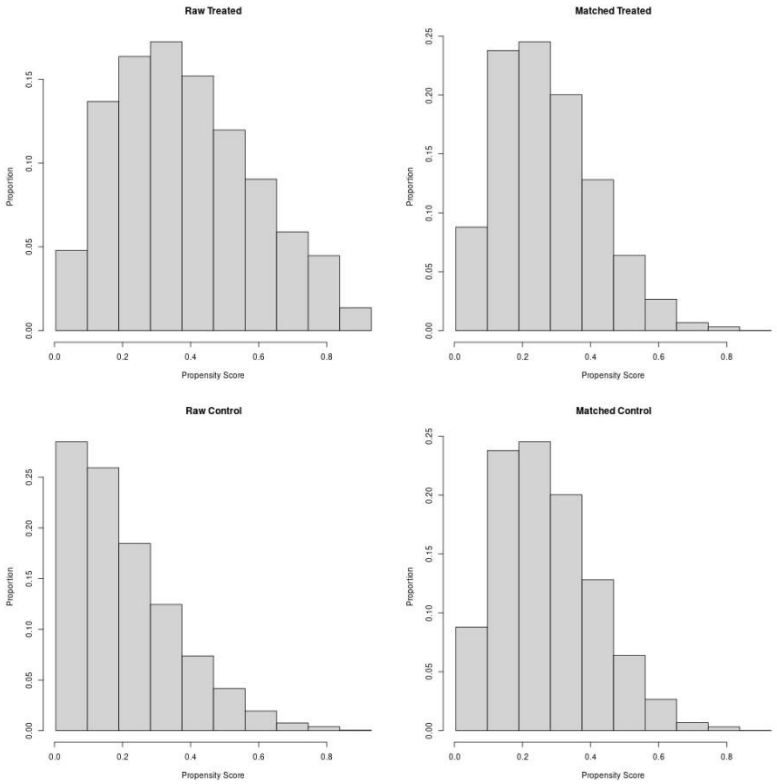
Figure 1 uses the econometric model estimated in Table 3 to plot the distribution of the propensity scores predicted for the 10,333 student-supervisor pairs having an AI supervisor before the student's enrollment (*Raw treated*, Figure 1) and the distribution of the propensity scores of the 30,519 pairs having a non-AI supervisor (*Raw control*, Figure 1). Moreover, Figure

[11] We were not able to match the gender of all students and supervisors, resulting in 8.5% and 4.7% of the pairs having missing gender information, respectively. However, we have controlled for this missing information by including *Missing gender* dummy variables in our regression.

1 reposts the distributions of the propensity scores after the matching exercise. Specifically, it reports the distributions of the propensity scores of the 5,245 pairs having an AI supervisor before the student's enrollment (*Matched treated*, Figure 1) and the matched 5,245 highly similar control pairs having a non-AI-supervisor before the student enrollment (*Matched control*, Figure 1). The reduction in the number of pairs with AI supervisors after the matching, from 10,333 (*Raw treated*) to 5,245 (*Matched treated*), is due to the introduction of a tolerance level, or "caliper," setting the maximum acceptable distance between the propensity scores of pairs with an AI supervisor and the control pairs without an AI supervisor. Introducing a maximum tolerance level aims to prevent the matching of student-supervisor pairs with an AI supervisor with highly different controls in terms of propensity score values.[12] Figure 1 shows that the distributions of the propensity scores of the 5,245 pairs with an AI supervisor (*Matched treated*, Figure 1) and the 5,245 pairs with a non-AI supervisor (*Matched control*, Figure 1) are highly similar after applying "caliper." This latter evidence is coherent with a high-quality matching exercise.

[12] We adopt a commonly used approach to implement "caliper" in our analysis. First, we define a threshold for identifying and discarding highly different combinations of treated student-supervisor pairs and their controls. To do so, we set a threshold corresponding to the value of the maximum acceptable difference between the propensity scores of the 10,333 treated student-supervisor pairs and the propensity scores of their matched controls. Then, we follow an iterative process considering a smaller threshold in each iteration until we reach a point where the treated pairs are not statistically different from their controls according to their observable characteristics. All the treated pairs with AI supervisors differing in the propensity score values to their controls more than the threshold have been discarded. This latter restriction reduces the number of student-supervisor pairs with AI supervisors considered in our analysis from 10,333 to 5,245. As a robustness check, in Appendix L, we rerun our analyses without applying the "caliper" restriction when identifying the control sample. The results of our analyses do not change.

Figure 1: *Supervisor-to-student* knowledge transfer. Propensity score matching distribution before and after the matching.



To further assess the quality of our matching exercise, Table 4 shows the average characteristics of the student-supervisor pairs before and after the matching exercise. According to the characteristics observed, the pairs having supervisors with and without AI knowledge do not show statistically significant differences after the matching (*Post-matching* panel in Table 4). The only exception is students who graduated from Lille, in which both pairs still show statistically significant differences even after the matching using "caliper" (P-value 0.09). Although statistically significant, the shares of pairs having AI and non-AI supervisors in Lille are similar[13], i.e., 0.04 versus 0.04.

[13]Appendix K performs an exact matching on the graduation location of the students. The results of our analyses do not change.

Table 4: *Supervisor-to-student* knowledge transfer. Pre- and Post-matching observable characteristics.

	Pre-matching			Post-matching		
	AI Supervisor before = 0 (30,519 pairs)	AI Supervisor before = 1 (10,333 pairs)	P-value	AI Supervisor before = 0 (5,245 pairs)	AI Supervisor before = 1 (5,245 pairs*)	P-value
Supervisor characteristics						
Female	0.32	0.25	0.00	0.29	0.28	0.33
Missing gender	0.05	0.05	0.13	0.05	0.05	0.75
N. publications	43.62	55.43	0.00	50.92	50.97	0.96
Seniority	21.76	24.84	0.00	24.51	24.38	0.52
Grant	0.17	0.20	0.00	0.19	0.19	0.71
Experience	1.21	1.84	0.00	1.27	1.28	0.77
Interdisciplinarity	9.33	12.14	0.00	11.03	11.00	0.60
Filed						
Biology	0.22	0.15	0.00	0.20	0.19	0.36
Chemistry	0.13	0.03	0.00	0.06	0.05	0.21
Engineering	0.20	0.29	0.00	0.26	0.27	0.48
Geology	0.05	0.04	0.00	0.06	0.06	0.80
Mathematics	0.06	0.07	0.04	0.07	0.07	0.97
Medicine	0.11	0.10	0.00	0.11	0.12	0.74
Physics	0.15	0.11	0.00	0.15	0.16	0.28
Computer science	0.07	0.22	0.00	0.09	0.09	0.73
Location						
Bordeaux	0.04	0.04	0.52	0.04	0.04	0.80
Lille	0.04	0.04	0.07	0.04	0.04	0.09
Lyon	0.07	0.07	0.14	0.08	0.07	0.74
Marseille	0.04	0.03	0.49	0.04	0.04	0.25
Montpellier	0.05	0.03	0.00	0.04	0.04	0.72
Nantes	0.02	0.02	0.95	0.02	0.02	0.89
Others	0.34	0.36	0.02	0.35	0.35	0.90
Rennes	0.03	0.04	0.00	0.03	0.03	0.12
Strasbourg	0.05	0.02	0.00	0.03	0.03	0.87
Toulouse	0.07	0.08	0.02	0.07	0.07	0.91
Paris	0.27	0.27	0.97	0.26	0.26	0.96
Year						
2010	0.10	0.07	0.00	0.08	0.08	0.74
2011	0.10	0.08	0.00	0.09	0.09	0.73
2012	0.11	0.09	0.00	0.10	0.11	0.15
2013	0.11	0.11	0.18	0.12	0.11	0.62
2014	0.12	0.11	0.28	0.11	0.12	0.48
2015	0.12	0.12	0.10	0.13	0.12	0.77
2016	0.12	0.13	0.00	0.13	0.13	1.00
2017	0.12	0.14	0.00	0.14	0.13	0.46
2018	0.10	0.13	0.00	0.11	0.11	0.64

NOTE: The level of observation is the student-supervisor pair. *The reduction in the number of observations is due to the use of a "caliper".

5.3.2 *Student-to-supervisor* knowledge transfer

Table 5 reports the estimates of Equation 2 used to predict the propensity score values corresponding to the probability of observing a student-supervisor pair with a student writing a thesis in AI (*AI student*=1). We find no statistically significant relationship between the student's and supervisor's gender and the student's probability of writing an AI thesis. We also find that student-supervisor pairs in which supervisors are both senior and productive show a lower probability of observing a student writing a thesis in AI. As expected, student-supervisor pairs in *Computer Science* exhibit a higher probability of observing a student writing an AI thesis than pairs in other fields, as well as students located in universities in Paris. Moreover, we observe a positive trend in the probability of writing a thesis in AI in recent years, although the coefficients of the *Defense year* dummy variables are not statistically significant.

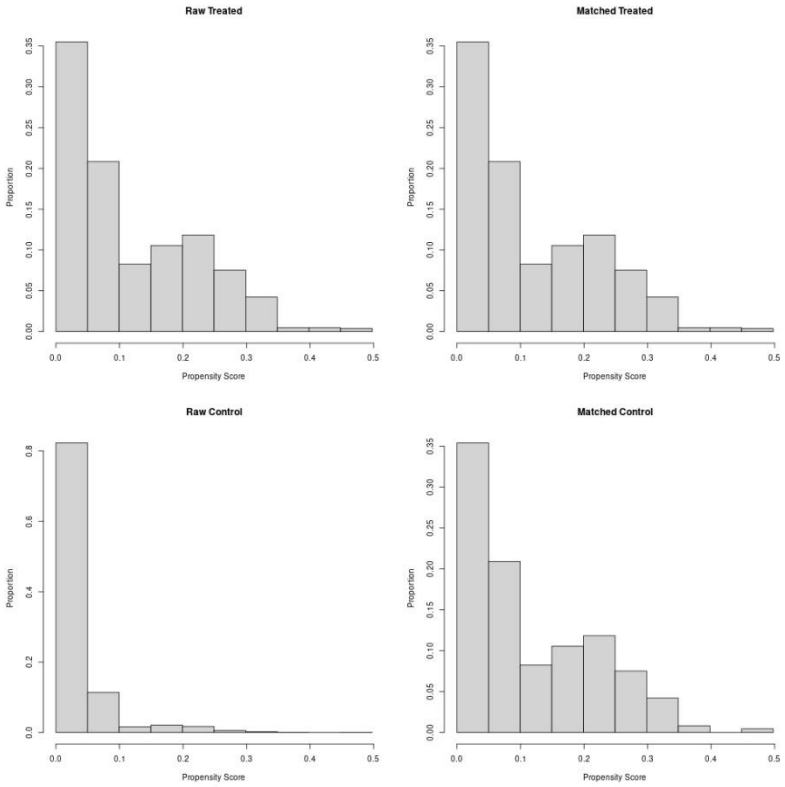
Table 5: *Student-to-supervisor* knowledge transfer. Regression predicting the probability of observing a student-supervisor pair in which the student is writing a thesis in AI (*AI student*). Logit estimates.

Dependent variable: AI student	Estimate	Std.error	t-statistic	P-value
(Intercept)	-1.168	0.148	-7.899	0.000
Student characteristics				
Female	-0.038	0.075	-0.514	0.607
Missing gender	0.094	0.108	0.869	0.385
Team size	0.046	0.018	2.506	0.012
Supervisor characteristics				
Female	0.041	0.073	0.563	0.573
Missing gender	0.048	0.146	0.328	0.743
N. publications	-0.012	0.002	-6.107	0.000
Seniority	-0.016	0.004	-3.719	0.000
Grant	0.092	0.095	0.974	0.330
Experience	0.029	0.016	1.820	0.069
Interdisciplinarity	-0.002	0.011	-0.216	0.829
Filed				
Biology	-2.639	0.138	-19.058	0.000
Chemistry	-3.671	0.279	-13.177	0.000
Engineering	-1.283	0.084	-15.250	0.000
Geology	-2.150	0.190	-11.287	0.000
Mathematics	-1.285	0.124	-10.400	0.000
Medicine	-2.330	0.164	-14.181	0.000
Physics	-2.125	0.130	-16.340	0.000
Computer science	ref	ref	ref	ref
Location				
Bordeaux	-0.167	0.208	-0.803	0.422
Lille	0.215	0.175	1.226	0.220
Lyon	-0.110	0.145	-0.761	0.447
Marseille	0.142	0.203	0.698	0.485
Montpellier	0.102	0.188	0.544	0.586
Nantes	-0.107	0.219	-0.488	0.626
Others	0.035	0.083	0.423	0.672
Rennes	-0.526	0.242	-2.172	0.030
Strasbourg	-0.050	0.220	-0.229	0.819
Toulouse	0.057	0.130	0.440	0.660
Paris	ref	ref	ref	ref
Defense year				
2010	ref	ref	ref	ref
2011	-0.197	0.145	-1.355	0.175
2012	-0.222	0.144	-1.543	0.123
2013	-0.003	0.139	-0.020	0.984
2014	-0.245	0.145	-1.690	0.091
2015	-0.006	0.139	-0.041	0.967
2016	0.032	0.142	0.223	0.824
2017	0.133	0.144	0.925	0.355
2018	0.243	0.149	1.632	0.103
Pseudo R2	0.170			
Student-supervisor pairs	30,519			

Figure 2 uses the prediction of the econometric model estimated in Table 4 to plot the distribution of the propensity predicted for the 1,090 student-supervisor pairs with an AI student (*Raw treated*, Figure 2) and the distribution of the propensity scores of the 29,429 student-supervisor pairs with a non-AI student (*Raw control*, Figure 2). Moreover, Figure 2 reports the distributions of the propensity scores after the matching exercise. Specifically, it reports the distribution of the propensity scores for the 1,090 student-supervisor pairs with an AI student (*Matched treated*, Figure 2), and the distribution of the propensity scores for the 1,090 student-supervisor pairs with a non-AI student (*Matched control*, Figure 2). After the matching, we find that the distribution of the propensity score values for the student-supervisor pairs with students writing a thesis in AI (*Matched treated*, Figure 2) is almost identical to one of the pairs with

students not writing a thesis in AI (*Matched control*, Figure 2). This latter evidence is coherent with a high-quality matching exercise.

Figure 2: *Student-to-supervisor* knowledge transfer. Propensity score matching distribution before and after the matching.



To further assess the quality of our matching exercise, Table 6 shows the average characteristics of the student-supervisor pairs before and after the matching exercise. According to the characteristics observed, the pairs with students writing a thesis in AI and those with students not writing a thesis in AI do not show statistically significant differences. Unlike the matching exercise conducted in section 5.3.1, in Table 6, we do not observe any reduction in the number of student-supervisor pairs before and after the matching because applying the "caliper" restriction was not necessary.

Table 6: Pre- and post-matching characteristics.

	Pre-matching			Post-matching		
	AI student = 0 (29,429 pairs)	AI student = 1 (1,090 pairs)	P-value	AI student = 0 (1,090 pairs)	AI student = 1 (1,090 pairs)	P-value
Student characteristics						
Female	0.39	0.29	0.00	0.27	0.29	0.44
Missing gender	0.08	0.11	0.01	0.11	0.11	0.94
Team size	2.12	2.50	0.00	2.60	2.50	0.26
Supervisor characteristics						
Female	0.32	0.30	0.12	0.29	0.30	0.81
Missing gender	0.05	0.05	0.28	0.06	0.05	0.31
N. publications	44.55	18.50	0.00	19.95	18.50	0.19
Seniority	21.99	15.58	0.00	15.44	15.58	0.79
Grant	0.17	0.14	0.00	0.14	0.14	0.80
Experience	1.19	1.54	0.00	1.58	1.54	0.75
Interdisciplinarity	9.41	7.21	0.00	7.25	7.21	0.86
Filed						
Biology	0.22	0.07	0.00	0.07	0.07	0.80
Chemistry	0.14	0.01	0.00	0.01	0.01	0.71
Engineering	0.20	0.28	0.00	0.28	0.28	0.74
Geology	0.06	0.03	0.00	0.02	0.03	0.42
Mathematics	0.06	0.08	0.00	0.07	0.08	0.38
Medicine	0.11	0.04	0.00	0.04	0.04	0.75
Physics	0.15	0.08	0.00	0.09	0.08	0.19
Computer science	0.06	0.40	0.00	0.41	0.40	0.63
Location						
Bordeaux	0.04	0.03	0.02	0.02	0.03	0.79
Lille	0.03	0.04	0.44	0.04	0.04	0.82
Lyon	0.07	0.06	0.07	0.07	0.06	0.54
Marseille	0.04	0.03	0.11	0.02	0.03	0.15
Montpellier	0.05	0.03	0.02	0.04	0.03	0.72
Nantes	0.02	0.02	0.56	0.02	0.02	0.29
Others	0.34	0.41	0.00	0.41	0.41	0.86
Rennes	0.03	0.02	0.02	0.02	0.02	0.62
Strasbourg	0.05	0.02	0.00	0.02	0.02	0.77
Toulouse	0.07	0.08	0.08	0.08	0.08	0.75
Paris	0.27	0.25	0.26	0.27	0.25	0.24
Year						
2010	0.09	0.10	0.29	0.10	0.10	0.57
2011	0.10	0.09	0.44	0.09	0.09	0.60
2012	0.12	0.10	0.04	0.11	0.10	0.32
2013	0.11	0.12	0.86	0.14	0.12	0.14
2014	0.12	0.10	0.04	0.09	0.10	0.37
2015	0.12	0.13	0.45	0.12	0.13	0.60
2016	0.12	0.12	0.93	0.12	0.12	0.64
2017	0.12	0.13	0.37	0.13	0.13	0.85
2018	0.10	0.12	0.11	0.12	0.12	0.89

NOTE: The level of observation is the student-supervisor pair.

6. Results

This section presents the results of the empirical analysis aiming to investigate the bidirectional knowledge transfer (Section 6.1), provides further empirical evidence on the mechanisms by which the knowledge transfer happens (Section 6.2), and shows how the bidirectional knowledge transfer works in different fields (Section 6.3).

6.1 Bidirectional knowledge transfer

6.1.1 *Supervisor-to-student* knowledge transfer

Table 7 assesses the probability of writing an AI thesis for students in student-supervisor pairs having an AI supervisor and the probability of writing an AI thesis for students for a control sample of similar pairs having a non-AI supervisor. We find that students in pairs with an AI supervisor have a probability of 14% of writing a thesis in AI, compared to the 4% probability in pairs with non-AI supervisors. The statistically significant difference between the two probabilities, i.e., 10 percentage points, indicates an AI knowledge transfer from the supervisor to the student.

Table 7: Student's probability of writing a thesis in AI.

	Pr(AI student = 1)
AI Supervisor before = 1 (5,245 student-supervisor pairs)	0.14
AI Supervisor before = 0 (5,245 student-supervisor pairs)	0.04
Difference	0.10
P-value	<0.01

6.1.2 *Student-to-supervisor* knowledge transfer

Table 8 assesses the probability that in student-supervisor pairs in which supervisors have no previous AI knowledge, supervisors initiate AI research within the three years after their student's defense year. We find that for pairs in which the student writes an AI thesis, the probability of the supervisor publishing in AI within the three-year window following the student's defense is 32%, compared to a lower probability of 13% for supervisors in pairs where the student does not write an AI thesis. The statistically significant difference between the two probabilities, i.e., 19 percentage points, indicates knowledge transfer from the student to the supervisor.[14]

[14] In our analysis, we proxy the student-to-supervisor knowledge transfer with the variable *AI Supervisor after* that measures the knowledge transfer in a time window of three years after the student's defense (*t*), i.e., from *t+1* to *t+3*. The three-year time window is arbitrarily defined. Therefore, in Table C1 in Appendix C, we investigate the year-by-year student-to-supervisor knowledge transfer from *t-3* to *t+5*. We find that the student-to-supervisor knowledge transfer starts soon after the student's enrollment and persists until t+5.

Table 8: Supervisor's probability of having an AI publication three years after the student's defense year.

	Pr(AI Supervisor after = 1)
AI student = 1 (1,090 student-supervisor pairs)	0.32
AI student = 0 (1,090 student-supervisor pairs)	0.13
Difference	0.19
P-value	<0.01

6.2 Mechanisms behind the bidirectional knowledge transfer

When analyzing the student-to-supervisor knowledge transfer, it is difficult to distinguish whether the AI knowledge transfer to the supervisor results from hiring a PhD student with previous AI knowledge or from an AI co-learning of the supervisor and student during the PhD training period. The data do not allow us to disentangle the two types of learning processes clearly. However, we can use the timing of the learning to conjecture which of the two learning processes, *learning-by-hiring* or *co-learning*, is at play. Observing a supervisor publishing articles with AI content during their students' training period may suggest a learning-by-hiring process, as the student enrolled is already equipped with ready-to-use AI knowledge. On the contrary, if the supervisor publishes articles with AI content only after the student's PhD defense, the supervisor's AI knowledge acquisition is more likely to result from co-learning with the student during the PhD period since the learning process takes time. To assess the timing of the supervisor's AI knowledge acquisition, in Appendix C, we conduct a year-by-year analysis, assessing the student-to-supervisor knowledge transfer from the year of the student's enrollment (*t-3*) to five years after the student's defense (*t+5*). Table C1 shows that there is already a significant student-to-supervisor AI knowledge transfer soon after the student's enrollment (*t-3*) and during the PhD period, probably highlighting a learning-by-hiring knowledge acquisition process. After graduation from *t* to *t+5*, the student-to-supervisor AI knowledge transfer further increases, probably highlighting a co-learning knowledge acquisition process that took time to transfer the student's AI knowledge to the supervisor. Taken together, our results do not allow us to distinguish between *learning-by-hiring* and *co-learning* clearly. Nonetheless, they are consistent with both learning mechanisms.

The ideal way to identify if a *learning-by-hiring* process is in place would be to assess whether the student's AI knowledge is acquired before being enrolled in the PhD program. This approach would require obtaining information about AI knowledge acquisition before the PhD, for instance, during their master's studies. Unfortunately, we do not have precise information on students before enrolling in the PhD program. Nonetheless, in Appendix N, we use publication data to calculate a proxy for the student's AI knowledge acquisition before enrollment. Specifically, we calculate the variable *Student AI knowledge before enrollment,*

which equals one if the student has at least one AI publication before enrollment[15] and zero otherwise. Among the 40,852 student-supervisor pairs in our study sample used to assess the supervisor-to-student AI knowledge transfer, only 0.7% of the students have a publication in AI before enrollment in the PhD program. When we restrict the study sample to the 30,519 pairs used to assess the student-to-supervisor AI knowledge transfer (i.e., the sample in which *AI supervisor before* = 0), only 0.5% of the students have a publication in AI before enrollment in the PhD program. This latter figure shows that, according to our proxy, having AI knowledge before enrollment is rare and, consequently, the potential pool of students from which a supervisor could learn from is small. Hence, we expect that learning-by-hiring through students with previous AI knowledge is likely to be limited on average.. Although rare, we expect that the few students with AI knowledge before enrollment will leverage this knowledge for writing their PhD thesis. In Table N1, we add the variable *Student AI knowledge before enrollment* to Equation 2. As expected, the estimated coefficient of the variable *Student AI knowledge before enrollment* shows a positive and statistically significant relationship between having AI knowledge before PhD and the probability of writing an AI thesis. Moreover, Table N2 shows that when assessing the student-to-supervisor knowledge transfer, including the variable *Student AI knowledge before enrollment* among the matching criteria does not change the result of our analysis presented in Table 8.

Another crucial aspect of student-to-supervisor knowledge transfer concerns the supervisor's actual acquisition of AI knowledge from the student. Indeed, the supervisor can publish articles with AI content after the student's graduation by continuing to collaborate with the student but without actively acquiring AI knowledge. To rule out this possibility, in Appendix F, we conduct a set of analyses in which we exclude all the supervisors' publications co-authored with their students in the three years after the defense. Table F1 shows evidence of an actual acquisition of AI knowledge by the supervisor since, after excluding the publications co-authored with the students, the difference between the probability of publishing an AI article remains 11 percentage points higher for those supervisors mentoring an AI student than for supervisors mentoring a non-AI student. Therefore, we can conclude that the use of AI in the supervisor's work is not the result of the prosecution of the thesis work with the student, but it is an actual acquisition of AI knowledge.

Another concern regards the possibility that the observed AI knowledge flows between student and supervisor originate from other researchers. For instance, the supervisor's co-authors may

[15] We intend the period before enrollment as the period before *t-3*, where *t* is the defense date.

be a source of AI knowledge. To rule out this possibility we conduct the analysis reported in Appendix M in which we consider the AI knowledge of the supervisor's co-authors as a matching criterion. We observe that our results are consistent with the ones observed in Tables 7 and 8.

One of the main contributions of our analysis to the literature is tracing the content of the knowledge learned by the supervisor and the student. We trace the flow between individuals of a specific piece of knowledge, i.e., AI. However, AI is a broad concept encompassing different types of AI such as *Artificial Neural Networks*, *Machine Learning*, *Naive Bayes Classifiers*, *Natural Language Processing*, *Symbolic Reasoning*, and many others. In our analysis, we classified AI documents using 272 keywords representing different types of AI but without focusing on which type of AI is transferred. Appendix D investigates to what extent the supervisor-student bidirectional learning regards the exact same piece of AI knowledge. In other words, we investigate to what extent the supervisors and students use the same AI keyword when they learn from each other. Table D1 shows that 33.26% of the students writing a thesis in AI use the same AI keyword appearing in their supervisors' publications before their enrollment in the PhD program. We also find that 33.89% of the supervisors who learn AI from their students use the same AI keywords that appear in their students' theses. These two relatively high percentages show that the learning process often regards the exact same piece of AI knowledge.

6.3 Bidirectional knowledge transfer in different fields

We analyze the supervisor-student bidirectional learning process across research fields over a long period. In doing so, we do consider the research field and time heterogeneity, which might provide new nuances to our results.

Appendix G assesses how supervisor-student bidirectional learning varies across research fields. In Table G1, we find a high probability of supervisor-to-student knowledge transfer in *Computer Science*, a research field in which AI is more likely to be the central research interest, and a low probability of student-to-supervisor knowledge transfer (Table G2). On the contrary, in *Biology*, a research field in which AI is more likely to be used as a tool for research, we find no evidence of supervisor-to-student knowledge transfer and a high probability of student-to-supervisor knowledge transfer. One possible explanation of this result is that in fields where AI is likely to be the supervisors' central research interest, supervisors transfer AI knowledge to the students. In contrast, in fields in which AI is used as a tool, AI is less likely to be the supervisor's central research interest, making supervisors more likely to learn AI knowledge

from their students. Crucial to this interpretation is to show that in *Computer Scienc*e, AI knowledge is more likely to be the central research interest than in other fields like *Biology*, in which AI is more likely to be used as a tool. To differentiate how AI is considered across research fields, Appendix I assesses the intensity of AI keywords in theses in different fields. The most intensive use of AI keywords is in the full text of theses in *Computer Science*, coherently with observing theses focusing on AI as the central research interest, while the intensity of use of AI keywords is lower in *Biology*, coherently with observing theses using AI as a tool.

In addition to the research field, another dimension according to which the bidirectional supervisor-student knowledge transfer might vary is time. The type and relevance of the use of AI in scientific research has evolved over time, booming in more recent years. Appendix H assesses the bidirectional supervisor-student knowledge transfer by the student's graduation cohort. Tables H1 and H2 show that both supervisor-to-student and student-to-supervisor knowledge transfers do not show any time trend. The absence of trends might be interpreted as evidence that our results are not a consequence of AI's recent surge.

7. Conclusion

This paper studies the bidirectional knowledge transfer between students and supervisors. Specifically, it focuses on the transfer of AI knowledge between French PhD students and their supervisors for 40,852 student-supervisor pairs during the period 2010-2018. To assess knowledge transfer, we identified student-supervisor pairs characterized by students with AI knowledge, i.e., AI students, by looking at the content of their theses, and pairs characterized by supervisors with AI knowledge, i.e., AI supervisors before the student's enrollment, by looking at the content of their publications. Then, we estimated the *supervisor-to-student* knowledge transfer by calculating the probability of observing a student writing a thesis in AI. We find that pairs with AI supervisors are 10 percentage points more likely to show a student writing a thesis in AI than pairs with no AI supervisors, denoting an AI knowledge transfer from supervisors to students. We also estimated the *student-to-supervisor* knowledge transfer by calculating the probability that the supervisor publishes AI articles after the student's thesis defense. We find evidence that supervisors in pairs with AI students are 19 percentage points more likely to publish an AI article in the three years following the student's defense than supervisors in pairs with non-AI students, denoting an AI knowledge transfer from students to supervisors.

This paper adds to the unidirectional knowledge transfer perspective prevalent in the higher education literature. It shows that students can bring knowledge to their supervisors, illustrating the bidirectionality of knowledge transfer in the PhD training context. It also contributes to the literature on academic organizations, showing that PhD students can be a vehicle to explore and update the knowledge basis of their supervisors and, in general, of a research team. Our findings can also be extended to the organizational learning literature, which studies "learning-by-hiring" as an exploration route. Specifically, we draw on microlens to illustrate how learning-by-hiring functions at the individual level in the unexplored case of trainees' hiring. Moreover, instead of focusing on the supervision styles and skills as in the previous higher education literature, this paper measures the transfer of a specific piece of knowledge, i.e., AI knowledge, between supervisors and students. Finally, our paper contributes to the literature on the AI labor market, explicitly investigating the production of AI talents in a big country for the first time.

Our results are relevant for policymakers who want to steer the research agenda of a country toward specific topics. They show two main channels through which AI can permeate the research system of a country. One channel is to incentivize established researchers to develop AI knowledge, for instance, designing thematic research grants[16], setting an AI-centered research agenda for existing labs[17], or creating new labs specialized in AI (Savage, 2020). Responding to these incentives, established researchers will develop AI knowledge that will be transmitted to new generations of university researchers, i.e., their PhD students who pursue an academic career, and industrial scientists, i.e., their PhD students hired in private companies. This supervisor-to-student knowledge transfer will spread AI research within the country's public and private systems. Another channel through which policymakers can permeate a country's research system with AI is the AI education of young generations of undergraduate students. Indeed, our results show that individuals acquiring AI knowledge before enrolling in a PhD program or during the PhD period can transmit this knowledge to their supervisors, steering supervisors' research efforts toward AI. This student-to-supervisor knowledge transfer can foster the adoption of AI by established researchers.

Our paper is not exempt from limitations. The first limitation regards the use of keywords to identify AI documents. Although easy to implement and used in preceding works (Cockburn et

[16] In the French context, the *Agence Nationale de Recherche* (ANR) is co-financing 200 PhD contracts in AI at the national level in 2022: https://anr.fr/en/call-for-proposals-details/call/call-for-programmes-phd-contracts-in-artificial-intelligence/

[17] In the French context, the Interdisciplinary Institutes for Artificial Intelligence (3IA) is securing support for research specialized centers in AI: https://instituts-3ia.fr/

al., 2018; Baruffaldi et al., 2020; Calvino et al., 2023), this approach might undermine the identification of AI documents if the words used to identify AI concepts are not reported in our list (Appendix A). A future extension of this work is to use machine learning techniques to interpret the semantic meaning of documents independently from the specific words used to express AI-related concepts. The second limitation concerns our proxy for determining whether a supervisor had AI knowledge before their student's enrollment. Our approach relies on publication records, which may not be sufficient to capture the supervisor's AI expertise. A third limitation regards the lack of student information before their enrollment in the PhD program. For instance, we do not know if students acquired AI knowledge before starting their PhD. A possible way to address the second and third limitations is to conduct interviews or survey a subset of students and supervisors to ask where and when they acquired their AI knowledge.

References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. Journal of Labor Economics, 40(S1), S293-S340.

Amara, N., Landry, R., Becheikh, N., & Ouimet, M. (2008). Learning and novelty of innovation in established manufacturing SMEs. *Technovation*, *28*(7), 450-463.

Austin, A. E. (2002). Preparing the next generation of faculty: Graduate school as socialization to the academic career. *The journal of higher education*, *73*(1), 94-122.

Ayoubi, C., Pezzoni, M., & Visentin, F., (2017). At the origins of learning: Absorbing knowledge flows from within the team. *Journal of Economic Behavior & Organization* 134.

Baruffaldi, S., van Beuzekom, B., Dernis, H., Harhoff, D., Rao, N., Rosenfeld, D., & Squicciarini, M. (2020). Identifying and measuring developments in artificial intelligence: Making the impossible possible.

Bianchini, S., Müller, M., & Pelletier, P. (2022). Artificial intelligence in science: An emerging general method of invention. Research Policy, 51(10), 104604.

Bozeman, B., & Corley, E. (2004). Scientists' collaboration strategies: implications for scientific and technical human capital. *Research policy*, *33*(4), 599-616.

Buenstorf, G., & Heinisch, D. P. (2020). When do firms get ideas from hiring PhDs?. *Research Policy*, *49*(3), 103913.

Caliendo, M., & Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of economic surveys*, *22*(1), 31-72.

Calvino, F., Criscuolo, C., Dernis, H., & Samek, L. (2023). What technologies are at the core of AI?: An exploration based on patent data.

Campbell, R. A. (2003). Preparing the next generation of scientists: The social process of managing students. *Social studies of Science*, *33*(6), 897-927.

Carayol, N., & Matt, M. (2006). Individual and collective determinants of academic scientists' productivity. *Information Economics and policy*, *18*(1), 55-72.

Cockburn, I. M., Henderson, R., & Stern, S. (2018). The impact of artificial intelligence on innovation: An exploratory analysis. In *The economics of artificial intelligence: An agenda* (pp. 115-146). University of Chicago Press.

Colbeck, C. L. (1998). Merging in a seamless blend: How faculty integrate teaching and research. *The journal of higher education*, *69*(6), 647-671.

Conti, A., & Liu, C. C. (2015). Bringing the lab back in: Personnel composition and scientific output at the MIT Department of Biology. *Research Policy*, *44*(9), 1633-1644.

Cox, R., 1988. Postgraduate Research Training, in: Cox, R. (Ed.), Reviews of Literature and Data Sources: The Characteristics of the Training Process and Those Undergoing the Research Training. Institute of Education, London, pp. 37-45.

Dasgupta, P., & David, P. A. (1994). Toward a new economics of science. *Research policy*, *23*(5), 487-521.

Delamont, S., & Atkinson, P. (2001). Doctoring uncertainty: Mastering craft knowledge. *Social studies of science*, *31*(1), 87-107.

Delamont, S., Parry, O., & Atkinson, P. (1997). Critical mass and pedagogic continuity: Studies in academic habitus. *British Journal of Sociology of Education*, *18*(4), 533-549.

Flanagan, K., Ribeiro, B., & Ferri, P. (2023). AI and scientific productivity: Considering policy and governance challenges.

Fleming, L., & Sorenson, O. (2003). Navigating the technology landscape of innovation. *MIT Sloan Management Review*.

Foster, J. G., Rzhetsky, A., & Evans, J. A. (2015). Tradition and innovation in scientists' research strategies. *American sociological review*, *80*(5), 875-908.

Franzoni, C., Scellato, G., & Stephan, P. (2018). Context factors and the performance of mobile individuals in research teams. *Journal of Management Studies*, *55*(1), 27-59.

Gupta, A. K., Smith, K. G., & Shalley, C. E. (2006). The interplay between exploration and exploitation. *Academy of management journal*, *49*(4), 693-706.

Gurr, G. M. (2001). Negotiating the" Rackety Bridge"—a dynamic model for aligning supervisory style with research student development. *Higher Education Research & Development*, *20*(1), 81-92.

Hacket, E. J. (1990). Science as a Vocation in the 1990s. The Changing Organizational Culture of Academic Science. *The Journal of Higher Education*, *61*(3), 241-279.

Halse, C. (2011). 'Becoming a supervisor': the impact of doctoral supervision on supervisors' learning. Studies in higher education, 36(5), 557-570.

Heinze, T., Shapira, P., Rogers, J.D., Senker, J.M., 2009. Organizational and Institutional Influences on Creativity in Scientific Research. *Research Policy* 38, 610-623.

Hilmer, M. J., & Hilmer, C. E. (2011). Is it where you go or who you know? On the relationship between students, Ph. D. program quality, dissertation advisor prominence, and early career publishing success. Economics of Education Review, 30(5), 991-996.

Hon Kam, B. (1997). Style and quality in research supervision: the supervisor dependency factor. *Higher education*, *34*(1), 81-103.

Igami, M., Nagaoka, S., & Walsh, J. P. (2015). Contribution of postdoctoral fellows to fast-moving and competitive scientific research. *The Journal of Technology Transfer*, *40*, 723-741.

Jain, A. (2016). Learning by hiring and change to organizational knowledge: Countering obsolescence as organizations age. *Strategic Management Journal*, *37*(8), 1667-1687.

Kahn, S., & Ginther, D. K. (2017). The impact of postdoctoral training on early careers in biomedicine. *Nature biotechnology*, *35*(1), 90-94.

Katz, J. S., & Martin, B. R. (1997). What is research collaboration?. Research policy, 26(1), 1-18.

Knorr-Cetina, K., 1999. Epistemic Cultures : How the Sciences Make Knowledge. Harvard University Press, Cambridge, MA.

Kyvik, S., & Smeby, J. C. (1994). Teaching and research. The relationship between the supervision of graduate students and faculty research performance. *Higher Education*, *28*, 227-239.

Larivière, V. (2012). On the shoulders of students? The contribution of PhD students to the advancement of knowledge. Scientometrics, 90(2), 463-481.

Latour, B., Woolgar, S., 1979. Laboratory Life: The Construction of Scientific Facts. Princeton University Press, Princeton, NJ.

Laudel, G., & Glaser, J. (2008). From apprentice to colleague: The metamorphosis of early career researchers. Higher education, 55, 387-406.

Lee, A. (2008). How are doctoral students supervised? Concepts of doctoral research supervision. *Studies in Higher education*, *33*(3), 267-281.

Manuel, N. (2024). Does the labour market value field of study specific knowledge? An alignment score based approach. Economics of Education Review, 101, 102561.

Marsh, H.W., Rowe, K.J., Martin, A., 2002. Phd Students' Evaluations of Research Supervision - Issues, Complexities, and Challenges in a Nationwide Australian Experiment in Benchmarking Universities. *Journal of Higher Education* 73, 313-348.

McFadyen, M. A., & Cannella Jr, A. A. (2004). Social capital and knowledge creation: Diminishing returns of the number and strength of exchange relationships. *Academy of management Journal*, *47*(5), 735-746.

Owen-Smith, J. (2001). Managing laboratory work through skepticism: Processes of evaluation and control. *American Sociological Review*, 427-452.

Paglis, L.L., Green, S.G., Bauer, T.N., 2006. Does Adviser Mentoring Add Value? A Longitudinal Study of Mentoring and Doctoral Student Outcomes. *Research in Higher Education* 47, 451-476.

Patsali, S., Pezzoni, M., & Visentin, F. (2024). Research independence: drivers and impact on PhD students' careers. *Studies in Higher Education*, 1-24.

Pommier, S., Talby, M., Auffray-Seguette, M., Dalaut, M., Eijsberg, H., Elshawish, P., & Muller, H. (2022). Le doctorat en France Regards croisés sur la formation doctorale (Doctoral dissertation, Réseau National de Collèges Doctoraux).

Priem, J., Piwowar, H., & Orr, R. (2022). OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts. ArXiv. https://arxiv.org/abs/2205.01833.

Savage, N. (2020). The race to the top among the world's leaders in artificial intelligence. *Nature*, *588*(7837), S102-S102.

Shibayama, S. (2019). Sustainable development of science and scientists: Academic training in life science labs. *Research Policy*, *48*(3), 676-692.

Shibayama, S. (2022). Development of originality under inbreeding: A case of life science labs in Japan. *Higher Education Quarterly*, *76*(1), 63-75.

Shibayama, S., Baba, Y., & Walsh, J. P. (2015). Organizational design of university laboratories: Task allocation and lab performance in Japanese bioscience laboratories. *Research Policy*, *44*(3), 610-622.

Slavova, K., Fosfuri, A., & De Castro, J. O. (2016). Learning by hiring: The effects of scientists' inbound mobility on research performance in academia. *Organization Science*, *27*(1), 72-89.

Smith, H. L. (1997). Matching with multiple controls to estimate treatment effects in observational studies. *Sociological methodology*, *27*(1), 325-353.

Smith, J. A., & Todd, P. E. (2005). Does matching overcome LaLonde's critique of nonexperimental estimators?. *Journal of econometrics*, *125*(1-2), 305-353.

Stephan, P. (2012). *How economics shapes science*. Harvard University Press, Cambridge, MA.

Stephan, P. E. (2001). Educational implications of university–industry technology transfer. *The Journal of Technology Transfer*, *26*, 199-205.

Stephan, P., Levin, S., 1992. Striking the Mother Lode in Science: The Importance of Age, Place, and Time. Oxford University Press, New York.

Sumell, A. J., Stephan, P. E., & Adams, J. D. (2009). Capturing knowledge: The location decision of new Ph. Ds working in industry. In *Science and engineering careers in the United States: An analysis of markets and employment* (pp. 257-287). University of Chicago Press.

Tzabbar, D. (2009). When does scientist recruitment affect technological repositioning?. *Academy of Management Journal*, *52*(5), 873-896.

Vannuccini, S., & Prytkova, E. (2023). Artificial Intelligence's new clothes? A system technology perspective. Journal of Information Technology, 02683962231197824.

Wang, J., & Shibayama, S. (2022). Mentorship and creativity: Effects of mentor creativity and mentoring style. *Research Policy*, *51*(3), 104451.

Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science*, *316*(5827), 1036-1039.

Yoshioka-Kobayashi, T., & Shibayama, S. (2021). Early career training and development of academic independence: a case of life sciences in Japan. *Studies in Higher Education*, *46*(12), 2751-2773.

Zanchetta, M. S., Bailey, A., Kolisnyk, O., Baku, L., Schwind, J., Osino, E., ... & Yu, L. (2017). Mentors' and mentees' intellectual-partnership through the lens of the Transformative Learning Theory. *Nurse education in practice*, *25*, 111-120.

Zucker, L. G., & Darby, M. R. (1997). Present at the biotechnological revolution: transformation of technological identity for a large incumbent pharmaceutical firm. *Research Policy*, *26*(4-5), 429-446.

Zuckerman, H., 1988. The Sociology of Science, in: Smelser, N.J. (Ed.), Handbook of Sociology. Sage, Newbury Park, CA, pp. 511-574.

Online Material

Appendix A: AI theses identification methodology

Economists and technologists have proposed different definitions of AI (Agrawal et al., 2019). The complexity in defining AI lies in the fact that AI is not a homogeneous technology, but it encompasses a broad range of technologies (Vannuccini and Prytkova 2023). For instance, Cockburn et al. (2018) identified three technological fields within AI, i.e., symbolic systems, robotics, and machine learning, while other definitions encompass a broader set of technologies (Baruffaldi et al., 2020; Calvino et al., 2023). In this paper, we identify AI thesis documents using keywords. Specifically, we define a document as having AI content if it has at least one AI keyword in the title or abstract.

The selection of the AI keywords is crucial for our strategy of identification of AI documents. To create a reliable list of AI keywords we merge two widely recognized existing lists. The first list of 189 AI keywords is obtained from a recent OECD report aiming to develop a methodology to identify publications and patents with AI content (Baruffaldi et al., 2020; Calvino et al., 2023). The second list of 39 AI keywords is obtained from a highly relevant economic study aiming to identify publications in AI (Cockburn et al., 2018). Merging these two lists, we obtain a list of 215 unique AI keywords. We refine the list of AI keywords by adding plural versions of the keywords that were only included in their singular form. For instance, we added to the list "symbolic error analyses" as the plural of the keyword "symbolic error analysis" included in the Cockburn et al. (2018) list. We ended up with a list of 281 AI keywords. The complete list of keywords is reported in Table A1.

We use the 281 AI keywords to identify AI documents in the universe of 40,852 PhD theses in STEM in France from 2010 to 2018. We find that 8,564 theses include at least one AI keyword, 20.96% of the overall number of theses. This high proportion of theses classified as AI is due to the inclusion in the AI keywords of generic terms not always pertinent to AI. For instance, the term "robotics" might encompass documents with limited AI content (Baruffaldi et al., 2020). Therefore, we need to refine our AI keyword selection. To do so, we assess three figures emerging from our analysis of the thesis data. Specifically, for each AI keyword, we calculate the number (and share) of theses in which the keyword appears, the probability of co-occurrence of an AI keyword with at least another AI keyword in the same document, and the share of AI keywords by field. The rationale of analyzing these three figures is that an AI keyword that allows the identification of a high number of theses, not often co-occurring with other AI

keywords, and is frequent in fields other than computer science, is a keyword that is too generic to be used to identify AI documents and, consequently, the keyword must be removed from our list.

We use Table A2 to identify keywords that are too generic to be used to identify AI documents. To calculate Table A2, we first aggregate the keywords that refer to the same concept, calculating the figures for each concept. We label the concept by using the simplest singular keyword among those referring to it. For instance, the terms “robotics”, “robot”, “robots”, and “robotic” refer to the same concept and, in Table A2, we label the concept as “robot”. Columns 1 and 2 in Table A2, show the number of theses classified as AI using the focal AI concept, and the proportion of AI theses identified with that concept over all the AI theses. Column 3 measures the co-occurrence among concepts, showing the proportion of theses in which the focal AI concept appears with at least another AI concept. Finally, Columns 4-11, show the proportion of theses identified as AI by the focal concept, by field. We define the thesis field using the official classification used by ABES.

According to Table A2, we select five concepts that include keywords too generic to identify AI documents: "algorithm", "network", "mapping", "robot", and "decision making". The concepts "algorithm" and "network" are present in the 33.72% and 19.78% of the theses identified as AI theses (Column 2), respectively. These high values show that "algorithm" and "network" are frequently used to identify AI theses. Moreover, they co-occur with other AI concepts less often than other concepts (Column 3) and are frequent also in fields other than Computer Science, e.g., in Engineering, Physics, Mathematics, and Medicine (Columns 4-11). These figures, lead us to consider "algorithm" and "network" as generic concepts and, consequently, to remove the keywords associated with these concepts from the list of AI keywords. Although the other three concepts, i.e., “robot”, “mapping”, and “decision making” are less frequently used to identify AI theses (Columns 1 and 2), they present the same pattern of relatively low co-occurrence with other AI concepts, and a wide distribution across fields as "algorithm" and "network". Therefore, we remove also the keywords associated with the concepts “robot”, “mapping”, and “decision making” from our AI keyword list. The 6th most frequent concept used to identify AI documents is “machine learning”. Contrary to the previous concepts, it is used to identify 2.29% of the AI theses, it often co-occurs with other AI concepts (85.45% of the time) and more than half (60.37%) of the AI theses identified with this concept are in Computer Science. Therefore, we keep the keywords associated with “machine learning”

in our AI keyword list as well as all the keywords appearing in all the other less frequent concepts reported in Table A2.

After removing, all the keywords related to the concepts "algorithm", "network", "mapping", "robot", and "decision making", we ended up with a list of 272 keywords. Using this latter list of keywords to identify AI theses, we find 3,012 theses including at least one AI keyword, 7.37% of the overall number of theses.

Table A1: Full list of keywords

action recognition	cognitive computing	feature extraction	image matching	machine learning	particle swarm optimisation	similarity learning
activity recognition	cognitive modelling	feature learning	image processing	machine translation	pattern analysis	simultaneous localisation
adaboost	collaborative filtering	feature selection	image recognition	machine vision	pattern recognition	social robot
adaptive boosting	collaborative systems	fuzzy c	image retrieval	mapping	pedestrian detection	social robots
adversarial network	collision avoidance	fuzzy logic	image segmentation	mapreduce	physical symbol system	sparse representation
adversarial networks	community detection	fuzzy number	independent component analyses	markovian	q learning	spectral clustering
algorithm	computational intelligence	fuzzy numbers	independent component analysis	memetic algorithm	random field	speech recognition
algorithms	computer vision	fuzzy set	inductive monitoring	memetic algorithms	random fields	speech recognitions
ambient intelligence	convolutional neural network	fuzzy sets	industrial robot	meta learning	random forest	speech to text
ant colony	convolutional neural networks	gaussian mixture model	industrial robots	motion planning	random forests	stacked generalization
ant colony optimisation	crowdsourcing and human computation	gaussian mixture models	instance based learning	multi agent system	rankboost	statistical relational learning
artificial bee colony algorithm	cyber physical system	gaussian process	intelligence augmentation	multi agent systems	recommender system	stochastic gradient
artificial intelligence	cyber physical systems	generative adversarial network	intelligent agent	multi label classification	recommender systems	supervised learning
artificial neural network	data mining	generative adversarial networks	intelligent agents	multi label classifications	recurrent neural network	support vector machine
artificial neural networks	decision making	genetic algorithm	intelligent classifier	multi layer perceptron	recurrent neural networks	support vector machines
association rule	decision tree	genetic algorithms	intelligent classifiers	multi objective evolutionary	regression tree	support vector regression
autoencoder	decision trees	genetic programming	intelligent infrastructure	multi objective optimisation	regression trees	support vector regressions
autoencoders	deep belief network	gesture recognition	intelligent infrastructures	multi objective optimisations	reinforcement learning	swarm intelligence
autonomic computing	deep belief networks	gradient boosting	intelligent software agent	multi sensor fusion	relational learning	swarm optimisation
autonomous vehicle	deep convolutional neural	gradient tree boosting	intelligent software agents	multi sensor fusions	robot	symbol processing
autonomous vehicles	deep learning	graphical model	k means	multi task learning	robot systems	symbolic error analyses
backpropagation	deep neural network	graphical models	kernel learning	naive bayes classifier	robotic	symbolic error analysis
bayesian belief network	deep neural networks	hebbian learning	knowledge representation and reasoning	naive bayes classifiers	robotics	symbolic reasoning
bayesian belief networks	dictionary learning	hidden markov model	latent dirichlet allocation	natural gradient	robots	systems and control theory
bayesian learning	differential evolution	hidden markov models	latent semantic analyses	natural language generation	rough set	temporal difference learning
bayesian network	differential evolution algorithm	hierarchical clustering	latent semantic analysis	natural language processing	rule based learning	text mining
bayesian networks	dimensionality reduction	high dimensional data	latent variable	natural language understanding	rule learning	text to speech
biped robot	dynamic time warping	high dimensional feature	latent variables	natural languages	self organising map	topic model
biped robots	emotion recognition	high dimensional input	layered control systems	nearest neighbour algorithm	semantic web	trajectory planning
blind signal separation	ensemble learning	high dimensional space	learning automata	nearest neighbour algorithms	semi supervised learning	trajectory tracking
bootstrap aggregation	evolutionary algorithm	human action recognition	legged robot	network	sensor data fusion	transfer learning
brain computer interface	evolutionary algorithms	human activity recognition	legged robots	neural network	sensor data fusions	unmanned aerial vehicle
brain computer interfaces	evolutionary computation	human robot interaction	link prediction	neural networks	sensor fusion	unmanned aerial vehicles
chatbot	extreme learning machine	human robot interactions	link predictions	neural turing	sensor fusions	unsupervised learning
chatbots	extreme learning machines	humanoid robot	logic theorist	neuromorphic computing	sensor network	variational inference
classification tree	face recognition	humanoid robotics	logitboost	non negative matrix factorisation	sensor networks	vector machine
classification trees	facial expression recognition	humanoid robots	long short term memory	object detection	sentiment analyses	vector machines
cluster analyses	factorisation machine	image alignment	long short term memory lstm	object recognition	sentiment analysis	virtual assistant
cluster analysis	factorisation machines	image classification	lstm	obstacle avoidance	service robot	virtual assistants
cognitive automation	feature engineering	image grammars	machine intelligence	optimal search	service robots	visual servoing
						xgboost

Table A2: AI concept statistics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Concepts*	**N. of theses ****	**Share**	**Keyword co-occurrence**	**Biology**	**Chemistry**	**Computer science**	**Engineering**	**Geology**	**Mathematics**	**Medicine**	**Physics**
algorithm	4747	33.72%	42.31%	2.20%	0.00%	39.62%	29.66%	2.31%	12.96%	2.49%	10.01%
network	2785	19.78%	30.74%	17.13%	0.01%	27.50%	24.31%	5.57%	3.23%	7.97%	8.19%
robot	878	6.24%	67.40%	2.98%	0.01%	37.38%	48.31%	0.00%	1.39%	2.78%	6.76%
mapping	688	4.89%	35.90%	18.17%	0.03%	22.24%	21.37%	8.43%	6.40%	7.41%	12.06%
decision making	391	2.78%	45.52%	14.32%	0.04%	35.81%	27.37%	2.30%	2.56%	14.58%	2.81%
machine learning	323	2.29%	85.45%	4.95%	0.02%	60.37%	16.10%	1.24%	9.29%	3.41%	3.41%
sensor network	298	2.12%	73.88%	0.41%	0.00%	55.10%	39.59%	0.41%	1.63%	0.41%	2.45%
image processing	250	1.78%	69.20%	2.80%	0.01%	32.80%	32.40%	0.80%	12.00%	3.20%	15.60%
neural network	221	1.57%	70.79%	13.37%	0.07%	24.26%	30.69%	1.49%	5.94%	13.37%	10.89%
computer vision	165	1.17%	84.85%	0.61%	0.00%	56.97%	29.09%	0.61%	7.88%	1.21%	3.64%
genetic algorithm	159	1.13%	57.42%	3.23%	0.02%	21.94%	52.90%	2.58%	2.58%	1.94%	12.26%
data mining	113	0.80%	77.88%	3.54%	0.03%	67.26%	15.04%	4.42%	3.54%	4.42%	1.77%
support vector machine	92	0.65%	87.78%	3.33%	0.04%	36.67%	40.00%	3.33%	3.33%	5.56%	6.67%
random field	88	0.63%	59.49%	0.00%	0.00%	15.19%	54.43%	6.33%	17.72%	2.53%	3.80%
artificial intelligence	87	0.62%	81.61%	2.30%	0.03%	57.47%	27.59%	1.15%	3.45%	5.75%	2.30%
markovian	86	0.61%	45.35%	1.16%	0.01%	17.44%	15.12%	0.00%	50.00%	3.49%	11.63%
multi agent system	81	0.58%	58.33%	0.00%	0.00%	62.50%	31.94%	0.00%	2.78%	1.39%	1.39%
fuzzy logic	77	0.55%	76.62%	1.30%	0.02%	38.96%	51.95%	0.00%	2.60%	2.60%	2.60%
semantic web	76	0.54%	53.95%	0.00%	0.00%	93.42%	3.95%	0.00%	0.00%	2.63%	0.00%
supervised learning	68	0.48%	88.24%	2.94%	0.04%	66.18%	14.71%	2.94%	10.29%	1.47%	1.47%
decision tree	65	0.46%	74.19%	9.68%	0.16%	43.55%	16.13%	6.45%	1.61%	8.06%	11.29%
image segmentation	62	0.44%	79.03%	1.61%	0.03%	33.87%	33.87%	3.23%	12.90%	4.84%	9.68%
pattern recognition	62	0.44%	56.45%	17.74%	0.29%	43.55%	11.29%	4.84%	4.84%	8.06%	4.84%
bayesian network	61	0.43%	78.43%	3.92%	0.08%	47.06%	43.14%	0.00%	5.88%	0.00%	0.00%
artificial neural network	58	0.41%	74.55%	5.45%	0.10%	21.82%	43.64%	7.27%	1.82%	7.27%	12.73%
natural language processing	58	0.41%	67.24%	0.00%	0.00%	86.21%	8.62%	0.00%	0.00%	5.17%	0.00%
feature extraction	54	0.38%	96.30%	0.00%	0.00%	53.70%	29.63%	0.00%	5.56%	7.41%	3.70%
random forest	54	0.38%	89.80%	6.12%	0.12%	36.73%	24.49%	18.37%	12.24%	0.00%	2.04%
deep learning	51	0.36%	98.04%	1.96%	0.04%	52.94%	27.45%	0.00%	7.84%	9.80%	0.00%
graphical model	50	0.36%	80.43%	0.00%	0.00%	56.52%	26.09%	0.00%	15.22%	0.00%	2.17%
evolutionary algorithm	49	0.35%	73.33%	4.44%	0.10%	46.67%	33.33%	0.00%	2.22%	0.00%	8.89%
recommender system	49	0.35%	82.93%	0.00%	0.00%	95.12%	2.44%	0.00%	2.44%	0.00%	0.00%
feature selection	47	0.33%	95.74%	2.13%	0.05%	46.81%	27.66%	2.13%	17.02%	4.26%	0.00%
autonomous vehicle	46	0.33%	83.33%	0.00%	0.00%	50.00%	42.86%	0.00%	4.76%	0.00%	2.38%
reinforcement learning	46	0.33%	91.30%	10.87%	0.24%	56.52%	15.22%	0.00%	2.17%	13.04%	2.17%
humanoid robot	45	0.32%	88.89%	5.56%	0.15%	36.11%	47.22%	0.00%	2.78%	2.78%	5.56%
unsupervised learning	42	0.30%	95.24%	4.76%	0.11%	54.76%	16.67%	0.00%	19.05%	0.00%	4.76%
convolutional neural network	40	0.28%	97.22%	2.78%	0.08%	52.78%	36.11%	0.00%	5.56%	2.78%	0.00%
object detection	40	0.28%	82.50%	0.00%	0.00%	55.00%	30.00%	0.00%	10.00%	0.00%	5.00%
gaussian mixture model	39	0.28%	76.32%	0.00%	0.00%	50.00%	26.32%	0.00%	10.53%	7.89%	5.26%
dimensionality reduction	35	0.25%	71.43%	11.43%	0.33%	40.00%	22.86%	0.00%	17.14%	5.71%	2.86%
human robot interaction	35	0.25%	96.88%	0.00%	0.00%	68.75%	28.13%	0.00%	0.00%	0.00%	3.13%
image classification	35	0.25%	91.43%	0.00%	0.00%	68.57%	11.43%	8.57%	5.71%	0.00%	5.71%
object recognition	35	0.25%	54.29%	17.14%	0.49%	45.71%	11.43%	0.00%	14.29%	11.43%	0.00%
k means	33	0.23%	87.88%	6.06%	0.18%	36.36%	24.24%	3.03%	15.15%	9.09%	6.06%
speech recognition	32	0.23%	62.50%	0.00%	0.00%	78.13%	15.63%	0.00%	3.13%	0.00%	3.13%
trajectory tracking	31	0.22%	83.87%	0.00%	0.00%	29.03%	61.29%	0.00%	0.00%	0.00%	9.68%
hidden markov model	30	0.21%	73.33%	6.67%	0.22%	50.00%	26.67%	0.00%	16.67%	0.00%	0.00%
hidden markov models	30	0.21%	93.33%	0.00%	0.00%	46.67%	30.00%	0.00%	23.33%	0.00%	0.00%
visual servoing	30	0.21%	83.33%	0.00%	0.00%	23.33%	66.67%	0.00%	3.33%	0.00%	6.67%
high dimensional data	28	0.20%	96.43%	0.00%	0.00%	39.29%	10.71%	0.00%	35.71%	7.14%	0.00%
independent component analysis	28	0.20%	75.00%	17.86%	0.64%	10.71%	39.29%	3.57%	10.71%	10.71%	7.14%
community detection	27	0.19%	100.00%	0.00%	0.00%	81.48%	0.00%	0.00%	7.41%	0.00%	11.11%
gaussian process	27	0.19%	37.04%	0.00%	0.00%	14.81%	33.33%	0.00%	51.85%	0.00%	0.00%
unmanned aerial vehicle	27	0.19%	69.23%	0.00%	0.00%	34.62%	46.15%	7.69%	3.85%	0.00%	7.69%
brain computer interface	25	0.18%	55.00%	5.00%	0.25%	45.00%	25.00%	0.00%	0.00%	20.00%	5.00%
industrial robot	25	0.18%	100.00%	0.00%	0.00%	10.53%	78.95%	0.00%	5.26%	0.00%	5.26%

latent variable	25	0.18%	72.73%	4.55%	0.21%	36.36%	13.64%	0.00%	36.36%	9.09%	0.00%
motion planning	25	0.18%	88.00%	4.00%	0.16%	32.00%	48.00%	0.00%	12.00%	0.00%	4.00%
ambient intelligence	23	0.16%	73.91%	0.00%	0.00%	73.91%	21.74%	0.00%	4.35%	0.00%	0.00%
face recognition	23	0.16%	73.91%	0.00%	0.00%	52.17%	30.43%	0.00%	0.00%	17.39%	0.00%
image retrieval	23	0.16%	91.30%	0.00%	0.00%	82.61%	13.04%	0.00%	0.00%	0.00%	4.35%
machine translation	23	0.16%	52.17%	0.00%	0.00%	95.65%	0.00%	0.00%	4.35%	0.00%	0.00%
sparse representation	23	0.16%	86.96%	0.00%	0.00%	26.09%	65.22%	0.00%	8.70%	0.00%	0.00%
cyber physical system	22	0.16%	52.38%	0.00%	0.00%	76.19%	23.81%	0.00%	0.00%	0.00%	0.00%
hierarchical clustering	22	0.16%	59.09%	13.64%	0.62%	31.82%	13.64%	9.09%	9.09%	13.64%	4.55%
obstacle avoidance	20	0.14%	85.00%	0.00%	0.00%	45.00%	45.00%	0.00%	0.00%	5.00%	5.00%
trajectory planning	20	0.14%	90.00%	0.00%	0.00%	25.00%	45.00%	0.00%	20.00%	0.00%	10.00%
collision avoidance	19	0.13%	68.42%	0.00%	0.00%	36.84%	57.89%	0.00%	0.00%	0.00%	5.26%
mapreduce	19	0.13%	78.95%	0.00%	0.00%	78.95%	21.05%	0.00%	0.00%	0.00%	0.00%
cluster analysis	18	0.13%	55.56%	22.22%	1.23%	5.56%	22.22%	11.11%	11.11%	27.78%	0.00%
recurrent neural network	18	0.13%	83.33%	0.00%	0.00%	50.00%	16.67%	0.00%	11.11%	5.56%	16.67%
ant colony	16	0.11%	87.50%	0.00%	0.00%	56.25%	25.00%	0.00%	12.50%	0.00%	6.25%
dictionary learning	16	0.11%	87.50%	0.00%	0.00%	37.50%	31.25%	0.00%	18.75%	12.50%	0.00%
deep neural network	14	0.10%	100.00%	7.14%	0.51%	50.00%	35.71%	0.00%	0.00%	7.14%	0.00%
collaborative filtering	13	0.09%	84.62%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
dynamic time warping	13	0.09%	84.62%	0.00%	0.00%	46.15%	38.46%	7.69%	0.00%	7.69%	0.00%
gesture recognition	13	0.09%	69.23%	0.00%	0.00%	69.23%	23.08%	0.00%	7.69%	0.00%	0.00%
differential evolution	12	0.09%	66.67%	16.67%	1.39%	16.67%	41.67%	8.33%	8.33%	0.00%	8.33%
emotion recognition	12	0.09%	75.00%	0.00%	0.00%	50.00%	25.00%	0.00%	0.00%	16.67%	8.33%
fuzzy set	12	0.09%	70.00%	0.00%	0.00%	70.00%	20.00%	0.00%	10.00%	0.00%	0.00%
kernel learning	12	0.09%	83.33%	8.33%	0.69%	58.33%	25.00%	0.00%	8.33%	0.00%	0.00%
memetic algorithm	12	0.09%	75.00%	8.33%	0.69%	91.67%	0.00%	0.00%	0.00%	0.00%	0.00%
transfer learning	12	0.09%	91.67%	0.00%	0.00%	50.00%	41.67%	0.00%	8.33%	0.00%	0.00%
semi supervised learning	11	0.08%	100.00%	9.09%	0.83%	45.45%	27.27%	0.00%	9.09%	0.00%	9.09%
stochastic gradient	11	0.08%	100.00%	9.09%	0.83%	36.36%	18.18%	0.00%	36.36%	0.00%	0.00%
biped robot	10	0.07%	100.00%	0.00%	0.00%	33.33%	66.67%	0.00%	0.00%	0.00%	0.00%
fuzzy c	10	0.07%	100.00%	0.00%	0.00%	30.00%	40.00%	10.00%	0.00%	0.00%	20.00%
text mining	10	0.07%	70.00%	0.00%	0.00%	80.00%	10.00%	0.00%	0.00%	10.00%	0.00%
action recognition	9	0.06%	88.89%	0.00%	0.00%	55.56%	22.22%	0.00%	22.22%	0.00%	0.00%
adaboost	9	0.06%	100.00%	0.00%	0.00%	55.56%	44.44%	0.00%	0.00%	0.00%	0.00%
autonomic computing	9	0.06%	66.67%	0.00%	0.00%	88.89%	11.11%	0.00%	0.00%	0.00%	0.00%
ensemble learning	9	0.06%	100.00%	0.00%	0.00%	66.67%	11.11%	11.11%	0.00%	11.11%	0.00%
image matching	9	0.06%	88.89%	0.00%	0.00%	44.44%	22.22%	0.00%	33.33%	0.00%	0.00%
image recognition	9	0.06%	100.00%	11.11%	1.23%	55.56%	22.22%	0.00%	11.11%	0.00%	0.00%
pattern analysis	9	0.06%	22.22%	44.44%	4.94%	0.00%	0.00%	11.11%	0.00%	11.11%	33.33%
regression tree	9	0.06%	55.56%	22.22%	2.47%	0.00%	0.00%	33.33%	44.44%	0.00%	0.00%
service robot	9	0.06%	100.00%	0.00%	0.00%	77.78%	11.11%	0.00%	0.00%	0.00%	11.11%
spectral clustering	9	0.06%	77.78%	0.00%	0.00%	44.44%	22.22%	0.00%	22.22%	0.00%	0.00%
support vector regression	9	0.06%	77.78%	0.00%	0.00%	33.33%	55.56%	0.00%	0.00%	0.00%	11.11%
activity recognition	8	0.06%	87.50%	0.00%	0.00%	50.00%	25.00%	0.00%	0.00%	0.00%	25.00%
deep convolutional neural	8	0.06%	100.00%	0.00%	0.00%	62.50%	25.00%	0.00%	12.50%	0.00%	0.00%
link prediction	8	0.06%	100.00%	12.50%	1.56%	75.00%	12.50%	0.00%	0.00%	0.00%	0.00%
pedestrian detection	8	0.06%	87.50%	0.00%	0.00%	62.50%	25.00%	0.00%	0.00%	12.50%	0.00%
genetic programming	7	0.05%	85.71%	0.00%	0.00%	57.14%	28.57%	0.00%	0.00%	0.00%	14.29%
sensor fusion	7	0.05%	85.71%	0.00%	0.00%	57.14%	42.86%	0.00%	0.00%	0.00%	0.00%
association rule	6	0.04%	83.33%	0.00%	0.00%	83.33%	0.00%	16.67%	0.00%	0.00%	0.00%
collaborative systems	6	0.04%	66.67%	0.00%	0.00%	83.33%	0.00%	0.00%	0.00%	0.00%	16.67%
generative adversarial network	6	0.04%	100.00%	16.67%	2.78%	0.00%	66.67%	0.00%	16.67%	0.00%	0.00%
humanoid robotics	6	0.04%	100.00%	16.67%	2.78%	16.67%	66.67%	0.00%	0.00%	0.00%	0.00%
intelligent agent	6	0.04%	83.33%	0.00%	0.00%	83.33%	16.67%	0.00%	0.00%	0.00%	0.00%
multi label classification	6	0.04%	66.67%	0.00%	0.00%	83.33%	16.67%	0.00%	0.00%	0.00%	0.00%
multi sensor fusion	6	0.04%	50.00%	0.00%	0.00%	66.67%	33.33%	0.00%	0.00%	0.00%	0.00%
sentiment analysis	6	0.04%	50.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
autoencoder	5	0.04%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
backpropagation	5	0.04%	80.00%	0.00%	0.00%	40.00%	40.00%	0.00%	0.00%	0.00%	20.00%
high dimensional space	5	0.04%	100.00%	0.00%	0.00%	60.00%	0.00%	0.00%	20.00%	20.00%	0.00%
image alignment	5	0.04%	80.00%	0.00%	0.00%	20.00%	20.00%	0.00%	20.00%	40.00%	0.00%
lstm	5	0.04%	80.00%	20.00%	4.00%	60.00%	20.00%	0.00%	0.00%	0.00%	0.00%

machine vision	5	0.04%	100.00%	0.00%	0.00%	40.00%	40.00%	0.00%	0.00%	0.00%	20.00%
natural language generation	5	0.04%	60.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
text to speech	5	0.04%	80.00%	0.00%	0.00%	60.00%	20.00%	0.00%	0.00%	0.00%	20.00%
bayesian learning	4	0.03%	75.00%	25.00%	6.25%	25.00%	50.00%	0.00%	0.00%	0.00%	0.00%
facial expression recognition	4	0.03%	100.00%	0.00%	0.00%	75.00%	25.00%	0.00%	0.00%	0.00%	0.00%
long short term memory	4	0.03%	100.00%	0.00%	0.00%	75.00%	25.00%	0.00%	0.00%	0.00%	0.00%
meta learning	4	0.03%	50.00%	0.00%	0.00%	75.00%	0.00%	0.00%	0.00%	25.00%	0.00%
multi layer perceptron	4	0.03%	100.00%	0.00%	0.00%	50.00%	25.00%	0.00%	0.00%	25.00%	0.00%
natural languages	4	0.03%	75.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
neuromorphic computing	4	0.03%	75.00%	0.00%	0.00%	25.00%	50.00%	0.00%	0.00%	0.00%	25.00%
sensor data fusion	4	0.03%	100.00%	0.00%	0.00%	50.00%	25.00%	0.00%	25.00%	0.00%	0.00%
simultaneous localisation	4	0.03%	100.00%	0.00%	0.00%	25.00%	75.00%	0.00%	0.00%	0.00%	0.00%
topic model	4	0.03%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
classification tree	3	0.02%	66.67%	0.00%	0.00%	33.33%	0.00%	33.33%	0.00%	33.33%	0.00%
cognitive computing	3	0.02%	100.00%	0.00%	0.00%	0.00%	33.33%	0.00%	0.00%	33.33%	33.33%
high dimensional feature	3	0.02%	100.00%	0.00%	0.00%	66.67%	33.33%	0.00%	0.00%	0.00%	0.00%
knowledge representation and reasoning	3	0.02%	66.67%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
multi task learning	3	0.02%	66.67%	0.00%	0.00%	66.67%	0.00%	0.00%	33.33%	0.00%	0.00%
naive bayes classifier	3	0.02%	100.00%	0.00%	0.00%	33.33%	66.67%	0.00%	0.00%	0.00%	0.00%
natural gradient	3	0.02%	100.00%	0.00%	0.00%	66.67%	33.33%	0.00%	0.00%	0.00%	0.00%
natural language understanding	3	0.02%	33.33%	0.00%	0.00%	66.67%	33.33%	0.00%	0.00%	0.00%	0.00%
variational inference	3	0.02%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
virtual assistant	3	0.02%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
deep belief network	2	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
fuzzy number	2	0.01%	100.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%
human action recognition	2	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
latent dirichlet allocation	2	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
latent semantic analysis	2	0.01%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	50.00%	50.00%
multi objective optimisation	2	0.01%	100.00%	0.00%	0.00%	50.00%	50.00%	0.00%	0.00%	0.00%	0.00%
optimal search	2	0.01%	50.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%
particle swarm optimisation	2	0.01%	100.00%	0.00%	0.00%	50.00%	0.00%	0.00%	0.00%	0.00%	50.00%
q learning	2	0.01%	100.00%	0.00%	0.00%	50.00%	50.00%	0.00%	0.00%	0.00%	0.00%
statistical relational learning	2	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
swarm intelligence	2	0.01%	100.00%	0.00%	0.00%	50.00%	0.00%	0.00%	0.00%	0.00%	50.00%
symbolic reasoning	2	0.01%	50.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
vector machine	2	0.01%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%
adversarial network	1	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
cognitive automation	1	0.01%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%
cognitive modelling	1	0.01%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%
evolutionary computation	1	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
extreme learning machine	1	0.01%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%
gradient boosting	1	0.01%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%
hebbian learning	1	0.01%	100.00%	100.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
high dimensional input	1	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
human activity recognition	1	0.01%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%
machine intelligence	1	0.01%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%
relational learning	1	0.01%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%
rough set	1	0.01%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
similarity learning	1	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
social robot	1	0.01%	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	0.00%

NOTE: *Concepts are sets of similar keywords. For instance, the terms “robotics”, “robot”, “robots”, and “robotic” refer to the same concept “robot”. **The keyword may co-occur with other keywords in the same document. Therefore, the same theses might be counted for different keywords introducing double counting for the variable “*N. of theses*”.

Appendix B: Matching OpenAlex authors with PhD supervisors

To collect bibliometric information about the supervisors, we matched their names with the OpenAlex authors' names. To do so, we converted names to lowercase, removing punctuation, and handling special characters in both supervisors and OpenAlex datasets. After cleaning, we matched the supervisors' and OpenAlex authors' names. We expected to find false positives among the matches, i.e., supervisors with common names that are matched to OpenAlex authors due to homonymy. To exclude matches likely to be false positives, we applied a filter dropping supervisor-author profiles if the author's first publication was less than five years before the matched supervisor's first PhD thesis supervision. The rationale for applying this filter relies on the fact that an HDR (*Habilitation à Diriger des Recherches*) is required to supervise PhD students in France. Therefore, we excluded supervisors matched with OpenAlex authors' profiles whose first publication occurred too close to the first matched supervisor's PhD supervision. In other words, we considered the supervisor-author matches valid if a minimum reasonable period elapsed between the start of the research activity and the HDR. After filtering the false positive supervisor-author matches, we obtained 15,906 supervisors with a unique OpenAlex author profile, 8,069 supervisors with multiple OpenAlex author profiles, and 1,381 supervisors not matched with any OpenAlex author profile. To define our study sample of 17,287 supervisors, we adopted a conservative approach. We included the 15,906 supervisors with a unique OpenAlex author profile and the 1,381 supervisors not matched with any OpenAlex author profile, i.e., the supervisors with no publications. We dropped the 8,069 supervisors for which we observed multiple OpenAlex author profiles because we were unable to attribute to the supervisor with certainty one of the multiple OpenAlex authors (all having the same name). In other words, the risk encountered was to attribute to the supervisor the publications of different homonym authors.

Appendix C: Year-by-year analysis of the student-to-supervisor knowledge transfer

In this appendix, we break down the student-to-supervisor knowledge transfer analysis year-by-year, starting from the year of the student's enrollment (*t-3*) to five years after the student's defense (*t+5*). Table C1 presents the results. Between *t-3* and *t-1*, the difference between pairs with and without AI students in the probability of the supervisor engaging in AI research ranges from 5 to 8 percentage points, reaching its maximum of 12 percentage points only in *t*. Finally, we observe a slight decline to 11 and 10 percentage points in *t+1* and *t+2* and a further decline to 8 and 7 percentage points in *t+3* until *t+5*. This analysis allows us to assess the timing of the supervisor's AI knowledge acquisition, shedding light on the possible mechanism leading to

knowledge acquisition. Indeed, we argue that knowledge transfer during the PhD period is more likely to be the result of hiring a PhD student with ready-to-use AI knowledge. On the contrary, the supervisor's use of AI knowledge after graduation is more likely to result from a time-consuming co-learning process with the student during the PhD period.

Table C1: Supervisor's year-by-year probability of having an AI publication

	Probability of observing an AI Supervisor								
	t-3	*t-2*	*t-1*	*t*	*t+1*	*t+2*	*t+3*	*t+4*	*t+5*
AI student = 1 (1,090 student-supervisor pairs)	0.08	0.09	0.13	0.17	0.17	0.16	0.14	0.16	0.16
AI student = 0 (1,090 student-supervisor pairs)	0.03	0.04	0.05	0.05	0.06	0.06	0.06	0.09	0.08
Difference	0.05	0.05	0.08	0.12	0.11	0.10	0.08	0.07	0.08
P-value	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

Appendix D: Use of specific keywords in tracing the bidirectional learning

This appendix investigates the AI knowledge transfer by checking whether supervisors and students use the same AI keyword when learning from each other. When we consider the supervisor-to-student knowledge transfer, Table D1 shows that 33.26% of students use in their theses the exact AI keyword used by their supervisors in their publication before the students' enrollment. When considering the student-to-supervisor knowledge transfer, Table D1 shows that 33.89% of supervisors publish articles after their students' defense year using the same AI keywords found in their students' theses.

Table D1: Learning traced using specific AI keywords

	Share
Supervisor-to-student knowledge transfer: Share of students using supervisors' keywords	33.26%
Student-to-supervisor knowledge transfer: Share of supervisors using students' keywords	33.89%

Appendix E: Considering alternative lists of AI keywords

In this Appendix, we consider alternative ways to identify AI documents. Since there is no consensus in the literature on a list of AI keywords that should be used to identify AI documents, we rerun our main empirical analyses identifying AI documents using separately two AI keyword lists: one keyword list by Cockburn et al. (2018) and one keyword list by OECD (Baruffaldi et al., 2020; Calvino et al., 2023). Tables E1 and E2 present the results of assessing the bidirectional AI knowledge transfer when identifying AI documents using the Cockburn et al. (2018) list, while Tables E3 and E4 show the results when using the OECD list. The results obtained in the four tables are qualitatively the same as those reported in the paper's main text, confirming that our analyses are not sensitive to the list of keywords used to identify AI documents.

Table E1: Student's probability of doing a thesis in AI, using the Cockburn et al. (2018) list

	Pr(*AI student = 1*)
AI Supervisor before = 1 (3,396 student-supervisor pairs)	0.13
AI Supervisor before = 0 (3,396 student-supervisor pairs)	0.03
Difference	0.10
P-value	<0.01

Table E2: Supervisor's probability of having an AI publication in the three years after the student's defense year, using the Cockburn et al. (2018) list

	Pr(*AI supervisor after = 1*)
AI student = 1 (1,090 student-supervisor pairs)	0.29
AI student = 0 (1,090 student-supervisor pairs)	0.14
Difference	0.15
P-value	<0.01

Table E3: Student's probability of doing a thesis in AI, using the OECD list

	Pr(*AI student = 1*)
AI Supervisor before = 1 (6,775 student-supervisor pairs)	0.20
AI Supervisor before = 0 (6,775 student-supervisor pairs)	0.08
Difference	0.12
P-value	<0.01

Table E4: Supervisor's probability of having an AI publication in the three years after the student's defense year, using the OECD list

	Pr(*AI supervisor after = 1*)
AI student = 1 (1,090 student-supervisor pairs)	0.29
AI student = 0 (1,090 student-supervisor pairs)	0.14
Difference	0.15
P-value	<0.01

Appendix F: Removing the supervisor's publications co-authored with the student after the defense year

In this appendix, we assess student-to-supervisor knowledge transfer as in the main analysis but removing the supervisors' publications co-authored with the students after the students' defense year. Table F1 shows results similar to those reported in Table 8.

Table F1: Supervisor's probability of having an AI publication in the three years after the student's defense year, excluding the publications co-authored with the student

	Pr(*AI supervisor after = 1*)
AI student = 1 (1,090 student-supervisor pairs)	0.25
AI student = 0 (1,090 student-supervisor pairs)	0.14
Difference	0.11
P-value	<0.01

Appendix G: Bidirectional learning by research field

In this appendix, we explore field heterogeneity. Table G1 shows that the AI knowledge transfer from supervisor to student with the largest extent occurs in *Computer Science*, while Table G2 shows that the transfer from student to supervisor in *Computer Science* is quite limited if compared to other fields. We find the opposite empirical evidence in *Biology*, in which we do not observe AI knowledge transfer from supervisor to student (Table G1), while we observe a positive and statistically significant AI knowledge transfer from student to supervisor (Table G2).

Table G1: Supervisor-to-student knowledge transfer by research field.

	Pr(AI Student = 1)							
	Biology	*Chemistry*	*Engineering*	*Geology*	*Mathematics*	*Medicine*	*Physics*	*Computer Science*
AI supervisor before = 1 (A)	0.04	0.03	0.21	0.05	0.25	0.08	0.09	0.38
AI supervisor before = 0 (B)	0.03	0.06	0.04	0.02	0.04	0.03	0.03	0.06
Difference	0.01	-0.03	0.17	0.03	0.21	0.05	0.06	0.32
P-value	0.72	0.06	<0.01	0.08	<0.01	<0.01	<0.01	<0.01
N. of pairs (A)	1,014	273	1,395	301	389	606	820	447
N. of pairs (B)*	1,014	273	1,395	301	389	606	820	447

NOTE: *Not all the student-supervisor pairs with a non-AI supervisor before the student's enrollment (B) are in the same research field as the student-supervisor pairs with an AI supervisor (A). Indeed, the matching is based on a propensity score similarity, allowing for matching controls in different fields. This latter event is quite rare, being the field one of the variables used to calculate the propensity score. However, it must be noted that the number of observations of sample (B) is equivalent to the one of sample (A), because it indicates the number of controls matched with sample (A), despite their research field.

Table G2: Student-to-supervisor knowledge transfer by research field.

	Pr(AI Supervisor after = 1)							
	Biology	*Chemistry*	*Engineering*	*Geology*	*Mathematics*	*Medicine*	*Physics*	*Computer Science*
AI student = 1 (A)	0.20	0.14	0.37	0.19	0.23	0.33	0.21	0.29
AI student = 0 (B)	0.05	0.07	0.15	0.09	0.14	0.14	0.18	0.15
Difference	0.15	0.07	0.22	0.10	0.09	0.19	0.03	0.14
P-value	0.01	0.56	<0.01	0.29	0.13	0.03	0.70	<0.01
N. of pairs (A)	75	14	307	32	91	49	82	440
N. of pairs (B)*	75	14	307	32	91	49	82	440

NOTE: *Not all the student-supervisor pairs with a non-AI student (B) are in the same research field as the student-supervisor pairs with an AI student (A). Indeed, the matching is based on a propensity score similarity, allowing for matching controls in different fields. This latter event is quite rare, being the field one of the variables used to calculate the propensity score. However, it must be noted that the number of observations of sample (B) is equivalent to the one of sample (A), because it indicates the number of controls matched with sample (A), despite their research field.

Appendix H: Bidirectional learning by cohort of students

In Appendix H, we conduct our analyses by student's cohort of graduation. Table H1 shows that the supervisor-to-student knowledge transfer ranges from 0.09 to 0.12 percentage points, with no trend over time. Similarly, Table H2 shows that the student-to-supervisor knowledge transfer ranges from 0.12 to 0.20 percentage points, with no particular trend over time.

Table H1: Student's probability of doing a thesis in AI, by cohort

	Pr(AI Student = 1)								
	2010	*2011*	2012	2013	2014	*2015*	*2016*	*2017*	*2018*
AI supervisor before = 1	0.13	0.15	0.16	0.13	0.14	0.14	0.14	0.13	0.14
AI supervisor before = 0	0.02	0.04	0.04	0.04	0.05	0.03	0.04	0.03	0.05
Difference	0.11	0.11	0.12	0.09	0.09	0.11	0.10	0.10	0.09
P-value	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
N. of pairs	410	472	574	594	612	652	672	689	570

Table H2: Supervisor's probability of having an AI publication in the three years after the student's defense year, by cohort

	Pr(AI Supervisor = 1)								
	2010	*2011*	2012	2013	2014	*2015*	*2016*	*2017*	*2018*
AI student = 1	0.30	0.28	0.33	0.31	0.24	0.28	0.28	0.32	0.30
AI student = 0	0.15	0.16	0.14	0.17	0.10	0.15	0.16	0.12	0.12
Difference	0.15	0.12	0.19	0.14	0.14	0.13	0.12	0.20	0.18
P-value	0.01	0.04	<0.01	0.01	0.01	0.01	0.02	<0.01	<0.01
N. of pairs	114	101	105	127	106	137	133	139	128

Appendix I: Research on AI versus the use of AI as a tool

In this appendix, we aim to identify theses having AI as the main research subject, from theses using AI as a tool. To do so, for a subset of theses in AI for which the full text is available (1,210 theses out of 3,012 theses), we calculate the intensity of the use of AI keywords in the text. We assume that when there is a high intensity of use of AI keywords throughout the whole text, the thesis is likely to be a thesis having AI as its leading research subject. In contrast, when the intensity of AI keywords is low, the thesis is expected to use AI as a tool. Table I1 presents our findings and reveals that the intensity of use of AI keywords in theses in *Computer Science* is the highest, i.e., 33.18 AI keywords every 10,000 words. The intensity of AI keywords is lower in other research fields such as *Biology*, i.e., 12.79 AI keywords every 10,000 words. We interpret this finding as empirical evidence that AI tends to be a research subject in fields such as *Computer Science*. In contrast, in other fields such as *Biology*, researchers use AI as a tool.

Table I1: Intensity of AI keywords in the text by research field.

	Intensity of AI keywords	N. of theses
Biology	12.79 [AI keywords/10,000 words]	58
Chemistry	8.95 [AI keywords/10,000 words]	7
Computer Science	33.18 [AI keywords/10,000 words]	560
Engineering	27.90 [AI keywords/10,000 words]	313
Geology	14.21 [AI keywords/10,000 words]	18
Mathematics	16.45 [AI keywords/10,000 words]	140
Medicine	18.10 [AI keywords/10,000 words]	52
Physics	16.74 [AI keywords/10,000 words]	62

Appendix J: Validating the use of titles and abstracts to identify AI theses

We use the full text[18] of 12,382 theses to validate the use of titles and abstracts in identifying AI documents using keywords. Table J1 shows that, on average, theses that we classify as AI theses using titles and abstracts, show a high intensity of AI keywords, i.e., 26.99 AI keywords every 10,000 words. On the contrary, theses classified as non-AI theses using titles and abstracts, show a significantly lower intensity in the use of AI keywords, i.e., 1.18 AI keywords every 10,000 words. Those figures show that relying on titles and abstracts assures a good identification of AI documents. In other words, the AI content of the full text is coherent with the AI content found in the title and abstract.

Table J1: Intensity of AI keywords in the text.

Thesis classification according to titles and abstracts	Intensity of AI keywords in the full text	N. of theses
AI theses = 1	26.99 [AI keywords/10,000 words]	1,210
AI theses = 0	1.18 [AI keywords/10,000 words]	11,172

[18] ABES provided us with the full texts of a subset of theses that we used to conduct this robustness check.

Appendix K: Exact matching on the graduation location

In this appendix, we perform an exact matching on the graduation location. Tables K1 and K2 show that the results are consistent with those reported in Tables 7 and 8 in the main text.

Table K1: Supervisor-to-student knowledge transfer. Student's probability of doing a thesis in AI, with exact matching on the graduation location.

	Pr(*AI student = 1*)
AI Supervisor before = 1 (10,333 student-supervisor pairs)	0.19
AI Supervisor before = 0 (10,333 student-supervisor pairs)	0.05
Difference	0.14
P-value	<0.01

Table K2: Student-to-supervisor knowledge transfer. Supervisor's probability of having an AI publication in the three years after the student's defense year, with exact matching on the graduation location.

	Pr(*AI supervisor = 1*)
AI student = 1 (1,090 student-supervisor pairs)	0.29
AI student = 0 (1,090 student-supervisor pairs)	0.14
Difference	0.15
P-value	<0.01

Appendix L: Matching without caliper

In this appendix, we assess *Supervisor-to-student* knowledge transfer without applying the "caliper" restriction when identifying a control sample of similar pairs without an AI supervisor before the student's enrollment. Table L1 shows results similar to those of Table 7 in the main text.

Table L1: Supervisor-to-student knowledge transfer. Student's probability of doing a thesis in AI, excluding the "caliper" restriction when identifying the control sample.

	Pr(*AI student*)
AI Supervisor before (10,333 student-supervisor pairs)	0.19
Non-AI Supervisor before (10,333 student-supervisor pairs)	0.05
Difference	0.14
P-value	<0.01

Appendix M: Supervisor's co-authors with AI knowledge

In this appendix, we modify Equations 1 and 2 including a proxy for the presence of co-authors with AI knowledge in the supervisor's network. The rationale for including a proxy for AI knowledge in the supervisor's network is that the bidirectional flow of AI knowledge observed in our main analyses might also flow from the supervisor's co-authors. We proxy the presence of co-authors with AI knowledge in the supervisor's network by calculating the variable *Share of AI co-authors* as the ratio between the number of supervisor's co-authors with AI knowledge and the total number of supervisor's co-authors before the enrollment of the student, i.e., before *t-3*, where *t* is the defense year. We define a co-author with AI knowledge by looking at the presence in the co-author's articles of AI keywords. If the co-author has at least one article with AI keywords, we define her as an AI co-author. We find that it is common for the supervisors to have an AI co-author. Indeed, 86.97% of the 40,852 PhD students in our sample have a supervisor with at least one AI co-author. The variable *Share of AI co-authors*, included in Equations 1 and 2, measures the intensity of the presence of AI coauthors in the supervisors' network. The average *Share of AI co-authors* equals 0.17, meaning that, on average, 17% of the coauthors in the supervisors' networks before the student enrollment are AI co-authors.

Tables M1 and M2 show the results of our analysis when we include the variable *Share of AI co-authors* in Equations 1 and 2. The results are consistent with the ones observed in Tables 7 and 8 in the main text.

Table M1: Supervisor-to-student knowledge transfer. Student's probability of doing a thesis in AI, including *Share of AI co-authors* in Equation 1.

	Pr(*AI student = 1*)
AI Supervisor before = 1 (10,333 student-supervisor pairs)	0.19
AI Supervisor before = 0 (10,333 student-supervisor pairs)	0.05
Difference	0.14
P-value	<0.01

Table M2: Student-to-supervisor knowledge transfer. Supervisor's probability of having an AI publication in the three years after the student's defense year, including *Share of AI co-authors* in Equation 2.

	Pr(*AI supervisor = 1*)
AI student = 1 (1,090 student-supervisor pairs)	0.29
AI student = 0 (1,090 student-supervisor pairs)	0.13
Difference	0.16
P-value	<0.01

Appendix N: Student's AI knowledge before enrollment

In this appendix, we estimate an extension of Equation 2 in which we include a proxy for the student's AI knowledge before their enrollment in the PhD program. Specifically, we define *Student AI knowledge before enrollment* as a dummy variable that equals one if the student has published at least one article with AI content before enrollment.

Table N1 reports our results. We consider a study sample of 30,519 student-supervisor pairs whose supervisors lack AI knowledge (*AI supervisor before* = 0), as in Table 5 in the main text. As expected, the results show a positive and statistically significant association between the variable *Student AI knowledge before enrollment* and the probability of observing an *AI student*. This result shows that students with AI knowledge before enrolling in the PhD program leverage this knowledge to write their PhD thesis.

Table N2 reports the estimate of the supervisor's probability of having an AI publication in the three years after the student's defense, comparing similar student-supervisor pairs that differ only by the presence of a student writing a thesis in AI or not. The similarity between student-supervisor pairs is assessed considering the observable characteristics listed in Equation 2 and adding the variable *Student AI knowledge before enrollment* among the regressors. The results are consistent with those observed in Table 8 in the main text.

Table N1: *Student-to-supervisor* knowledge transfer. Regression predicting the probability of observing a student-supervisor pair in which the student is writing a thesis in AI (*AI student*). We include *Student AI knowledge before enrollment* as a proxy for the student's AI knowledge before enrollment in the PhD program. Logit estimates.

Dependent variable: AI student	Estimate	Std.error	t-statistic	P-value
(Intercept)	-1.183	0.148	-7.996	0.000
Student characteristics				
Student AI knowledge before enrollment	0.910	0.270	3.375	0.001
Female	-0.035	0.075	-0.474	0.635
Missing gender	0.095	0.108	0.873	0.382
Team size	0.046	0.018	2.509	0.012
Supervisor characteristics				
Female	0.040	0.073	0.548	0.584
Missing gender	0.042	0.146	0.290	0.771
N. publications	-0.012	0.002	-6.103	0.000
Seniority	-0.017	0.004	-3.769	0.000
Grant	0.096	0.095	1.014	0.310
Experience	0.027	0.016	1.688	0.091
Interdisciplinarity	-0.002	0.011	-0.133	0.894
Filed				
Biology	-2.632	0.139	-18.995	0.000
Chemistry	-3.664	0.279	-13.151	0.000
Engineering	-1.283	0.084	-15.231	0.000
Geology	-2.146	0.191	-11.265	0.000
Mathematics	-1.281	0.124	-10.365	0.000
Medicine	-2.324	0.164	-14.139	0.000
Physics	-2.120	0.130	-16.292	0.000
Computer science	ref	ref	ref	ref
Location				
Bordeaux	-0.157	0.208	-0.757	0.449
Lille	0.219	0.175	1.251	0.211
Lyon	-0.115	0.145	-0.790	0.429
Marseille	0.135	0.203	0.664	0.507
Montpellier	0.111	0.188	0.590	0.555
Nantes	-0.096	0.219	-0.441	0.660
Others	0.037	0.083	0.449	0.653
Rennes	-0.519	0.242	-2.144	0.032
Strasbourg	-0.046	0.220	-0.207	0.836
Toulouse	0.058	0.130	0.445	0.656
Paris	ref	ref	ref	ref
Defense year				
2010	ref	ref	ref	ref
2011	-0.199	0.145	-1.373	0.170
2012	-0.220	0.144	-1.526	0.127
2013	-0.006	0.139	-0.046	0.963
2014	-0.243	0.145	-1.676	0.094
2015	-0.007	0.139	-0.050	0.960
2016	0.035	0.142	0.243	0.808
2017	0.135	0.144	0.938	0.348
2018	0.241	0.149	1.617	0.106
Pseudo R2	0.171			
Student-supervisor pairs	30,519			

Table N2: Student-to-supervisor knowledge transfer. Supervisor's probability of having an AI publication in the three years after the student's defense year, including *Student AI knowledge before enrollment* in Equation 2.

	Pr(*AI supervisor = 1*)
AI student = 1 (1,090 student-supervisor pairs)	0.29
AI student = 0 (1,090 student-supervisor pairs)	0.14
Difference	0.15
P-value	<0.01