AI Integration in Portuguese Higher Education: Evidence from Curricular Changes Before and After ChatGPT

Ana Ferreira

Dissertation written under the supervision of professor Miguel Godinho de Matos.

Dissertation submitted in partial fulfilment of requirements for the MSc in Business Analytics, at the Universidade Católica Portuguesa,

January 2026

AI Integration in Portuguese Higher Education: Evidence from Curricular Changes Before and After ChatGPT

Ana Ferreira

January 2026

Abstract

The integration of artificial intelligence into labor markets worldwide is creating demand for new competencies among workers. Universities face pressure to respond by preparing students with relevant AI skills, and ChatGPT's release in November 2022 intensified this discussion. We investigate this response by analyzing whether and how Portuguese universities have integrated AI content into their curricula. We built a dataset combining data from three official sources: the higher education registry, accreditation agency records, and official government publications. The resulting data includes detailed curricular plans for over 3,000 distinct programs accredited between 2011 and 2025. We classified curricular units by AI content and computed metrics of AI integration for each program in each year. This allowed us to compare programs re-accredited before versus after ChatGPT's release. Our results show that programs re-accredited post-ChatGPT are more likely to include AI content, though effects remain modest overall. The response is highly uneven across fields, with strong increases in engineering but essentially no change in arts, humanities, and health. Additionally, we observe considerable growth in new AI-specialized programs created since 2020. Overall, the results suggest that Portuguese universities are responding by integrating AI content, though the response remains in its early stages.

Keywords: Artificial Intelligence, ChatGPT, University Curricula, Portugal, AI Education

AI Integration in Portuguese Higher Education: Evidence from Curricular Changes Before and After ChatGPT

Ana Ferreira

January 2026

Resumo

A integração da inteligência artificial nos mercados de trabalho em todo o mundo está a criar procura por novas competências dos trabalhadores. As universidades enfrentam pressão para responder preparando os estudantes com competências relevantes em IA, e o lançamento do ChatGPT em novembro de 2022 intensificou esta discussão. Neste estudo, investigamos esta resposta analisando se e como as universidades portuguesas integraram conteúdo de IA nos seus currículos. Construímos um conjunto de dados combinando informação de três fontes oficiais: os registos do ensino superior, registos da agência de acreditação e publicações oficiais do governo. Os dados resultantes incluem planos curriculares detalhados para mais de 3000 programas distintos acreditados entre 2011 e 2025. Classificámos as unidades curriculares de acordo com o conteúdo de IA e calculámos métricas de integração de IA para cada programa em cada ano. Isto permitiu-nos comparar programas reacreditados antes e depois do lançamento do ChatGPT. Os nossos resultados mostram que os programas reacreditados pós-ChatGPT têm maior probabilidade de incluir conteúdo de IA, embora os efeitos permaneçam modestos no geral. A resposta é altamente desigual entre áreas, com aumentos fortes em engenharia mas essencialmente nenhuma mudança nas artes, humanidades e saúde. Adicionalmente, observamos o crescimento considerável de novos programas especializados em IA criados desde 2020. No geral, os resultados sugerem que as universidades portuguesas estão a responder integrando conteúdo de IA, embora a resposta permaneça numa fase inicial.

Palavras-Chave: Inteligência Artificial, ChatGPT, Planos Curriculares Universitários, Portugal, Educação em IA

Table of Contents

List of Figures

List of Tables

List of Abbreviations

AI	Artificial Intelligence
A3ES	Agência De Avaliação e Acreditação Do Ensino Superior
DGES	Direção Geral do Ensino Superior
DR	Diário da República
ECTS	European Credit Transfer System
LPM	Linear Probability Model
ZINB	Zero-Inflated Negative Binomial
NegBin	Negative Binomial

1 Introduction

Artificial intelligence is fundamentally transforming labor markets worldwide. As AI technologies become increasingly integrated into business operations, the skills required for employment are shifting. Recent survey evidence indicates that employers expect 39% of workers' core skills to change by 2030 with AI and Big Data being at the top of the list of fastest-growing skills World Economic Forum (2025). This transformation places higher education institutions at a critical positioning, they must prepare students for a labor market where AI literacy is becoming essential, yet the nature of required AI competencies remains uncertain.

The release of ChatGPT in November 2022 marked a turning point in AI accessibility. For the first time, powerful AI tools became available to the general public, demonstrating capabilities that were previously confined to specialized technical domains. This democratization of AI has accelerated discussions about AI's role in education and intensified pressure on universities to integrate AI-related content into their curricula. However, while there is growing evidence about AI's impact on labor markets and widespread adoption of AI tools by students and educators, there remains limited empirical research on how higher education institutions are actually responding through curricular reform.

This thesis addresses this gap by investigating how the emergence of widely accessible generative AI technologies has influenced curricular plans across Portuguese higher education. We specifically focus on comparing programs updated or created after ChatGPT's release to those developed before. With this we aim to investigate whether institutions have responded by integrating more AI-related content into their curricula. Our analysis draws on a unique dataset covering Portuguese degree programs from 2011 to 2025, constructed by systematically collecting and integrating data from three official sources: *Direção-Geral do Ensino Superior* (DGES), *Agência de Avaliação e Acreditação do Ensino Superior* (A3ES), and *Diário da República de Portugal* (DR). This dataset provides detailed information about curricular plans for over 3,000 distinct programs across bachelor's, master's, and integrated master's degrees. We classify each curricular unit as containing Direct AI content (courses explicitly focused on artificial intelligence concepts), AI Related content (foundational skills supporting AI development), or neither, using a systematic keyword-based approach grounded in AI education frameworks.

Our analysis reveals three main findings. First, programs re-accredited after ChatGPT's release are 2.7 percentage points more likely to include Direct AI content compared to those re-accredited before, representing a relative increase given that less than 6% of program-year observations include such content. Second, this response varies across program types and fields, bachelor's programs show effects more than double those of master's programs, and while engineering programs show effects exceeding 11 percentage points, arts, humanities, and health programs show no significant response. Third, we observe substantial growth in AI content among newly created programs, with 19 programs explicitly incorporating "Artificial Intelligence" in their titles, all established in 2020 or later. Together, these patterns indicate that Portuguese higher education is in the early stages of AI curricular integration, with clear responses emerging in technical fields but limited diffusion to other disciplines.

This research contributes to the limited but growing empirical literature on how higher education institutions adapt curricula in response to technological disruption. While much attention has focused on AI's potential to transform teaching practices through AI-enabled tools, our focus is on Education for AI, meaning how universities are preparing students to work, understand, and critically engage with AI systems in their future careers.

The paper is structured as follows: Section 2 reviews the relevant literature on AI's impact on labor markets and skills demand, as well as existing research on AI integration in higher education. Section 3 presents our data sources, provides descriptive statistics, and explains our methodology. Section 4 presents our results, which are discussed in Section 5. In Section 6 we acknowledge the limitations of our analysis. Finally, the conclusion is presented in Section 7.

2 Literature Review

Education functions as the primary system through which individuals acquire the skills required by the labor market. As artificial intelligence reshapes what jobs require and what competencies matter across different sectors, it becomes increasingly important to understand how universities are adapting their curricula. However, we cannot look at curricular changes in isolation. They must be understood within the larger picture of how AI is restructuring labor markets, redefining which skills are valued, and redefining classic paths from education to employment. This review, therefore, brings together the literature on AI's impact on labor and skills, the growing demand for AI competencies, and what this means for higher education institutions, both in terms of using AI as a teaching tool and developing AI-related curriculum content.

2.1 Impact of AI on Labor and Skills

The most recent report from the Digital Education Council (2025a) on AI in the Workplace indicates substantial AI adoption among workers, with employers expecting AI's disruptive trajectory to continue. Those trends are confirmed in studies made in the United States and Portugal, where around 40% of working-age adults use AI tools (Bick et al., 2024; Instituto Nacional de Estatística, 2025). However, how this will affect workers is still largely unknown. While AI automates some tasks, it simultaneously increases labor demand for new ones (Acemoglu and Restrepo, 2018). This double-edged aspect is consistent with employer expectations: 72% believe AI adoption will reduce headcount, while 62% expect new roles to emerge (Digital Education Council, 2025a). Early evidence also suggests heterogeneous changes across worker groups and occupations, where early-career workers seem to be the most affected. Those experienced employment declines, while their more experienced peers see stable or growing employment (Brynjolfsson et al., 2025).

The critical constraint to AI adoption is human capital. Skills gaps have become the primary barrier to business transformation, with 63% of employers identifying them as a major obstacle (World Economic Forum, 2025). The mismatch between workers' existing skills and AI-augmented economy requirements thus emerges not simply as a factor slowing adjustment, but as a constraint on realizing AI's productivity promise (Acemoglu and Restrepo, 2018). As a result, there has been an increasing demand for AI-skilled workers. In the US, job postings requiring at least one AI skill more than tripled since 2010 (Galeano et al., 2025), with particularly high growth in IT occupations, followed by considerable increases across architecture, engineering, scientific, and management roles (Alekseeva et al., 2021). Globally, AI and Big Data rank among the fastest-growing skills (World Economic Forum, 2025). Employers anticipate that 39% of workers' core skills will change by 2030, rising to 44% in Portugal (World Economic Forum, 2025). In addition, evidence shows that AI's value is maximized when algorithmic skills are dispersed among workers rather than concentrated among technical specialists (Tambe, 2025). This suggests that AI literacy must become a broadly distributed competency across disciplines, a finding with significant effects for higher education curriculum design.

These labor market transformations expose significant disconnects between higher education institutions and employer expectations. The Digital Education Council (2025a) report about AI in the Workplace stated that over half of employers now expect graduates to have AI proficiency. Yet many believe universities are not adequately preparing students, with 80% reporting that higher education is failing to keep pace with industry change (Digital Education Council, 2025a). The future action calls for a fundamental curricular reorientation: the educational priority shifts from teaching students to build algorithms to also teaching them to guide and evaluate AI outputs, an education that must span all majors, not exclusively technical programs (Tambe, 2025). Institutions that successfully guide in this stage of transition will better prepare graduates for the transforming workforce needs.

2.2 Impact of AI on Higher Education

Given the impact of AI on labor markets and the changing skills landscape, higher education institutions are under pressure to prepare students for an AI-integrated future. While comprehensive research on AI's impact on the education sector remains limited, initial evidence suggests that universities are experiencing rapid, organic adoption of AI tools by students and educators. Yet, the implications and institutional response of this change lack in-depth analysis.

To understand this situation, we distinguish between two domains of AI integration in higher education: AI in Education and Education for AI. AI in Education refers to the use of AI tools to improve teaching and learning processes, improve teaching delivery, customize learning experiences, and support pedagogical practices. In contrast, Education for AI focuses on developing students' understanding of AI technologies, including AI literacy, technical competencies, and awareness of AI's ethics and social implications (Maslej et al., 2025; Schiff, 2022). Both approaches are necessary to fully understand AI integration in higher education, though they fulfill different purposes and require different institutional strategies.

AI in Education

AI tools are widely used across higher education. Recent survey data from the Digital Education Council (2024) reveals that 86% of students already incorporate AI in their studies and 54% use these tools at least weekly. Similarly, 61% of faculty have already integrated AI within their teaching practices (Digital Education Council, 2025b). This is consistent with

Anthropic's most recent education reports, which demonstrated that students and educators are actively using AI tools (Handa et al., 2025; Bent et al., 2025). However, adoption is not uniform, with significant disciplinary variation reported. STEM students, particularly those in computer science, constitute the early adopters, while students in business, health, and humanities show lower engagement (Handa et al., 2025). This uneven pattern suggests AI integration is not spreading evenly throughout universities but rather following disciplinary cultures and perceived relevance. The ways students and educators use AI also differ meaningfully. Students primarily use AI as a learning tool, to search for information (Digital Education Council, 2024), and to create and improve educational content (Handa et al., 2025). In contrast, educators primarily leverage AI for curriculum development, using these tools beyond simple chatbot functionality to create interactive teaching materials that can be immediately deployed in classrooms (Bent et al., 2025; Digital Education Council, 2025b). Notably, many educators recognize that AI tools are changing the way students learn, creating pressure to change their teaching ways accordingly (Bent et al., 2025). This evidence reveals the importance of establishing complete directives for AI integration in teaching scenarios, aligned with the development of AI literacy competencies for students and professors.

Education for AI

While AI in Education treats AI as a means to improve learning results, Education for AI positions AI as the subject of learning, preparing students to work with, develop, and thoughtfully interact with AI systems in their future careers. International recognition of the importance of AI education is widespread. An analysis of 24 national AI policy strategies demonstrated that Education for AI is a priority international conversation topic, with all analyzed policies discussing at least one of three key areas: training AI experts, preparing the workforce for AI, and developing public AI literacy (Schiff, 2022). This policy attention demonstrates rising awareness that populations need preparation both to work with AI and to comprehend its social implications. However, recognition of importance has not consistently translated to action. Despite policy commitments, relatively few countries have explicitly integrated it in their course offerings or developed detailed implementation plans (Maslej et al., 2025). However, there is already some evidence of initial progress. In the United States, there has been a significant growth in AI-specific degree offerings. The number of institutions offering AI bachelor's degrees nearly doubled between 2022 and 2023, and there were significant increases in AI master's programs (Maslej et al., 2025). However, these specialized degrees reach only a small fraction of students, and the larger challenge of integrating AI competencies across all programs continues largely unaddressed. An analysis of top-ranked Asian universities revealed that only 11 of 25 institutions had explicit AI policies, and few mentioned the full-scale curriculum reforms necessary to prepare students for the AI era (Wong et al., 2025). This implementation gap is clear in students' perceptions. A majority (58%) report insufficient AI knowledge and skills, 48% feel inadequately prepared for AI-enabled workplaces, and 80% believe their universities are not fully meeting expectations regarding AI integration (Digital Education Council, 2024). Overall, the literature indicates a considerable disconnect between policy ambitions and institutional capabilities, pointing to the need for empirical research on how universities translate AI education priorities into practice.

2.3 Research Question

For our analysis, we will focus on Education for AI since data on curricular structures is available, while systematic data on actual AI tool usage by Portuguese students and teachers in educational contexts is not. Addressing this requires understanding how higher education institutions incorporate AI-related curricular development. We focus on the Portuguese context given our familiarity with its institutional structure. Accordingly, we define our main research question as follows:

How is the emergence of AI transforming curricular plans in Portuguese higher education institutions?

More specifically, we aim to understand if there were any changes in Portuguese degree programs after the release of ChatGPT in November 2022. We focus on two possible ways through which institutions update curricula: re-accreditation of existing programs (mandatory periodic review processes) and creation of new programs (institutions designing curricula from scratch). In addition, we also investigate if the possible increase in AI integration is happening only in technical fields or is spread across areas. By investigating whether and how institutions are reforming curricula in response to AI, this research contributes to much-needed empirical evidence and contributes to ensure that Portuguese graduates remain competitive in an increasingly AI-driven economy.

3 Data and Empirical Strategy

3.1 Data

To answer our research question, we created a dataset containing detailed information about Portuguese degrees' curricula by collecting and integrating data from three different official sources: *Direção-Geral do Ensino Superior* (DGES), *Agência de Avaliação e Acreditação do Ensino Superior* (A3ES), and *Diário da República de Portugal* (DR).

Our primary data source was DGES, a central service of Portugal's Ministry of Education and Science. DGES is responsible for managing higher education policy, student admissions, institutional registration, and statistical information about Portugal's higher education system. From DGES, we retrieved information about registered study cycles across three higher education degree levels: bachelor's, master's, and integrated master's (bachelor's and master's), including those that are no longer active. For each study cycle registration, we collected information that included degree identification (name, institution, ECTS), registration details (code, date, status), academic structure (duration, scientific area classification), accreditation information, and available curricular plan details. The final dataset contained 10,545 observations, each representing one accreditation process for one program.[1] Overall the dataset covers 4582 distinct degrees, but we could only collect detailed curricular plans for 1888 observations.[2] To overcome this limitation we subsequently enriched this dataset with additional information collected from DR and A3ES.

We used the A3ES source to collected comprehensive accreditation information.[3] A3ES's website contains publicly accessible information about accreditation processes and the status of study programs across Portugal's higher education institutions. For each accreditation process, we extracted institutional and degree information, accreditation outcomes, dates, and process numbers. We matched this data with the dataset we collected from DGES to enrich it with additional attributes. This enrichment proved essential for enhancing the matching process with DR records, allowing for the capture of a significantly larger number of detailed curricular plans

[1]This means each unique program will have in the dataset as many observations as accreditation processes it went through. For example, a program created in 2012 and re-accredited in 2018 and 2024 will have 3 observations

[2]DGES website only had available detailed curricular plans for the most recent accreditation processes

[3]A3ES, an agency established by the Portuguese State with the mission of assessing and accrediting higher education institutions and their study programs, with the objective of guaranteeing the quality of Portuguese higher education

that were missing in our original data.

Finally, we used DR, an official governmental source where legislators publish laws, decrees, court decisions, and government acts. Until March 2024 [4], institutions needed to publish in DR any creation or modification of a study cycle, following its accreditation by A3ES and subsequent registration in DGES. Therefore, from the DR platform, we systematically collected these legal documents through seven distinct search queries: *"Criação de ciclo de estudos"* (study cycle creation), *"Criação da licenciatura"* (bachelor's degree creation), *"Criação do mestrado"* (master's degree creation), *"Alteração ao ciclo de estudos"* (study cycle modification), *"Alteração à licenciatura"* (bachelor's degree modification), *"Alteração ao mestrado"* (master's degree modification) and *"Plano de estudos"* (study plan). This automated search process retrieved more than 11000 documents published between 2007 and 2024. The documents we obtained are PDF files, which we processed using OpenAI's large language model API calls to extract the curricular structure and plan information from different degrees.

The final cleaned dataset, integrating information from all three sources, comprises 8,524 observations representing unique program-accreditation combinations across 4,491 distinct degree programs in Portuguese higher education. As shown in Table 1, detailed curricular plans are available for 5,414 observations (63.5%). The reduction from the initial 10,545 observations to the final sample is mainly due to the elimination of duplicates during data cleaning. In addition, some incomplete or inconsistent observations led to the exclusion of 91 programs.

Table 1: Dataset Composition by Degree Level

			With Curricular Plan	
Degree Level	**Total Obs.**	**Unique Prog.**	Total Obs.	Unique Prog.
Bachelor	3636	1740	2221 (61.1%)	1330 (76.4%)
Master	4621	2594	3079 (66.6%)	2016 (77.7%)
Integrated Master	267	157	114 (42.7%)	80 (51.0%)
Total	**8524**	**4491**	**5414 (63.5%)**	**3426 (76.3%)**

The dataset comprises 60 distinct variables that capture multiple dimensions of degree programs, including program identification, structural characteristics, accreditation information,

[4]Source: DGES website, https://www.dges.gov.pt/pt/pagina/alteracao-3

and institutional attributes. Table 2 provides detailed descriptions of the key variables used in this research.

Table 2: Definitions of Key Variables

Variable	**Definition**
Degree Name	Title of the study program
DGES Registration Code	Unique identifier assigned by DGES
Education Level	Bachelor, Master or Integrated Master degree
Institution Sector	Classification of institution as public or private
Education Type	Classification of institution as university (*Universitário*) or polytechnic (*Politécnico*)
Training Area	Major group from the CNAEF (National Classification of Education and Training Areas)
ECTS	Total ECTS credits required for degree completion according to official program structure
Total ECTS plan	Sum of ECTS credits from all courses listed in the available curricular plan (may differ from required ECTS)
Accreditation Date	Date of accreditation of the curricular plan
Number of years accredited	Duration of the accreditation period (in years)

In addition to those presented in Table 2, two critical variables capture detailed curricular information in JSON format. ***Curricular structure*** describes the organization of programs into scientific areas, specifying the distribution of mandatory and optional ECTS credits across this category. ***Curricular plan*** contains granular information about individual curricular units, including their names, ECTS credits, and semester placement.

While the dataset described above represents accreditation events, with each observation corresponding to a single accreditation process, our analysis also requires a panel structure to examine temporal patterns in AI integration. Therefore, we constructed a panel dataset where each degree program is represented once per year for the period during which it was active. To ensure all panel observations could be analyzed for AI content, we restricted the panel to registrations that had at least one accreditation with available curricular plan information. To create that dataset, the first step was to identify the first and last year each registration was active. We obtained the first year by extracting it from the DGES registration code, and the last year was calculated as either the documented discontinuation year or the accreditation year plus the number of years of accreditation validity, capped at 2025. Following that, we generated a complete time series for each program, creating one observation per year from its first registration to its last available year. We then matched each program-year observation to the corresponding accreditation information from the base dataset. When a program had multiple accreditations in the same year, we retained the observation with the most complete curricular plan information, prioritizing plans closest to the expected ECTS total and with active accreditation status. To populate program-year observations without direct accreditation matches, we applied a forward-fill strategy followed by backward-fill. This approach carries each curricular plan forward through subsequent years until a new accreditation with updated plan information appears and fills early years with the program's first available curricular plan in our data. The resulting panel dataset contains 34570 program-year observations across 3352 unique degree programs. Programs in the panel have varying time periods from 2010 to 2025 depending on their registration and discontinuation dates, creating an unbalanced panel structure.

3.2 AI Classification

Classifying each curricular unit according to its AI presence level was a critical step in this project, as this classification formed the basis for computing our dependent variables. Ideally, this would have been accomplished by analyzing detailed methodologies and learning outcomes for each curricular unit to identify those related to or including AI topics. However, without this information available, we decided to follow a keyword identification approach. The classification system distinguishes between three categories: *Direct AI*, *AI Related*, and *Not AI*. Courses are assigned to the *Direct AI* category when they explicitly address artificial intelligence concepts or applications as core subject matter. The *AI Related* category includes courses that provide foundational knowledge or complementary skills essential to AI development and implementation. Courses that fall outside both categories are classified as *Not AI*.

To implement this methodology we created two keyword dictionaries: one for Direct AI content and another for AI Related content. The construction of these dictionaries was grounded in established educational frameworks for artificial intelligence and computer science, specifically the AI Educational Taxonomy developed by the AI Doctoral Academy (AIDA) International AI

Doctoral Academy (nd) and the Computer Science Curricula 2023 (CS2023) guidelines produced jointly by the Association for Computing Machinery (ACM), IEEE Computer Society, and the Association for the Advancement of Artificial Intelligence (AAAI) Kumar et al. (2024). These sources ensured that the keyword dictionaries reflect current academic consensus on AI education and capture a wide spectrum of topics relevant to AI degree programs.

The Direct AI dictionary includes words related to fundamental areas of the artificial intelligence domain, machine learning and deep learning, natural language processing, computer vision, robotics, generative AI, probabilistic reasoning, and AI ethics. The AI Related dictionary captures foundational and complementary areas such as data science, statistics, linear algebra, algorithms, programming languages, data management, parallel computing, and visualization techniques. The complete keyword dictionaries, including both English and Portuguese terms, are provided in Appendix A.1.

The classification process consists of three steps. First, each curricular unit name is preprocessed to ensure robust keyword matching: conversion to lowercase, removal of accent marks, elimination of special characters, and whitespace normalization. Second, the preprocessed text is matched against both keyword dictionaries by checking if any keyword appears within the course name. Third, courses are classified hierarchically: *Direct AI* if matching at least one Direct AI keyword, *AI Related* if matching only AI Related keywords, and *Not AI* if no matches are found.

Based on the AI classification of curricular units, we created new variables to measure AI presence across degree programs. These variables capture both the absolute scale and relative proportion of AI content within curricula and serve as the dependent variables in our analysis. For both categories, Direct AI and AI related, we constructed absolute and percentage measures in terms of total number of courses and ECTS credits. The mentioned variables are:

- **Total Courses Direct AI / AI Related**: The count of curricular units classified as Direct AI or AI Related within each degree program.
- **Percentage Courses Direct AI / AI Related**: The proportion of Direct AI or AI Related courses relative to the total number of curricular units in the program, calculated as:

$$\text{Percentage Courses} = \frac{\text{Total Courses (Direct AI or AI Related)}}{\text{Total Curricular Units in the Plan}} \times 100 \tag{1}$$

- **Total ECTS Direct AI / AI Related**: The sum of ECTS credits assigned to Direct AI or AI Related courses.

- **Percentage ECTS Direct AI / AI Related**: The proportion of Direct AI or AI Related ECTS credits relative to the total ECTS credits in the program, calculated as:

$$\text{Percentage ECTS} = \frac{\text{Total ECTS (Direct AI or AI Related)}}{\text{Total ECTS in Curricular Plan}} \times 100 \tag{2}$$

These metrics provide complementary perspectives on AI integration: absolute measures reflect the scale and breadth of AI content offered, while percentage measures control for variation in program size and structure, enabling direct comparison across programs of different lengths and credit requirements. Together, they allow us to assess both the extent of AI education (how many courses and credits) and its relative importance within each curricular plan (what proportion of the program).

3.3 Sample Characteristics

This section presents the main characteristics of the panel dataset described above. We first examine the distribution of our dependent variables measuring AI presence in curricula, followed by an overview of the sample composition and key characteristics.

AI Metrics

Table 3 presents summary statistics for the eight variables measuring AI presence across degree programs as well as for the variables representing the total number of curricular units and ECTS credits for better interpretation. The most standout finding is the high prevalence of programs with zero AI content: 94.3% of program-year observations contain no Direct AI courses, and 60.4% have no AI Related courses. This pattern is consistent across all metrics indicating that AI education remains largely absent from the vast majority of Portuguese degree programs during the period studied. The mean values across all observations are consequently very low, 0.08 Direct AI courses representing 0.27% of ECTS credits, and 0.91 AI Related courses representing 2.91% of ECTS credits. However, the maximum values indicate substantial variation among programs that do include AI content, with some offering up to 11 Direct AI courses or 18 AI Related courses.

Table 3: Summary Statistics for AI Metrics

Variable	N	% = 0	Mean	SD	Min	Q1	Median	Q3	Max
Num. Curricular Units	34,570	0.0	25.57	12.29	1	14.00	25.00	34.00	86
Total ECTS Plan	34,570	0.0	185.33	96.21	38	120.00	180.00	198.00	1,066
Direct AI									
Total Courses	34,570	94.3	0.08	0.43	0.00	0.00	0.00	0.00	11.00
% Courses	34,570	94.3	0.33	1.99	0.00	0.00	0.00	0.00	68.75
Total ECTS	34,570	94.3	0.51	3.90	0.00	0.00	0.00	0.00	246.00
% ECTS	34,570	94.3	0.27	1.71	0.00	0.00	0.00	0.00	70.69
AI Related									
Total Courses	34,570	60.4	0.91	1.71	0.00	0.00	0.00	1.00	18.00
% Courses	34,570	60.4	3.31	6.12	0.00	0.00	0.00	5.00	54.55
Total ECTS	34,570	60.5	5.39	11.16	0.00	0.00	0.00	6.00	129.00
% ECTS	34,570	60.5	2.91	5.84	0.00	0.00	0.00	3.57	65.00

For context, the average program in our sample contains 25.57 curricular units totaling 185.33 ECTS credits. The maximum values of 86 curricular units and 1066 ECTS credits reflect programs that offer multiple specialization paths or include optional courses across different degree variants, not all of which are required for individual students to complete.[5] This indicates that the observed scarcity of AI content is not a result of incomplete data but rather reflects true patterns in Portuguese higher education curricula. The contrast between the complete coverage of program structure data and the extreme concentration of zeros in AI metrics highlights the selective and uneven integration of AI education across institutions and degree programs.

To further investigate the characteristics of AI content integration, Figures 1 and 2 provide detailed analyses of programs that include AI-related coursework. Figure 1 displays the distribution of AI content among program-year observations with non-zero values, revealing some heterogeneity in how institutions incorporate AI into their curricula. Among the 5.7% of observations with Direct AI courses, the distribution is heavily right-skewed: the median program

[5] For example, some master's programs allow students to choose between different specializations, each with its own set of optional courses, resulting in a large total number of courses listed in the curricular plan even though students complete only a subset.

allocates 6.00 ECTS credits to Direct AI content, while the mean of the distribution is 8.95 credits. The concentration of observations at lower percentages suggests that most programs with Direct AI content adopt a selective integration approach, offering one or two specialized courses rather than comprehensive AI training. In contrast, among the 39.5% of observations with AI Related content, programs allocate a median of 7.50 and a mean of 13.63 ECTS credits to these foundational subjects. The broader distribution and higher prevalence of AI Related content reflect the more spread integration of those supporting disciplines.

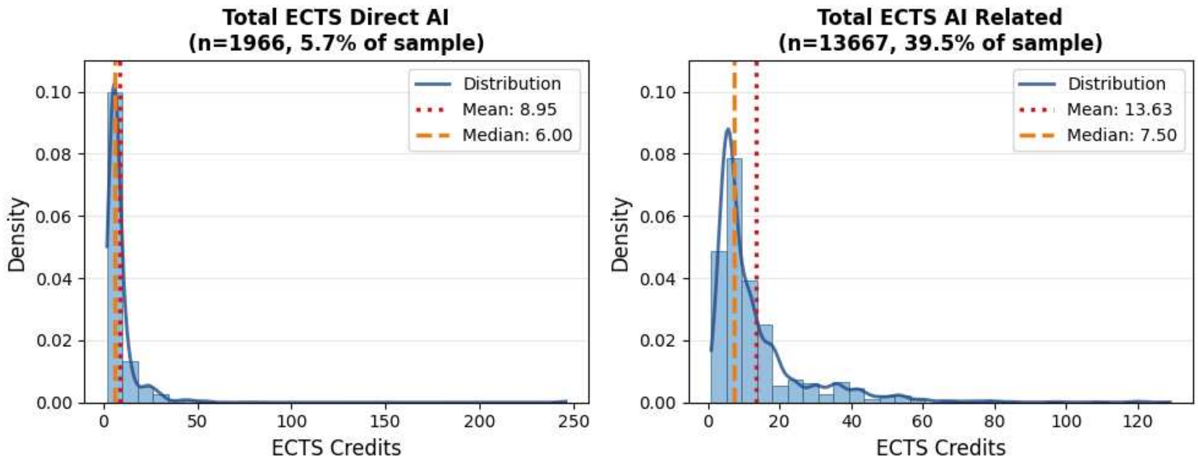

Figure 1: Distribution of AI Content (Non-Zero Observations Only)

Figure 2 illustrates the temporal evolution of AI content across Portuguese degree programs from 2010 to 2025. The left panel shows absolute numbers, while the right panel presents the normalized composition. First, throughout the 2011-2025 period, the proportion of programs with no AI content remains relatively stable around approximately 60%, indicating that AI-related coursework, while growing in absolute numbers, has not yet achieved widespread diffusion across Portuguese higher education. Second, programs with AI Related only content constitute the largest category of AI-inclusive programs, maintaining a consistent between approximately 30% to 40% share throughout the period. Third, programs with Direct AI Only content remain exceptionally rare, never exceeding 3% of the sample, which may suggest that explicit AI courses are typically offered alongside supporting foundational coursework rather than in isolation. Fourth, programs offering Both Direct and AI Related content show modest but notable growth, particularly after 2020, rising from approximately 3/4% to 7% of programs by 2025. This acceleration in recent years coincides with increased attention to AI following the emergence of widely accessible generative AI technologies, though the overall prevalence remains limited. These temporal patterns suggest that while foundational AI-enabling subjects may have been well-established in certain fields since at least 2011, comprehensive integration of explicit AI courses into degree programs remains concentrated in a minority of specialized programs.

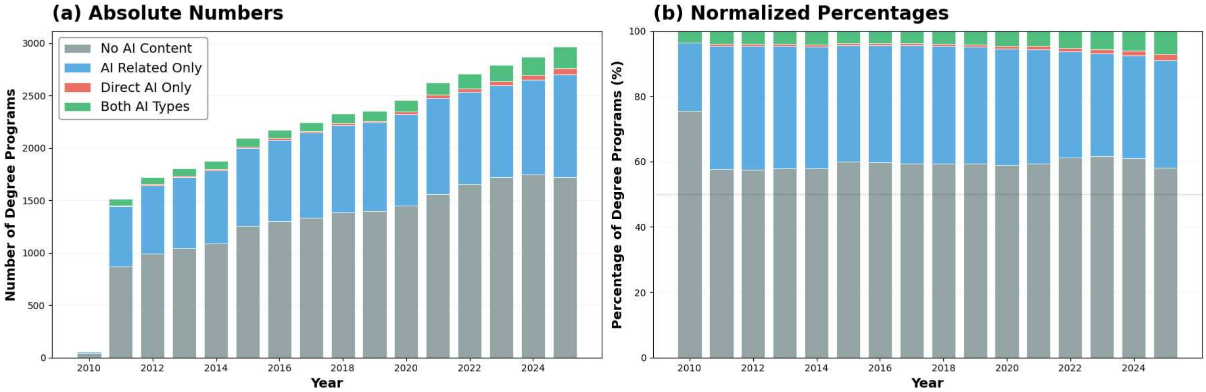

Figure 2: Distribution of Programs by AI Content Type Over Time

Other Characteristics

Having analyzed the distribution of AI content across programs, we now characterize the broader composition of our panel dataset to assess its alignment with the structure of Portuguese higher education. As previously stated, our dataset comprises 34,570 program-year observations including 3,352 unique degree programs observed between 2010 and 2025. Although, given the extremely small number of observations for the year 2010 (as seen in Figure 2), we will exclude it from the subsequent analyzes. The panel is unbalanced, with programs entering and exiting at different times due to new program creation and discontinuations. The number of programs observed per year grows from approximately 1,500 in 2011 to nearly 3,000 by 2025, which may reflect two things, the expansion of Portuguese higher education and improvements in data availability over time. The dataset includes 96 distinct institutions that operated at some point during the study period, including both currently active institutions and those that merged or closed during this time period.

Our analysis focuses on bachelor's, master's, and integrated master's (bachelor and master) programs, which represent the core of structured degree offerings in Portuguese higher education. We exclude two program types, CTeSP, which are short-cycle technical programs with distinct curricular structures, and doctoral programs, which emphasize research training rather than structured coursework. This scope restriction is consistent with our focus on AI integration in formal curricula, as CTeSP programs follow different pedagogical models and PhD programs lack standardized course structures comparable to bachelor's and master's degrees. Master's programs constitute the majority of the sample with 1,963 programs (58.6%), followed by bachelor's with 1,310 programs (39.1%), and integrated master's with only 79 programs (2.4%). Regarding institutional sector, public institutions account for 2,567 programs (76.6%), while private institutions have 782 programs (23.4%). By education type, university programs comprise 2,081 programs (62.1%), while polytechnic programs represent 1,270 programs (37.9%). Detailed breakdowns by degree level, institutional sector, and education type are provided in Appendix A.2.

In addition, we also analyzed the distribution of unique programs across training areas as shown in Figure 3. Social sciences, commerce, and law represent the largest category with over one-quarter of all programs (26.3%), followed by engineering and health programs. The distribution demonstrates that our panel extends well beyond STEM fields, with disciplinary diversity ensuring that our findings regarding AI integration reflect patterns across varied contexts rather than being limited to technical fields where AI content might be most expected.

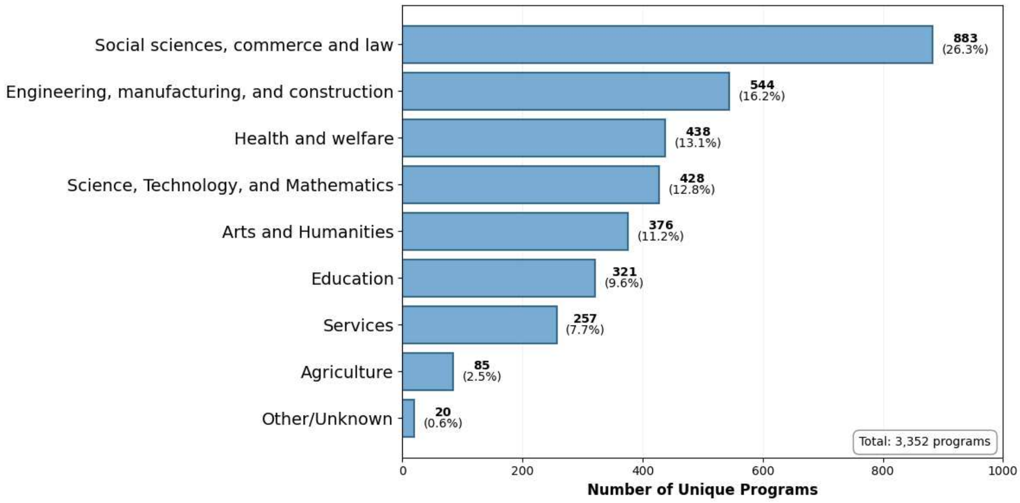

Figure 3: Distribution of Programs by Training Area

Lastly, we observed that programs are concentrated in Portugal's major metropolitan areas, with Lisboa (28.2%) and Porto (18.0%) representing 46.2% of all degree programs. Coimbra (9.1%), Braga (8.1%) and Aveiro (5.3%) represent secondary education hubs, while the remaining programs are distributed across smaller cities. The sample has representation for all 18 Portugal mainland districts as well as for the autonomous regions of Madeira and Açores.

3.4 Empirical Strategy

In this section we will explain the methods used to investigate how the emergence of widely accessible generative AI technologies, particularly ChatGPT, has transformed curricular plans in Portuguese higher education. Our empirical strategy compares programs that were accredited after 2022, when ChatGPT was released publicly, to those that were updated before, allowing us to understand if institutions have responded by integrating more AI content into their curricula. We focus on re-accreditation as the primary mechanism for curricular change. Re-accreditation involves updating the curricular plan of an existing program through formal review by A3ES, typically occurring on regular cycles every six years. During re-accreditation, institutions review and potentially revise course offerings and program structure. Since re-accreditation timing is predetermined by previous accreditation dates and regulatory requirements set by A3ES, the specific timing of re-accreditation is plausibly exogenous to recent trends in AI technology. Moreover, ChatGPT's release timing and impact were difficult to anticipate, making strategic delays to incorporate AI content unlikely. This exogeneity allows us to interpret our interaction coefficients as capturing the causal effect of post-ChatGPT re-accreditation timing on AI integration decisions.

To implement our identification strategy, we constructed two variables that capture the timing and nature of curricular changes:

- **Reaccredited** ($Reaccredited_{it}$): A binary indicator that is equal to 1 if program i was re-accredited in year t, and 0 otherwise. A program may undergo multiple re-accreditations during our observation period (2011-2025). This variable helps us identify the specific moments when institutions had the opportunity to update existing programs.
- **After** ($After_t$): A binary time indicator equal to 1 for years 2023-2025, and 0 for years 2011-2022, allowing us to identify the post-ChatGPT period.

Our primary model specification aims to understand if programs re-accredited after the ChatGPT release period (2023-2025) integrate more AI content than programs re-accredited before that period (2011-2022). It is represented by the following equation:

$$\text{Y}_{it} = \beta_0 + \beta_1(\text{Reaccredited}_{it} \times \text{After}_t) + \gamma_i + \delta_t + \epsilon_{it} \tag{3}$$

where γ_i represents program fixed effects and δ_t represents year fixed effects.

The interaction term β_1 is our coefficient of interest, it measures whether re-accreditations occurring after ChatGPT's release result in higher AI content than re-accreditations that occurred before. A positive and significant β_1 would indicate that institutions have responded to generative AI's emergence by integrating more AI content during the re-accreditation process. The program fixed effects (γ_i) control for time-invariant program characteristics, while year fixed effects (δ_t) account for general trends in AI integration affecting all programs. The interaction term thus identifies the incremental effect of re-accreditation timing on AI adoption, isolating the post-ChatGPT re-accreditation effect from both baseline program differences and general temporal trends.

We analyze two types of outcome measures, Y_{it}, capturing different dimensions of AI integration. For the probability of adoption (Section 4.2), we use *Has Any Direct AI*, a binary indicator equal to 1 if the program includes at least one Direct AI course and 0 otherwise. This captures whether programs adopt any AI content. For the quantity of adoption (Section 4.3), we use *Total ECTS Direct AI*, which counts the ECTS credits assigned to Direct AI courses. This measures how much AI content is adopted. Together, these measures distinguish between the decision to introduce AI content and the intensity of AI integration, providing a comprehensive view of institutional responses.

For binary outcomes (*Has Any Direct AI*), we employ two complementary estimation approaches, logistic regression and linear probability models (LPM). Logit models with program fixed effects suffer from the incidental parameter problem (Greene, 2004) and will drop observations without within-program variation in the outcome. The LPM model will estimate the probability of AI adoption directly in percentage point terms, using the full panel dataset with program and year fixed effects. While LPM can produce predictions outside the [0, 1] interval, it provides easily interpretable coefficients and utilizes all available data. For count outcomes (*Total ECTS Direct AI*), we use models that explicitly account for the extreme zero-inflation in our data, where 94.3% of observations have zero Direct AI ECTS. Zero-inflated negative binomial (ZINB) models serve as our primary specification for these outcomes. ZINB models estimate two separate processes: a logistic regression predicting whether a program contains any AI content (zero-inflation component) and a negative binomial regression predicting the amount of AI content conditional on its presence (count component). This two-part structure accommodates both structural zeros (programs in fields where AI content would never be relevant) and sampling zeros (programs in AI-relevant fields that happen to have zero AI content in a given year). For robustness, we also estimate standard negative binomial (NegBin) models that treat outcomes as overdispersed count variables without explicitly modeling zero-inflation. Comparing ZINB and NegBin results diagnoses whether findings are sensitive to assumptions about the zero-generating process.

To understand which types of programs and fields are most responsive to the emergence of AI, we estimate separate versions of our main models for key subgroups: education level and training areas. In the first case, we estimate separate models for bachelor's and master's degrees (excluding integrated masters due to their small sample size) to test whether the general conclusions are valid for both types of degrees. For the second case, we estimate separate models for five major categories: (1) Social Sciences, Commerce, and Law, (2) Engineering, Manufacturing, and Construction, (3) Arts and Humanities, (4) Health and Welfare, and (5) Science, Technology, and Mathematics. These heterogeneity analyses reveal which institutional contexts facilitate curricular adaptation and whether AI education is becoming broadly distributed or remains concentrated in technical programs.

4 Results

This section presents the empirical findings of our analysis. We organize our results into three parts. First, we present an exploratory analysis that documents trends in the integration of Direct AI content across Portuguese degree programs from 2011 to 2025. Second, we present regression results analyzing the probability that programs include Direct AI content, comparing programs accredited before and after the release of ChatGPT. Third, we examine the quantity of Direct AI content being integrated.

4.1 Descriptive Trends in AI Integration

We begin by examining temporal patterns in the integration of Direct AI content across Portuguese higher education. As reported in Section 3.3, Direct AI courses remain relatively rare, with only 5.7% of observations in our panel dataset containing those courses (Table 3). Despite this overall scarcity, the prevalence of Direct AI content has doubled over our observation period. The proportion of programs offering any Direct AI content rose from 4.6% in 2011 to 9.1% in 2025, as shown in Figure 4. Overall, there are 310 distinct programs that have AI curricular units during the observation period.

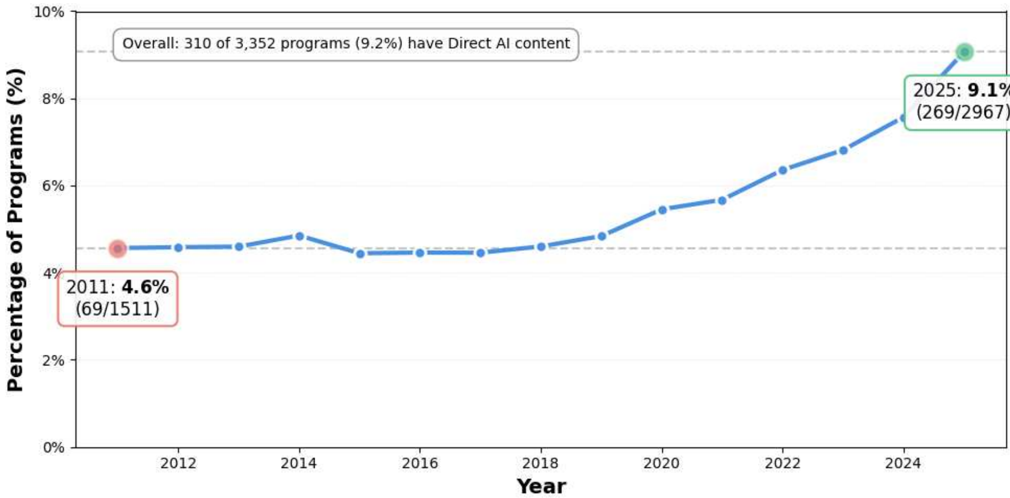

Figure 4: Percentage of Programs with Direct AI Content Over Time

Figure 4 suggests the upward trajectory started in 2020, with an acceleration of the trend in the most recent years. From 2011 to 2019, the proportion remained relatively flat around 4 to 5%, indicating minimal change in AI curricular integration during this period. This pre-2022 growth suggests that universities were already responding to the increasing prominence of AI technologies in industry and research, including advancements in machine learning, computer vision, and natural language processing. ChatGPT's November 2022 release appears to have accelerated an already-existing trend rather than initiating it from scratch. This growth reflects not only the addition of AI courses to existing programs but also the emergence of new programs more dedicated to artificial intelligence. We identified 19 programs with ”Artificial Intelligence” explicitly in their titles, 6 bachelor's degrees and 13 master's degrees. Interestingly, all of these programs were created in 2020 or after (see Appendix A.3 for a complete list and more details). This fact suggests universities are not only adding elective AI courses to existing programs but also creating dedicated AI degrees, signaling a substantial commitment to AI education.

In Portugal, institutions update their curricula through two primary mechanisms: the creation of new programs or the re-accreditation of existing ones. To better understand the source of the major AI introduction, in Figure 5 we compare the average Direct AI ECTS between the two options over time.

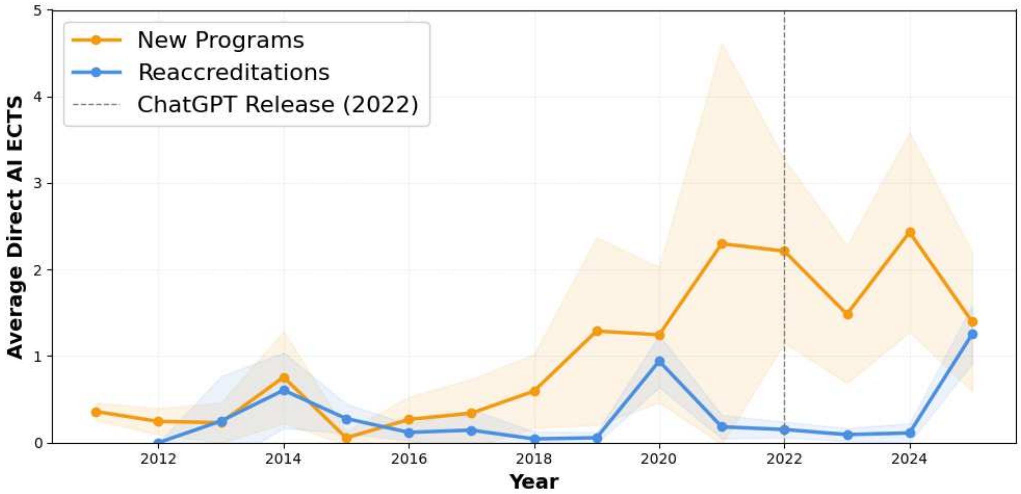

Figure 5: Average Direct AI ECTS: New Programs vs Reaccreditations
Note: Shaded areas represent 95% confidence intervals.

Until 2019, both new programs and re-accreditations show similarly low levels of Direct AI content, indicating minimal AI integration regardless of the pathway. However, starting in 2020, the trend changes. New programs begin allocating substantially more Direct AI content, rising from approximately 0.4 ECTS in 2011 to over 2.0 ECTS by 2021. It is important to note that these means are low due to the large number of programs with zero AI content overall. This trend shift suggests that institutions creating new degree programs from scratch began responding to AI's growing prominence well before the release of ChatGPT. In contrast, re-accreditations remained near zero through most of the period, showing only sporadic increases. The apparent spikes in re-accreditation AI content (in 2020 and 2025) coincide with years that saw higher volumes of re-accreditation activity (see Figure A1 in Appendix A.4 for yearly counts of re-accreditations and new programs). This pattern suggests an emerging trend towards increased AI integration during re-accreditation. Though the limited number of re-accreditations completed thus far in the post-ChatGPT period makes it difficult to establish a clear and consistent pattern. The wide confidence intervals, especially for new programs introduced after 2020, reflect the small sample sizes and underscore the importance of formal regression analysis to control for temporal trends and program characteristics.

Having established that both pathways exhibit increased AI content after 2020, we now focus specifically on whether the timing of re-accreditation relative to ChatGPT's release matters for AI integration. Figure 6 displays average Direct AI ECTS in the years surrounding ChatGPT's 2022 launch, separating programs that were re-accredited after 2022 from those that were not.

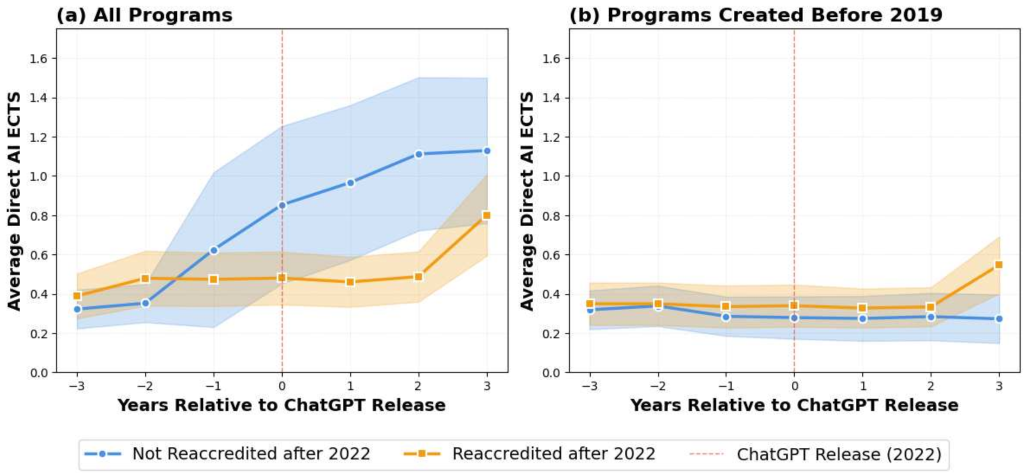

Figure 6: Average Direct AI ECTS around ChatGPT Release (2022)
Note: Shaded areas represent 95% confidence intervals.

Figure 6 (a) shows all programs in our sample. Programs not re-accredited after 2022 (blue line) show a steady upward trend throughout the period, while programs re-accredited after 2022 (orange line) remain flat until a sharp increase three years post-ChatGPT. However, the blue line's upward trend likely reflects the inclusion of recently created programs that entered the sample with AI content already incorporated from inception, rather than changes to existing programs' curricula. To isolate the effect of re-accreditation on existing programs, Figure 6 (b) restricts the sample to programs created before 2019, excluding any programs launched during the observation window. This restriction reveals a markedly different pattern. Programs not re-accredited after 2022 (blue line) now show no upward trend, remaining flat throughout the entire period. In contrast, programs re-accredited after 2022 (orange line) remain similarly flat pre-ChatGPT, then exhibit a substantial increase three years post-ChatGPT, like plot (a). This divergence, visible only when new programs are excluded, provides evidence that the re-accreditation process itself also serves as the mechanism through which existing programs integrate AI content. Analyzing both plots side by side reveals that the incorporation of AI in creating a new program began much earlier than through re-accreditation.

To further understand AI integration and explore if these aggregate patterns mask some heterogeneity across program characteristics, we examine differences by degree level and training area. Figure 7 shows the evolution of average Direct AI ECTS by degree level from 2011 to 2025. We excluded the integrated master's degrees due to the small number of those programs.

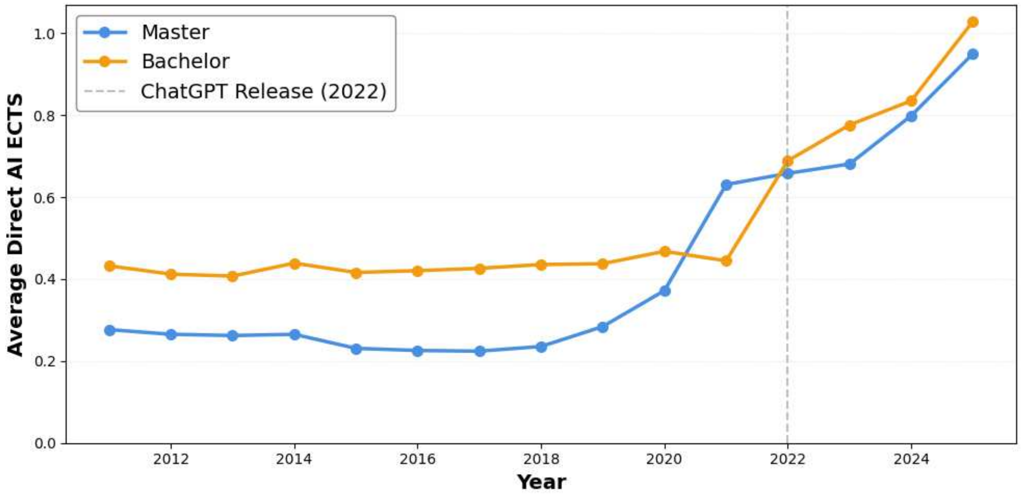

Figure 7: Average Direct AI ECTS by Year and Degree Level

Master's programs (blue line) begin from a lower baseline, averaging below 0.3 ECTS from 2011 to 2019. However, those appear to have anticipated AI integration, beginning their upward trajectory in 2020, two years before ChatGPT's release, and showing consistent year-over-year increases thereafter, reaching approximately 0.95 ECTS by 2025. Bachelor's programs (orange line) start from a higher baseline, remaining stable just above 0.4 ECTS from 2011 through 2021. From 2022 onward, bachelor's programs show steep, continuous year-over-year growth, ultimately converging with and slightly surpassing master's programs by 2024-2025, reaching approximately 1.02 ECTS. The two-year lead time in master's programs suggests that graduate education may have been more responsive to earlier AI developments prior to generative AI's public breakthrough. However, once bachelor's programs began integrating AI content, they did so rapidly, catching up within three years. Both degree levels show sustained acceleration with no signs of flattening through 2025, indicating ongoing curricular transformation at both undergraduate and graduate levels.

Next, we repeat the above analysis by training area, focusing on the top 5 biggest areas. With this we aim to determine if AI integration is being made solely in technical fields or is also changing other areas. That analysis is shown in Figure 8.

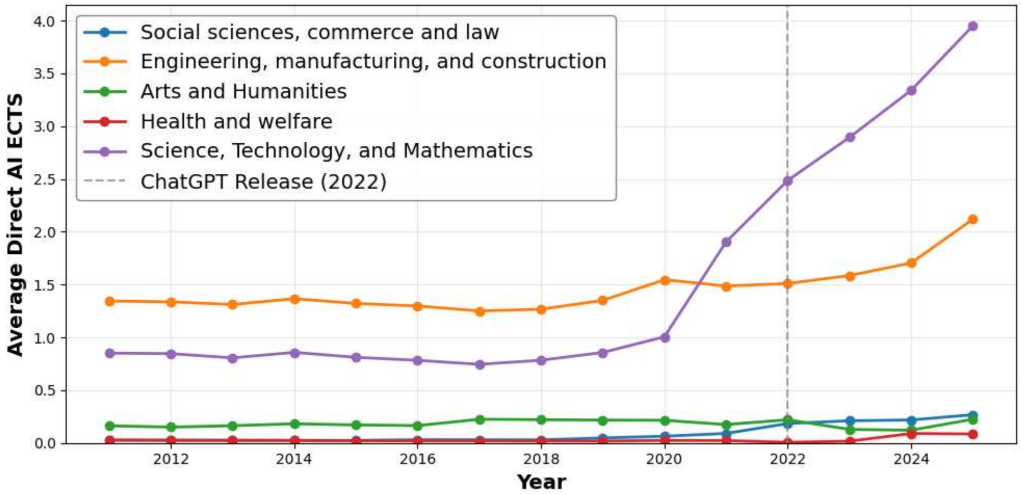

Figure 8: Average Direct AI ECTS by Year and Training Area

Disciplinary differences in AI integration are more pronounced than degree level differences. Science, technology, and mathematics programs (purple line) stand out as the area with biggest change. Engineering, manufacturing, and construction programs (orange line) follow a similar pattern with more modest increases. Social sciences, commerce and law (blue line) show a slight increase in the most recent years. In contrast, arts and humanities (green line), and health and welfare (red line) remain essentially flat near zero throughout the entire observation period, showing no visible response. This concentration reveals that Direct AI education remains almost entirely confined to technical fields, with minimal and very recent diffusion to other disciplines. Detailed plots for individual training areas are provided in Figure A2 of Appendix A.4.

These descriptive patterns reveal that Direct AI content, while still scarce, has been growing since 2020 with acceleration after ChatGPT's release. New programs incorporated AI earlier than re-accreditations, master's programs led bachelor's programs by two years, and AI integration remains heavily concentrated in technical fields. Our regression analysis tests whether these patterns hold when controlling for temporal trends and program characteristics.

4.2 Probability of AI Content

We now turn to formal regression analysis to test whether the descriptive patterns observed in the previous section reflect systematic institutional responses to AI's emergence. This section examines the probability that programs include Direct AI content, focusing on whether programs re-accredited after ChatGPT's release in 2022 are more likely to adopt AI courses than those re-accredited before. We present results from our main overall specification, followed by heterogeneity analyses by degree level and training area.

Table 4 presents our main results estimating the effect of post-ChatGPT re-accreditation on the probability of AI adoption. We report results from both logistic regression (Column 1) and a linear probability model (Column 2).

Table 4: Effect of Post-GPT Reaccreditation on AI Adoption

	AI Adoption (*has_any_ai*)	
	Logit	LPM
	(1)	(2)
Reaccredited × After	5.665***	0.027***
	(1.196)	(0.005)
Observations	961	34,371
R^2	0.551	0.936
Adjusted R^2	0.404	0.930
Program FE	Yes	Yes
Year FE	Yes	Yes

Notes: Standard errors are clustered at the program level.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

The logit model yields a positive and significant coefficient on the Reaccredited × After interaction, indicating that programs re-accredited after ChatGPT's release have higher odds of including Direct AI content than programs re-accredited before. However, this estimate should be interpreted with caution due to an important sample selection issue. The logit specification with program fixed effects uses only 961 observations, which is less than 3% of the full sample, because it necessarily drops all programs without within-program variation in AI adoption. This means the logit model identifies the effect exclusively among programs that changed their AI status at some point during the observation period. This selected sample represents programs already on the margin of AI adoption, making them fundamentally different from the broader population of programs that consistently offered no AI content throughout the entire period. The linear probability model, using almost the full sample, provides a more representative and directly interpretable estimate. The coefficient of 0.027 indicates that programs re-accredited after ChatGPT are 2.7 percentage points more likely to include Direct AI content than programs re-accredited before, controlling for program and year fixed effects. Given that only 5.7% of program-year observations in our panel include Direct AI content (Table 3), this 2.7 percentage point increase represents a substantial relative effect. Importantly, this estimate reflects the average effect across all programs, including those in fields where AI adoption remains rare.

In summary, both specifications point to the same conclusion, re-accreditation after ChatGPT's release is associated with statistically significant higher probability of AI adoption. However, the large difference in sample sizes means the logit estimate reflects a more selected population of programs with demonstrated capacity for curricular change, while the LPM estimate captures the average response across the entire higher education system.

Having analyzed the general effect of post-ChatGPT re-accreditation on AI adoption, we now examine whether this effect varies across program characteristics, beginning with degree level. Table 5 presents separate estimates for bachelor's and master's programs.

Table 5: Effect of Post-GPT Reaccreditation on AI Adoption by Degree Level

	AI Adoption (*has_any_ai*)			
	Bachelor		Master	
	Logit	LPM	Logit	LPM
	(1)	(2)	(3)	(4)
Reaccredited × After	8.153***	0.041***	4.371***	0.017**
	(1.284)	(0.008)	(0.906)	(0.006)
Observations	478	14,733	468	18,738
R^2	0.677	0.946	0.472	0.921
Adjusted R^2	0.498	0.941	0.306	0.912
Program FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: Standard errors are clustered at the program level.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Both degree levels show significant positive effects of post-ChatGPT re-accreditation on AI adoption across both model specifications. However, bachelor's programs exhibit a substantially larger response. The LPM estimates reveal that bachelor's programs experience a 4.1 percentage point increase in the probability of AI adoption, compared to 1.7 percentage points for master's programs. The logit specifications also confirm this pattern. This heterogeneity aligns directly with the temporal patterns observed in Figure 7. Master's programs began integrating AI content earlier, starting their upward trajectory in 2020, while bachelor's programs initiated their growth in 2022, coinciding with ChatGPT's release. The larger effect for bachelor's programs suggests that undergraduate education is catching up to graduate education in AI integration, with institutions using post-ChatGPT re-accreditation opportunities to introduce foundational AI literacy at the undergraduate level rather than reserving it exclusively for graduate study.

We proceed with our analysis by testing whether the post-ChatGPT re-accreditation effect varies across training areas. Table 6 presents separate LPM estimates for the five largest training area categories in our sample. We focus exclusively on LPM specifications for these separate models because dividing the already-restricted logit sample (961 observations) across training areas would yield insufficient observations for reliable estimation in each field.

Table 6: Effect of Post-GPT Reaccreditation on AI Adoption by Training Area

	AI Adoption (*has_any_ai*)				
	Law	Engineering	Arts	Health	STM
	(1)	(2)	(3)	(4)	(5)
Reaccredited × After	0.014**	0.115***	−0.007	0.003	0.058
	(0.005)	(0.026)	(0.007)	(0.003)	(0.033)
Observations	9,666	5,438	4,168	4,122	4,115
R^2	0.913	0.935	0.848	0.914	0.960
Adjusted R^2	0.905	0.927	0.833	0.904	0.955
Program FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Notes: Standard errors are clustered at the program level.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

The results reveal striking heterogeneity. Engineering, manufacturing, and construction programs show the largest response with a 11.5 percentage point increase. Social sciences, commerce, and law programs exhibit a modest but statistically significant effect of 1.4 percentage points. Science, technology, and mathematics (STM) programs show a coefficient of 0.058, though this does not reach statistical significance. However, the lack of significance for STM programs likely reflects that these programs had already begun substantial AI integration before the post-ChatGPT period, as observed in Figure 8, leaving less room for additional increases during post-ChatGPT re-accreditations. In contrast, arts and humanities and health and welfare show not significant coefficients near zero indicating no response. A triple interaction model including all training areas simultaneously (presented in Table A6 of Appendix A.5) confirms these patterns with nearly identical coefficient estimates.

These findings demonstrate that the post-ChatGPT re-accreditation effect is still heavily concentrated in technical fields where AI technologies most directly impact professional practice. The modest effect in social sciences suggests the start of some diffusion across fields, possibly reflecting growing recognition of AI's relevance for data-intensive research methods. However, the complete absence of effects in arts, humanities, and health indicates that AI curricular integration has not yet achieved broad interdisciplinary distribution across Portuguese higher education.

To conclude this section, the above regression results confirm that post-ChatGPT re-accreditation significantly increases the probability of AI adoption, with an average effect of 2.7 percentage points across all programs. However, this aggregate effect masks substantial heterogeneity. Bachelor's programs respond more aggressively than master's programs, and engineering fields show effects exceeding 11 percentage points while arts, humanities, and health show no response.

4.3 Quantity of AI Content

Thus far we have established that post-ChatGPT re-accreditation increases the probability of including Direct AI content. We now extend that analysis to quantify how much AI content is being integrated. Table 7 presents the results from count models using *Total ECTS Direct AI* as the outcome variable. We estimate three model specifications to account for the extreme zero-inflation in our data. Column 1 presents a negative binomial model without program fixed effects, using the full sample of 34,513 observations. Column 2 presents a negative binomial model with program fixed effects, which reduces the sample to 2,594 observations as it necessarily drops programs without variation in AI ECTS over time. Column 3 presents a zero-inflated negative binomial (ZINB) model with two components: a count model predicting the amount of AI content conditional on non-zero values, and a zero-inflation model predicting the probability of structural zeros.

Table 7: Effect of Post-ChatGPT Reaccreditation on Direct AI ECTS: Count Models

	Total ECTS Direct AI		
	NegBin		ZINB
	(1)	(2)	(3)
Count Model (Amount of AI Content):			
Reaccredited × After	−0.339	0.505*	−0.252
	(0.198)	(0.216)	(0.176)
Zero-Inflation Model (Probability of Zero):			
Reaccredited × After	—	—	−0.030
			(0.098)
Observations	34,513	2,594	34,513
R^2	0.003	0.302	0.028
Adjusted R^2	0.002	0.262	0.027
Program FE	No	Yes	No
Year FE	Yes	Yes	Yes

Notes: Standard errors are clustered at the program level.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

The results reveal no significant effects in the models using the full sample. Both the negative binomial model without program fixed effects (Column 1) and the ZINB model (Column 3) fail to detect significant post-ChatGPT re-accreditation effects on the quantity of AI ECTS. In contrast, the negative binomial model with program fixed effects (Column 2) yields a positive and significant coefficient of 0.505. However, this model uses only 2,594 observations because it requires within-program variation in AI ECTS. Similar to the logit model in Table 4, this specification identifies effects exclusively among programs that changed their AI ECTS content during the observation period, representing a selected sample of programs already demonstrating curricular flexibility.

These findings suggest that while post-ChatGPT re-accreditation increases the probability that programs adopt any AI content (as shown in Section 4.2), the quantity of AI ECTS being integrated remains too modest to produce detectable effects in count models using the full sample. The combination of extreme zero-inflation and the small amounts of AI content among programs that do adopt means that aggregate quantity effects remain difficult to detect statistically. The significant effect in Column 2 indicates that among the subset of programs actively changing their curricula, post-ChatGPT re-accreditation does lead to increased AI ECTS. However, across the broader population of programs, AI integration remains at an early stage where presence versus absence is more meaningful than the specific quantity of content.

5 Discussion

This study examined how Portuguese higher education institutions have responded to the emergence of widely accessible generative AI technologies through analyzing curricular changes in degree programs from 2011 to 2025. Our study focused on whether programs re-accredited after ChatGPT's release in November 2022 integrated more Direct AI content than those re-accredited before. Three main findings emerge from our study. First, programs re-accredited after ChatGPT are 2.7 percentage points more likely to include Direct AI content than those re-accredited before, representing a substantial relative increase given that only 5.7% of program-year observations include such content. Second, this effect varies across program types, bachelor's programs show effects more than double those of master's programs, and engineering programs show effects exceeding 11 percentage points while arts, humanities, and health programs show no significant response. Third, while re-accreditation increases the probability of AI adoption, the quantity of AI content being integrated remains modest, with no detectable effects on AI ECTS in models using the full sample. In addition, we also observed substantial growth in AI content among newly created programs, with average Direct AI ECTS in new programs rising from below 0.5 ECTS in the beginning of our observation period to over 2.0 ECTS by 2021. This suggests that AI content was being integrated into new programs before the public release of AI technologies. It was also observed the creation of 19 programs with "Artificial Intelligence" explicitly in their titles, all established in 2020 or later. These findings show that Portuguese higher education is in the early stages of AI curricular integration, with clear responses emerging in technical areas but limited diffusion to other disciplines.

These findings reveal a conservative but systematic institutional response to AI's emergence through the re-accreditation channel. The 2.7 percentage point aggregate effect, while statistically significant, is modest in magnitude. This reflects the structural built-in restrictions in higher education: curricula are not easily malleable, re-accreditation processes involve extensive deliberation, and institutions must balance competing priorities. Moreover, institutions may be responding through methods not captured in our data, such as integrating AI topics into existing courses rather than creating new AI-specific courses. The pronounced heterogeneity reveals that responses follow predictable trends based on perceived relevance. Bachelor's programs show stronger responses than master's programs, suggesting a shift toward treating AI literacy as a foundational competency to be introduced at the undergraduate level rather than reserved for graduate specialization. The extreme concentration in engineering and STM fields, with negligible responses in arts, humanities, and health, indicates that AI integration follows a logic of direct professional relevance. The modest social sciences effect suggests some initial diffusion outside technical domains, while the absence of effects in arts, humanities, and health poses questions regarding whether these fields perceive AI as not relevant or lack the resources to integrate it.

Research on how AI is transforming higher education curricula remains limited, with most existing studies focusing on AI adoption by students and faculty rather than institutional curricular responses. Nonetheless, our findings can be compared to the emerging evidence available. Our results confirm two separate channels of curricular change documented in preliminary international studies. First, the creation of entirely new AI-focused programs, we identified 19 programs with "Artificial Intelligence" explicitly in their titles, all created in 2020 or later. This is consistent with international trends, with the Stanford AI Index reporting that the number of US institutions offering AI-specific bachelor's degrees nearly doubled between 2022 and 2023 Maslej et al. (2025). Second, the integration of AI content into existing programs through re-accreditation. Our analysis suggests that this second pathway shows a modest aggregate effect. This finding indicates that although new AI-specialized programs are emerging rapidly, the wider challenge of integrating AI competencies across existing programs remains at an early stage. The results address a gap in the literature, providing quantitative evidence of how institutions are translating AI education priorities into practice.

This pattern resonates with international evidence showing a persistent gap between policy recognition and curricular implementation. Research on national AI strategies revealed that, despite widespread recognition of AI education's importance, relatively few countries have achieved comprehensive integration beyond specialized degree programs Maslej et al. (2025); Schiff (2022). Our quantitative evidence from Portugal confirms this implementation gap, demonstrating that AI-specific degrees are emerging, but broader curricular integration remains limited. The concentration of effects in STEM fields contradicts emerging labor market evidence that AI literacy must become a broadly distributed competency. Research demonstrates that AI's value is maximized when algorithmic skills are dispersed among workers with domain knowledge Tambe (2025), yet our findings show Portuguese universities preparing primarily technical students while leaving other fields largely unaddressed.

Portugal's national AI strategy emphasized developing AI competencies across the entire population, explicitly aiming to ensure "Portugal will be at the forefront of AI Education for all" INCoDe.2030 (2019). However, our evidence shows implementation remains concentrated in technical domains rather than spreading broadly across fields. The first signs of integration began in 2020, two years before ChatGPT, indicating institutions respond to gradual technological diffusion rather than discrete breakthroughs, though the pace of change remains slower and more uneven than policy aspirations envisioned.

6 Limitations

While this research provides an early comprehensive empirical analysis of AI integration in Portuguese higher education curricula, several limitations warrant acknowledgment. First, our analysis relies on merging data from three distinct sources, each with different reporting standards and temporal coverage. Some information may have been lost during the dataset construction process when variables differed across sources or when matching programs across databases. Additionally, our keyword-based classification approach, while ensuring scalability and objectivity across a large dataset, has inherent limitations. This automated method may not capture nuanced AI content expressed through non-standard terminology or courses that integrate AI concepts without explicit keyword references in their titles. Courses with ambiguous or generic names may be misclassified if they lack clear AI-related keywords, potentially underestimating the true extent of AI presence in curricula. Furthermore, we could not distinguish between mandatory and optional curricular units in program structures, meaning our data may include courses that are rarely offered or in which students rarely enroll, potentially overstating the AI content to which typical students are actually exposed.

In addition, our post-ChatGPT observation period from 2023 to 2025 includes fewer than three full years, making it relatively short for observing institutional change. The full impact of AI on curricula may require several additional accreditation cycles to materialize, as institutions typically review programs on multi-year schedules and curricular changes involve deliberation processes. Our findings therefore capture only the earliest institutional responses. Moreover, our study examines formal curricular content as reflected in course catalogues and accreditation documents. We do not observe how AI topics are taught within courses, the depth of AI instruction provided, or informal educational activities such as workshops or seminars. Institutions may be responding to AI's emergence through these mechanisms without formal curricular changes, meaning our analysis may underestimate the total institutional response.

The generalizability of our findings is also limited by our focus on Portugal's higher education system with its specific institutional characteristics. Countries with different accreditation systems, greater institutional autonomy, or different relationships between universities and employers may exhibit distinct patterns of AI integration.

Several extensions would address these limitations and deepen understanding of AI integration in higher education. Content analysis of course syllabus and learning outcomes would validate our keyword classification and assess the depth and quality of AI instruction. Cross-country comparative studies, particularly within Europe and with the United States, would illuminate how different regulatory and market contexts shape AI curricular adoption. Analysis connecting AI curricular content to graduate employment outcomes and earnings would test whether AI education translates into labor market advantages. Comparative examination of alignment between curricular AI content and actual AI skills demanded in job postings would assess whether higher education is preparing students for workplace needs. Finally, extended observation periods capturing full accreditation cycles post-ChatGPT would reveal whether initial AI adoption persists or fades over time, and whether the pace of integration accelerates or flattens.

7 Conclusion

This study investigated how Portuguese higher education institutions have responded to the rise of widely accessible generative AI by examining curricular changes in degree programs from 2011 to 2025. A major contribution of this work is the dataset itself. We built it from scratch by combining data from three official sources (DGES, A3ES and DR). While our analysis offers initial conclusions about institutional responses to AI, the dataset itself constitutes a valuable resource that can be extended, refined, and used to address other questions about the evolution of higher education.

Three main findings emerge from our analysis. Programs re-accredited after ChatGPT's release are more likely to include Direct AI content than those re-accredited before, but the effect is modest, just 2.7 percentage points. This reflects how slowly curricula changes. Universities need to deliberate, debate, and face real constraints in what they can realistically update. The response also varies by field and degree level. Bachelor's programs show a higher effect, but master's programs seem to have anticipated AI integration. Engineering programs show strong effects, with increases over 11 percentage points, while arts, humanities, and health programs show essentially no response. Clearly, AI integration is happening where it feels most technically relevant. Finally, we see growth through another channel, the creation of entirely new programs. Since 2020, 19 programs with 'Artificial Intelligence' in their official titles have been created. These findings provide some initial quantitative evidence of how a Portuguese higher education is translating AI education priorities into curricular practice. The response is systematic but conservative, with clear adoption in technical fields but limited diffusion to other disciplines. This pattern creates potential mismatches with labor market needs, as emerging evidence suggests AI's value is maximized when competencies are distributed across workers.

To conclude, these are initial conclusions based on the earliest observable institutional responses. The post-ChatGPT period in our data is less than three years, capturing only the beginning of what will likely be a longer process of curricular adaptation. As institutions complete additional accreditation cycles and as labor market signals about AI competencies strengthen, patterns of integration may shift. The dataset we have constructed provides the structure to track these changes over time. Future researchers can extend the temporal coverage, refine the classification of AI content through syllabus analysis, link curricular changes to employment outcomes, or compare Portugal's trajectory with other countries.

References

Acemoglu, D. and Restrepo, P. (2018). Artificial intelligence, automation and work. Working Paper 24196, National Bureau of Economic Research.

Alekseeva, L., Azar, J., Giné, M., Samila, S., and Taska, B. (2021). The demand for ai skills in the labor market. *Labour Economics*, 71:102002.

Bent, D., Handa, K., Durmus, E., Tamkin, A., McCain, M., Ritchie, S., Donegan, R., Martinez, J., and Jones], J. (2025). Anthropic education report: How educators use claude.

Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative ai. Working Paper 32966, National Bureau of Economic Research.

Brynjolfsson, E., Hui, X., and Liu, M. (2025). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Technical report, Stanford Digital Economy Lab, Stanford, CA.

Digital Education Council (2024). Global ai student survey.

Digital Education Council (2025a). Ai in the workplace 2025.

Digital Education Council (2025b). Global ai faculty survey.

Galeano, S., Hodge, N., and Ruder, A. (2025). By degree(s): Measuring employer demand for ai skills by educational requirements. *Federal Reserve Bank of Atlanta, Workforce Currents*.

Greene, W. (2004). The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects. *The Econometrics Journal*, 7(1):98–119.

Handa, K., Bent, D., Tamkin, A., McCain, M., Durmus, E., Stern, M., Schiraldi, M., Huang, S., Ritchie, S., Syverud, S., Jagadish, K., Vo, M., Bell, M., and Ganguli, D. (2025). Anthropic education report: How university students use claude.

INCoDe.2030 (2019). AI Portugal 2030: An innovation and growth strategy to foster artificial intelligence in Portugal in the European context. National strategy document, Portuguese National Initiative on Digital Skills, Lisbon, Portugal.

Instituto Nacional de Estatística (2025). Inquérito à utilização de tecnologias da informação e da comunicação pelas famílias 2025. Statistical report, Instituto Nacional de Estatística (Statistics Portugal), Lisbon, Portugal. ISSN 2184-4534.

International AI Doctoral Academy (n.d.). Ai educational taxonomy. `https://www.i-aida.org/aida-taxonomy/`. Accessed: 2025-11-20.

Kumar, A. N., Raj, R. K., Aly, S. G., Anderson, M. D., Becker, B. A., Blumenthal, R. L., Eaton, E., Epstein, S. L., Goldweber, M., Jalote, P., Lea, D., Oudshoorn, M., Pias, M., Reiser, S., Servin, C., Simha, R., Winters, T., and Xiang, Q. (2024). *Computer Science Curricula 2023*. Association for Computing Machinery, New York, NY, USA.

Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Lotufo, J. B., Rome, A., Shi, A., and Oak, S. (2025). The ai index 2025 annual report. Technical report, AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA.

Schiff, D. (2022). Education for ai, not ai for education: The role of education and ethics in national ai policy strategies. *International Journal of Artificial Intelligence in Education*, 32:527–563.

Tambe, P. (2025). Reskilling the workforce for ai: Domain expertise and algorithmic literacy. *Management Science*.

Wong, W., Aristidou, A., and Scheuermann, K. (2025). The future of learning or the future of dividing? exploring the impact of generative artificial intelligence on higher education. *Data & Policy*, 7:e44.

World Economic Forum (2025). The future of jobs report 2025.

A Appendix

A.1 AI Classification Dictionaries

This appendix presents the complete keyword dictionaries used to classify curricular units into Direct AI and AI Related categories, as described in Section 3.2. Keywords are stored and matched after text normalization (lowercase conversion, accent removal, and special character elimination). Both English and Portuguese terms are included to ensure comprehensive matching across bilingual course names.

A.1.1 Direct AI Keywords

The Direct AI dictionary comprises eight thematic categories encompassing core artificial intelligence topics. A curricular unit is classified as Direct AI if its name contains at least one keyword from any of these categories.

- **Machine Learning:** machine learning, deep learning, neural network, supervised learning, unsupervised learning, reinforcement learning, aprendizagem automatica, aprendizagem de maquina, aprendizagem maquina, aprendizagem profunda, redes neuronais, redes neurais, rede neuronal, aprendizagem supervisionada, aprendizagem nao supervisionada
- **AI Fundamentals:** artificial intelligence, intelligent agent, search algorithm, heuristic, expert system, inteligencia artificial, agente inteligente, heuristica, sistema especialista, sistemas especialistas, algoritmos de busca
- **Natural Language Processing:** natural language processing, NLP, text mining, sentiment analysis, language model, chatbot, dialogue system, processamento de linguagem natural, PLN, analise de sentimento, modelo de linguagem
- **Computer Vision:** computer vision, image recognition, object detection, image segmentation, face recognition, video analysis, pattern recognition, visao computacional, reconhecimento de imagem, detecao de objetos, reconhecimento facial, analise de video
- **Robotics:** robotics, robot, autonomous system, robotic vision, motion planning, robotica, robo, sistema autonomo

- **AI Ethics and Governance:** AI ethics, trustworthy AI, explainable AI, XAI, AI fairness, AI governance, responsible AI, etica em IA, IA confiavel, IA responsavel
- **Generative AI:** generative AI, generative adversarial, diffusion model, variational autoencoder, VAE, large language model, LLM, GPT, transformer, IA generativa
- **Probabilistic AI:** Bayesian network, Markov, probabilistic reasoning, decision theory, rede bayesiana, teoria de decisao

A.1.2 AI Related Keywords

The AI Related dictionary comprises nine thematic categories covering foundational and enabling topics for AI education. A curricular unit is classified as AI Related if its name contains at least one keyword from any of these categories but no keywords from the Direct AI dictionary.

- **Data Science:** data science, data mining, big data, data analytics, business intelligence, ciencia de dados, mineracao de dados, analise de dados
- **Statistics and Probability:** statistics, probability, statistical analysis, hypothesis test, stochastic, estatistica, probabilidade, analise estatistica, teste de hipotese, estocastico
- **Linear Algebra:** linear algebra, matrix, vector, eigenvalue, algebra linear, matriz, vetor, autovalor
- **Algorithms and Data Structures:** algorithm, data structure, complexity, optimization, graph theory, algoritmo, estrutura de dados, complexidade, otimizacao, teoria dos grafos
- **Programming:** Python, R programming, programming, software development, coding, programacao, desenvolvimento de software, codificacao
- **Data Management:** database, data management, SQL, NoSQL, data warehouse, base de dados, gestao de dados, armazem de dados
- **Computing:** parallel computing, distributed computing, GPU computing, CUDA, cloud computing, computacao paralela, computacao distribuida

- **Image Processing:** image processing, signal processing, digital signal, processamento de imagem, processamento de sinal, sinal digital
- **Data Visualization:** data visualization, graphics, visualization, visualizacao de dados, graficos, visualizacao

Note: All keywords shown are the normalized forms used for matching after preprocessing (lowercase conversion, accent removal and special character elimination). Keywords are matched using substring matching. The hierarchical classification prioritizes Direct AI over AI Related when a course name contains keywords from both dictionaries.

A.2 Panel Data Sample Characteristics

Table A1: Sample Composition by **Degree Level**

Level	N	Percentage (%)
Bachelor	1310	39.1
Master	1963	58.6
Integrated Master	79	2.4
Total	3352	100

Table A2: Sample Composition by **Institutional Sector**

Sector	N	Percentage (%)
Public	2567	76.6
Private	782	23.4
Total	3349	100

Table A3: Sample Composition by **Education Type**

Type	N	Percentage (%)
University	2081	62.1
Polytechnic	1270	37.9
Total	3351	100

A.3 AI Specific Programs

Table A4: Distribution of AI-Named Degree Programs by Level and Year

Year	Bachelor's	Master's	Total
2020	0	1	1
2021	1	2	3
2022	2	0	2
2023	2	2	4
2024	1	5	6
2025	0	3	3
Total	**6**	**13**	**19**

Table A5: Complete List of AI-Named Degree Programs

Program Name	Training Area	Level	Year
Engenharia de Inteligência Artificial	ENG	Master	2020
Computação Criativa e Inteligência Artificial	STM	Master	2021
Inteligência Artificial e Ciência de Dados	STM	Bachelor	2021
Inteligência Artificial Aplicada	STM	Master	2021
Tecnologias Digitais e Inteligência Artificial	STM	Bachelor	2022
Ciência de Dados e Inteligência Artificial	STM	Bachelor	2022
Inteligência Artificial e Ciência de Dados	STM	Bachelor	2023
Inteligência Artificial	STM	Master	2023
Inteligência Artificial e Ciência de Dados	STM	Bachelor	2023
Inteligência Artificial para Jogos	STM	Master	2023
Inteligência Artificial	STM	Master	2024
Inteligência Artificial e Ciência de Dados	STM	Master	2024
Inteligência Artificial para Sociedades Sustentáveis	STM	Master	2024
Inteligência Artificial	STM	Master	2024
Inteligência Artificial	STM	Master	2024

Continued on next page

Table A5 – *Continued from previous page*

Program Name	Training Area	Level	Year
Biomedicina Computacional e Inteligência Artificial	ENG	Bachelor	2024
Inteligência Artificial Aplicada ao Negócio	STM	Master	2025
Inteligência Artificial Aplicada e Computação	STM	Master	2025
Imagem Médica Avançada e Inteligência Artificial	HW	Master	2025

Training areas abbreviated as:

STM = Science, Technology, and Mathematics

ENG = Engineering, manufacturing, and construction

HW = Health and welfare.

A.4 Figures

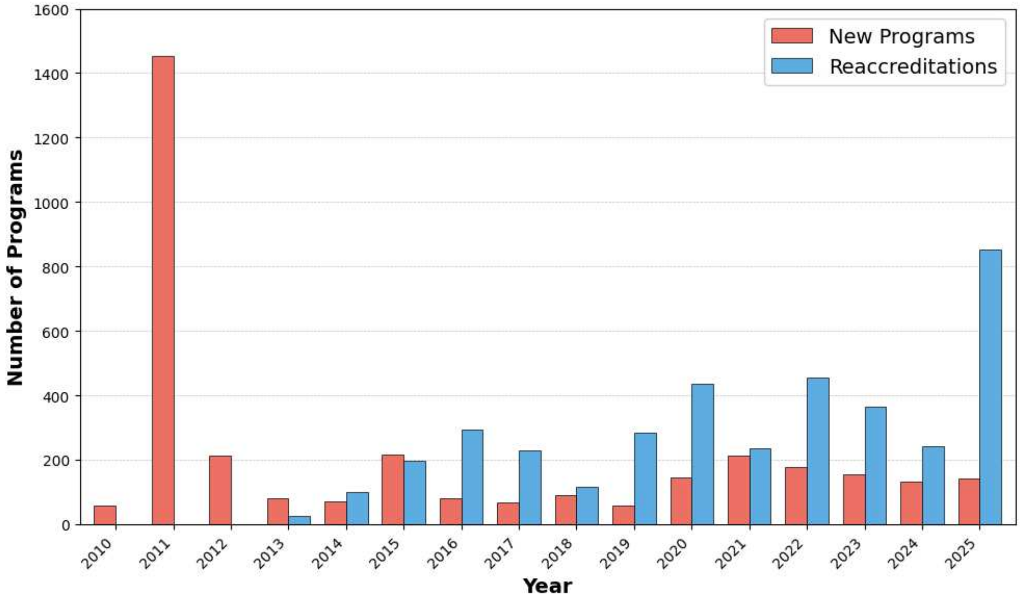

Figure A1: Number of Reaccreditations and New Programs by Year

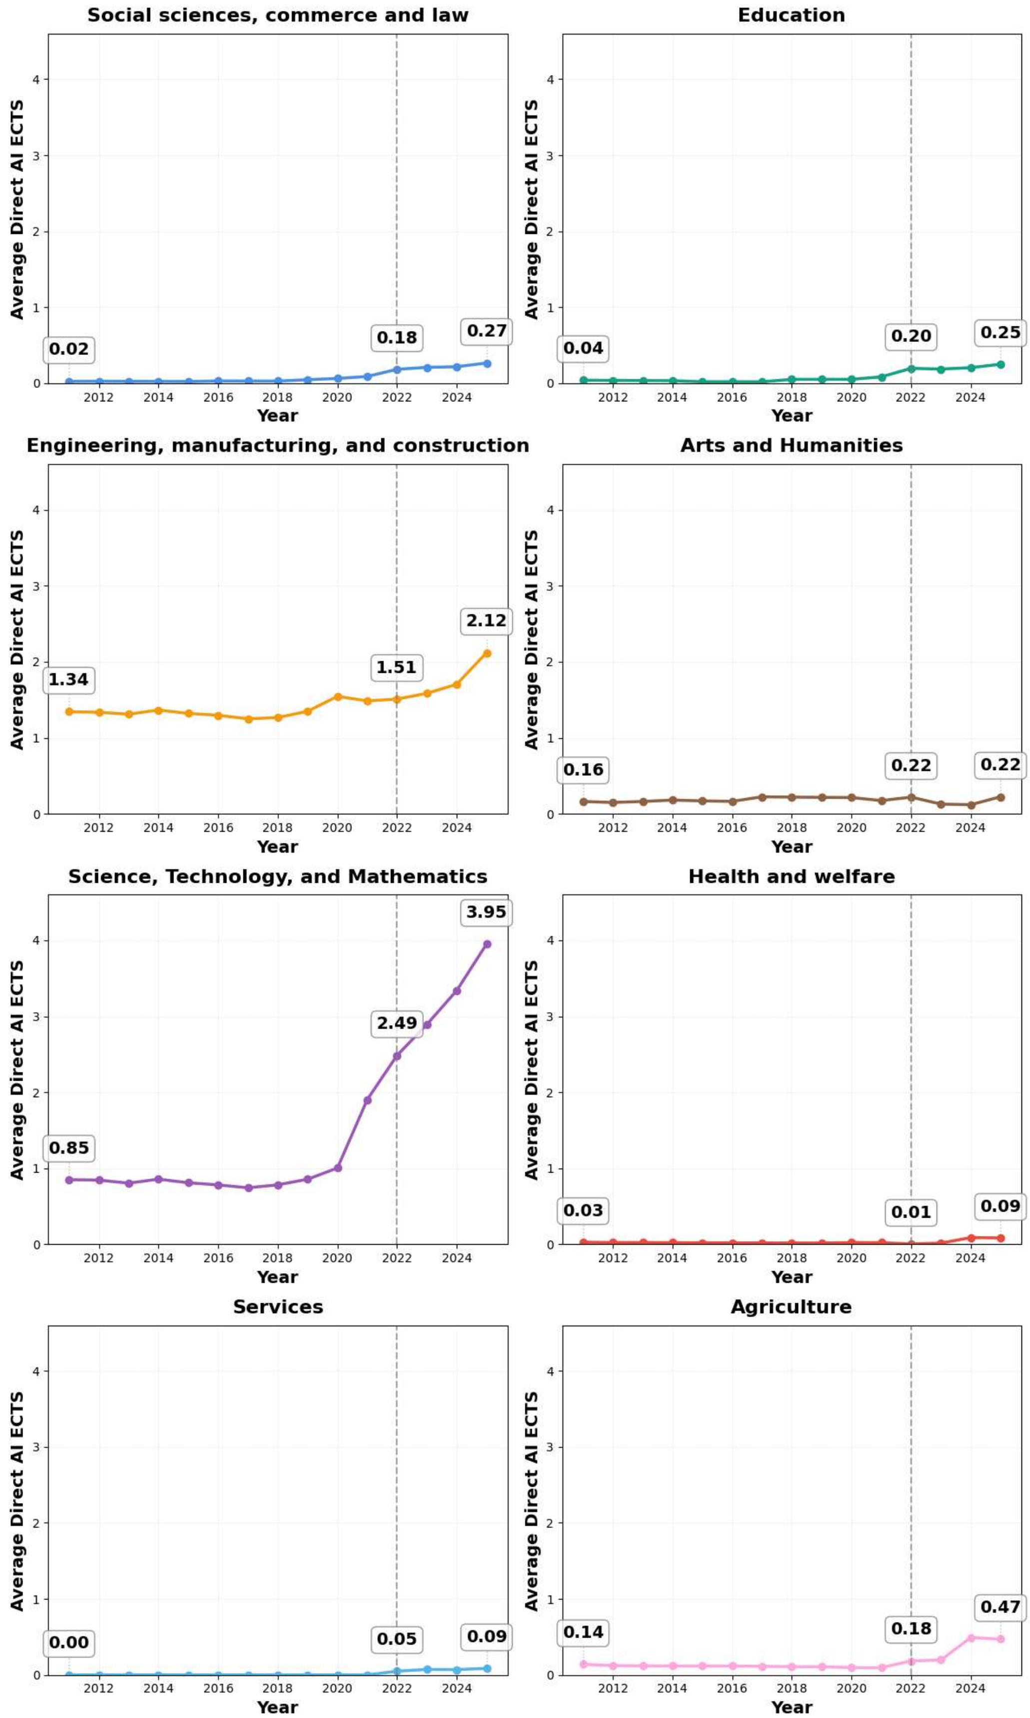

Figure A2: Average Direct AI ECTS by Training Area Over Time

A.5 Heterogeneous Effects by Training Area

Table A6: Heterogeneous Effects of Reaccreditation by Training Area (Triple Interaction)

	AI Adoption (*has_any_ai*)	
	Logit	LPM
	(1)	(2)
Reaccredited × After × Training Area:		
Agriculture	22.026***	0.052
	(0.816)	(0.047)
Arts and Humanities	−18.957***	−0.002
	(0.521)	(0.005)
Education	37.077***	0.015
	(0.869)	(0.008)
Engineering, manufacturing, construction	5.846***	0.105***
	(1.611)	(0.027)
Health and welfare	22.026***	0.009
	(0.816)	(0.005)
Science, Technology, Mathematics	4.540**	0.076*
	(1.741)	(0.035)
Services	22.026***	0.009
	(0.816)	(0.007)
Social sciences, commerce, law	37.226***	0.016**
	(1.022)	(0.006)
Observations	961	34,313
R^2	0.569	0.937
Adjusted R^2	0.411	0.930
Program FE	Yes	Yes
Year FE	Yes	Yes

Notes: Standard errors are clustered at the program level.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.