# Exploring Barriers to AI Doctoral Studies

**Reviewed by**

Professor Flavia De Luca
(and Dr Emily Wride?)

**Authors**

Jana Alloush, Erin A. F. Brady, Rosalba Castiglione,
Weilin He, Sven Hollowell and Adrianna Jezierska

**Date of Submission:** 31-03-2026

# Table of Contents

# Abstract

This study investigates potential barriers to doctoral studies in Artificial Intelligence (AI) and evaluates the extent to which an AI training programme, delivered by the Jean Golding Institute at the University of Bristol (UK), can mitigate them. The literature review identified three key domains shaping access to AI doctoral pathways: socio-economic inequalities, AI literacy, and institutional policy structures. The research was conducted between October 2025 and March 2026 by an interdisciplinary team of six researchers under the supervision of Professor Flavia De Luca, and examines whether such training can influence students' AI literacy, perceptions, and interest in postgraduate AI education.

Using a mixed-methods approach, data were collected through pre- and post-training surveys (n=13), qualitative interviews (n=5), and an automated analysis of PhD job listings in the UK. The findings show enhanced AI literacy following the training. However, no significant change was observed in participants' interest in pursuing postgraduate study in AI. Qualitative interviews indicate that participants primarily perceived AI as a tool to enhance employability within their current disciplines, rather than as a pathway to doctoral studies. This may reflect a gap between AI-related knowledge/skills acquisition and actual educational aspirations, possibly shaped by real and/or perceived relevance and general uncertainty about AI-related career pathways.

Given the short timeframe and small sample size, the findings are exploratory. Nevertheless, the study hopes to contribute to ongoing, timely discussions on equality, equity, diversity, and inclusion in AI education, highlighting that improving AI literacy alone may not be sufficient in order to enhance access to AI doctoral studies. Further research with larger and more diverse samples is needed to better understand and address the above barriers.

# 1. Introduction

The present review aims to identify potential barriers to PhD/Doctoral studies on AI (Artificial Intelligence) in Bristol and the United Kingdom. To achieve this, a team of researchers from the University of Bristol has conducted a literature review of relevant studies across disciplines. Following brainstorming (Figure 1), three main areas of focus were identified, that is, socio-economic, AI-literacy, and AI-policy related barriers.

To understand why the doctoral pipeline remains exclusive despite high demand, this report examines barriers across three interconnected dimensions. First, we examine *Socio-Economic Barriers* (Section 2), exploring how systemic inequalities in gender, race, and finance filter out talent before they apply. Second, we analyse *Individual Barriers* (Section 3), specifically the role of 'AI Literacy' and self-efficacy as psychological gatekeepers. Finally, we investigate *Structural Barriers* (Section 4), using a primary analysis of 2,643 live PhD listings to reveal how institutional policy and market mechanisms actively reinforce these exclusions.


*Figure 1. Brainstorming process for identifying barriers to AI PhD/Doctoral studies.*

Before delving into the literature review, it is important to briefly illustrate how 'AI' is understood as part of this work. Empirical definitions of AI remain scarce and inconsistent, which can affect both research findings and cross-study comparisons (Kelly et al., 2023). Nevertheless, the most widely used definition conceptualises AI as computer systems designed to perform tasks that typically require human intelligence, such as learning, problem-solving, and reasoning (MIOIR, 2025). Consistent with the above definition, some scholars suggest that AI can also be used to understanding natural language, recognising patterns, and making decisions (Kalim et al., 2025; Patole et al., 2024). For the purpose of the present review, it is also relevant to understand what it is meant by AI acceptance, or else, the intention or willingness to use, buy, or try AI systems (Kelly et al., 2023). AI is generally used to process existing data (normally referred to as processing AI) or to create new content (generative AI) (Banh and Strobel,2023; Ulnicane, 2025). While alternative socio-technical and critical definitions of AI may exist, the above functional definitions seem to underpin the literature reviewed here and reflect dominant usage of AI across education and policy research. In addition, this study adopts Long and Magerko's definition of AI literacy as the set of competencies needed to use and interact with AI tools in different professional and personal settings (2020).

To ground this review in the UK's current doctoral landscape, we complemented our literature analysis with a primary study of active recruitment data. Using a custom NLP pipeline on 2,643 listings from FindAPhD.com (), we mapped patterns in funding, eligibility, and entry requirements. This empirical approach tests the theoretical barriers identified in our brainstorming (Figure 1), revealing exactly how abstract policy issues manifest as concrete obstacles in the recruitment of future AI researchers.

# 2. Literature Review

## 2.1 Socio-Economic Barriers to Doctoral Studies In AI

This section examines barriers to PhD studies in Artificial Intelligence (AI), with a particular focus on socio-economic factors that may hinder access to and progression within doctoral education in AI. It also explores the bidirectional relationship between socio-economic inequalities in AI education and broader inequalities within the AI ecosystem, including those related to governance, labour markets, and research production. Specifically, inequalities in access to AI doctoral studies shape who produces AI research and who participates in AI-related decision-making, while existing power imbalances in AI research and industry, in turn, influence educational opportunities.

Drawing on a review of relevant literature, the chapter identifies key barriers to engagement with AI in doctoral contexts. It first discusses broad barriers to AI adoption among doctoral students before turning to specific socio-economic dimensions shaping engagement with AI, including gender, geography, and race and ethnicity. The chapter concludes with a summary of findings and a set of policy- and research-oriented recommendations.

The literature search was conducted with support from a University of Bristol subject librarian, who advised on search strategies, platforms, and terminology. Searches were carried out using Web of Science, Scopus, ERIC, Google Scholar, and Research Rabbit. Key search terms included socio-economic, poverty, financial inequality, doctoral, PhD, postgraduate, gender inequality, artificial intelligence, AI, ChatGPT, and Copilot.

The review revealed a substantial gap in studies directly examining socio-economic barriers to AI studies in higher education, particularly at the doctoral level, a limitation also noted by other scholars (e.g., Kalim et al., 2025). While a growing body of research addresses undergraduate and master's-level participation in AI, doctoral education remains comparatively underexplored. To address this gap, the review included studies on closely related topics (as revealed by the existing literature), such as STEM (Science, Technology, Engineering, Mathematics) education, AI-related authorship, and workforce participation. In addition to peer-reviewed academic articles, relevant reports and scholarly blogs were also included.

Most identified studies focused on skills gaps in relation to labour-market demands and proposed strategies to mitigate their impact. The literature was heavily dominated by studies conducted in the United States (e.g. Ferguson et al., 2023), with fewer contributions from the UK and other global contexts (e.g. Ulnicane, 2025). This geographic imbalance reflects broader concentrations of research funding, data availability, and publication practices, and limits the generalisability of findings across contexts.

Socio-economic barriers were most frequently discussed in relation to gender, typically operationalised as men and women, followed by race and or ethnicity. Other dimensions were often, if not always, overlooked, including age (e.g., mature students), disability, neurodiversity, migration status, or LGBTQ+ identities. This points to a significant gap in the current literature on Equity, Diversity, and Inclusion (EDI) in AI-related education. Consistent with these observations, a cross-cutting theme across the reviewed literature is the need for accurate data and robust methodologies to better understand EDI dynamics in AI and STEM higher education.

### 2.1.1. Socio-Economic Barriers to AI Adoption

One of the few studies explicitly examining AI adoption (as a tool to support their research) among doctoral students is that of Oliveira et al. (2024). The above study was conducted at the University of Texas at Tyler. Although context-specific, the study offers valuable insights given the broader lack of research in this area. While doctoral students generally recognised the efficiency and usefulness of AI tools, ethical concerns, particularly those related to academic integrity, ethical standards, and the unknown long-term impacts of AI, emerged as major barriers to adoption.

Technological dependence was another key concern. Participants expressed ambivalence towards AI's transformative potential. Although they acknowledged its utility, some reported negative emotional responses associated with a perceived loss of agency and autonomy in academic work. Scepticism also stemmed from the perception that AI-generated outputs lack genuine human thinking and articulation, raising concerns about originality and scholarly voice. While the authors do not further explore sociodemographic characteristics, they highlight how different demographic groups might be differently impacted by the academic culture around AI studies with unrepresented groups potentially being more sceptical about it due to culture (e.g., Kalim et al., 2025; Kelly, et al., 2023) and literacy.

The following sections examine how these general barriers intersect with specific socio-economic dimensions, beginning with gender.

### 2.1.2. Gender And Its Intersections with Other Socio-Cultural Factors

Gender inequality represents one of the most significant barriers to participation in AI doctoral studies. While gender disparities in education have long been documented (Van de Werfhorst, 2017), the growing integration of AI into academia may be reinforcing existing inequalities rather than mitigating them (Stanford University, 2023).

Evidence from multiple studies, reports, and institutional analyses (e.g. Frachtenberg and Kaner, 2022; Kalim et al., 2025; MIOIR, 2025; Stathoulopoulos et al., 2019) consistently shows that white men are significantly overrepresented in AI-related doctoral studies, conference presentations, and academic authorship, all of which are closely tied to early academic careers in AI. Although overall participation in AI PhD research increased between 2000 and 2020, STEM doctoral theses remain predominantly authored by men, particularly in AI-related fields (MIOIR, 2025).

One potential explanation lies in the gender composition of supervisory staff. Approximately 73.4% of AI PhD supervisors are male (based on UK data; MIOIR, 2025), a figure that reflects broader patterns within computer science and engineering. This imbalance may discourage women from pursuing AI doctorates and limit their access to professional networks, mentorship, and career opportunities (MIOIR, 2025). While the exact relationship between the gender composition of supervisory staff and that of doctoral students is not clear, it could be argue that male prevalence within the field may build the perception that AI is predominantly a male discipline.

AI is a specialised field but is embedded within broader STEM disciplines such as computer science, engineering, and mathematics (MIOIR, 2025). Patterns of gender segregation in STEM therefore provide valuable insight into women's underrepresentation in AI doctoral studies. Structural and cultural barriers identified in STEM appear to carry over into AI education and research. Van de Werfhorst (2017) provides empirical evidence of persistent gender segregation across post-secondary fields of study, particularly among students from lower socio-economic backgrounds. This highlights the intersection of gender and socio-economic status, suggesting that women from disadvantaged backgrounds are especially unlikely to pursue AI PhDs.

Additional factors further contribute to women's underrepresentation in AI. A systematic review by Kalim et al. (2025), focusing on Asian higher education contexts, found that women often report lower familiarity with AI tools, largely due to limited access to training, institutional support, and AI literacy programmes. Reduced familiarity, in turn, lowers AI adoption rates. Traditional gender roles and stereotypes also discourage women from entering AI-related studies and careers, as women are often expected to pursue occupations centred on care, communication, and human interaction (Kalim et al., 2025).

Gender bias embedded within AI algorithms themselves may further reinforce exclusion, particularly when biased systems shape educational assessment, recruitment, or research visibility. Moreover, women are frequently underrepresented in AI policy frameworks, limiting their influence on decision-making and the development of inclusive AI practices (Kalim et al., 2025). Women also tend to report greater concerns about privacy, data security, and the ethical implications of AI use (Kalim et al., 2025). These concerns, often linked to limited familiarity with AI tools, may further reduce engagement with AI technologies.

### 2.1.3. Authorship

Consistent with the patterns described above, studies examining AI-related authorship reveal substantial gender disparities (Ding et al., 2025; Frachtenberg and Kaner, 2022; Stathoulopoulos et al., 2019). Frachtenberg and Kaner (2022), using the female author ratio as a proxy measure, found that women accounted for only 10.26% of authors in leading AI-related systems conferences worldwide.

Similarly, Ding et al. (2025), drawing on the AI Scholar dataset, showed that women represented approximately 14 to 17% of AI researchers between 1980 and 2022. Although this proportion has increased slightly over time, the gender gap remains persistent. Male-authored publications also tend to receive more citations, partly due to men's greater participation in larger, male-dominated collaboration networks and citation practices.

It is important to note that these studies typically rely on binary gender classification based on names, which excludes non-binary researchers and may underestimate inequality. Furthermore, the "leaky pipeline" phenomenon, whereby women disproportionately exit academic careers at early stages, likely contributes to their continued underrepresentation in AI publications (Ding et al., 2025).

### 2.1.4. Location and Its Intersection with Gender

Geography plays a major role in shaping access to AI education, research, and careers. A study analysing AI professionals' LinkedIn profiles found that half of the sample was concentrated in just four countries, the USA, France, Germany, and the UK, while more than 50 countries were each represented by less than 1% (Young et al., 2023). This finding highlights a strong geographic concentration within the AI labour market.

Geographic disparities become particularly pronounced when intersecting with gender. For example, approximately 30% of AI papers in the Netherlands included at least one female co-author, compared with only 10 to 16% in Japan and Singapore. Countries such as Malaysia, Denmark, Norway, and Israel demonstrate relatively stronger female representation in AI research (Stathoulopoulos et al., 2019).

Within AI, certain subfields, such as machine learning and data science, show especially low probabilities of female authorship. Notably, research co-authored by women tends to focus more on applied and socially oriented topics, including fairness, health, mobility, and gender. This pattern reflects, and may also reinforce, the traditional gendered expectations discussed earlier.

### 2.1.5. Race And Ethnicity

Underrepresentation typically refers to groups that have experienced historical discrimination (Kuhlman et al., 2020). Across contexts, individuals from Black, Asian, and Latinx backgrounds face reduced access to AI and STEM education and employment (Roberts et al., 2021; Ferguson et al., 2023). A report by the UK Council for Graduate Education (UKCGE, 2024) found that 74% of postgraduate research students in the UK were white, compared with 5% Black and 9% Asian students. Economic barriers were identified as a major factor affecting completion and progression among ethnic minority students (Lindner, 2020).

Longitudinal research in STEM further shows that Black students are less likely to graduate with top degree classifications (Low and Kalender, 2024). While representation of ethnic minorities in computer science has improved over time, white men continue to dominate undergraduate, master's, and doctoral levels (Stanford University, 2023). Significant inequalities persist, particularly at the intersections of race, gender, poverty, and disability.

### 2.1.6. The Socio-Economic Inequalities Loop In AI

This paragraph attempts to explain the broader inequalities in the AI ecosystem how they might influence decisions related to future studies in AI and job market.

#### Governance, Politics, and Geography

Barriers to AI doctoral education reflect broader inequalities within the AI ecosystem (Kuhlman et al., 2020). A systematic review by Shams et al. (2025) highlights the underrepresentation of the Global South in AI research, governance, and labour markets. AI governance is shaped by political and economic power, with decision-making dominated by actors in the Global North (Vargas-Solar, 2025; Young et al., 2023). A limited number of countries, alongside a small group of large technology companies, appear to control much of the generative AI landscape (see study introduction for definition).

Democratic and participatory co-shaping of technology is therefore essential to address the unequal distribution of power in generative AI. Governments play a crucial role in ensuring equal access to AI through education policy, funding, regulation, and public infrastructure, and in addressing power imbalances both within societies and across socio-ecological levels (Ulnicane, 2025).

The aforementioned factors contribute to algorithmic bias reinforcing further inequalities, hence barriers. Although often framed as a technical problem requiring technical solutions, scholars argue that educational diversity and inclusion offer a more sustainable and just approach to addressing global systemic AI inequalities (Shams et al., 2025).

#### Job Market Inequalities and Education

Decisions to pursue AI careers are shaped by confidence in technical and interpersonal skills, perceptions of social impact, and professional identity (Ferguson et al., 2023). Women often report lower confidence in their abilities despite holding higher formal qualifications, a pattern widely observed across STEM and Information and Communication Technology (ICT) fields (Young et al., 2023). They also face stereotyping, lower pay, limited career advancement, and greater exposure to hostile work environments and microaggressions, all of which contribute to higher attrition rates (Ferguson et al., 2023; Young et al., 2023).

Evidence from STEM and ICT fields suggests that negative, gendered workplace cultures, excessive competition, and a lack of visible role models discourage women and minority groups from pursuing educational pathways that lead to AI careers ( Harehdasht and Saki, 2025; Williams et al., 2025).

The thematic mapping below (Figure 2) visualises the findings of the above literature, demonstrating how the different socio-economic barriers interrelate to form a ‘loop’ that hinders access to AI doctoral studies.

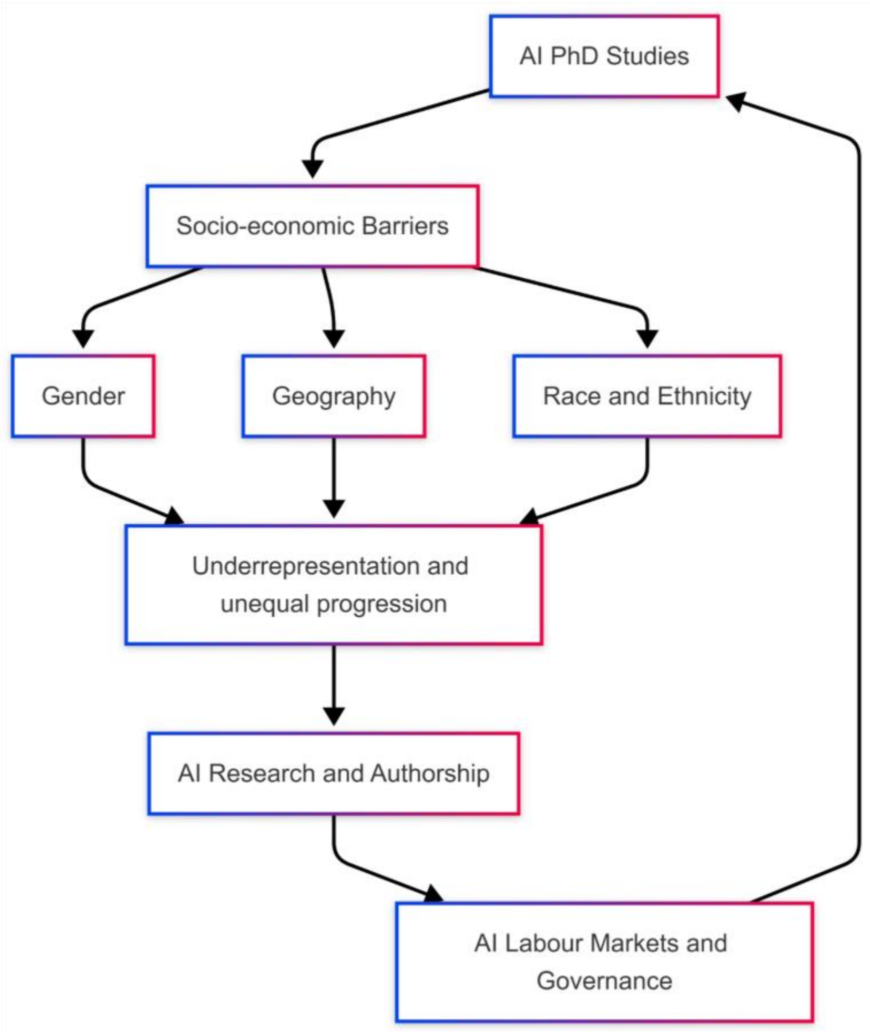


*Figure 2. A thematic mapping of socio-economics barriers interrelations*.

### 2.1.7. Conclusion and Recommendations

To reduce socio-economic barriers to AI doctoral studies, universities in the UK and the USA have proposed similar strategies, including diversifying admissions panels, offering bridging programmes to address skills gaps, and implementing structured mentorship schemes (Roberts et al., 2021). In countries such as Finland and Sweden, higher female participation in AI has been linked to targeted funding, gender quotas, and participation-focused initiatives (Vargas-Solar, 2025).

There is an urgent need to expand EDI-focused research in AI doctoral education to develop evidence-based, context-sensitive strategies. Educators integrating AI tools into doctoral training must not only promote technical proficiency but also actively challenge exclusionary academic cultures and stereotypes. Clear ethical guidelines, training that complements rather than replaces human skills, and strong support systems can help alleviate concerns related to loss of agency and academic integrity.

Research in this area should also examine the psychological and behavioural dimensions of AI adoption, including its impact on academic identity and professional trajectories. More empirical evidence is needed to assess AI's effects on learning outcomes, research quality, and critical thinking, as well as the risks associated with over-reliance on AI tools.

Given the links between secondary education pathways and postgraduate AI participation, scholars advocate for gender-sensitive career guidance, targeted mentoring for students from lower socio-economic backgrounds, and policies encouraging gender balance across technical and care-related fields. Systematic monitoring of gender and socio-economic segregation is essential to evaluate policy effectiveness over time.

Finally, although inequalities in access to AI doctoral studies can be highly context-specific, the use of global literature does not necessarily constitute a limitation in this chapter. UK higher education is inherently international and multicultural, with doctoral cohorts often comprising students and staff from diverse geographic, cultural, and socio-economic backgrounds (Universities UK International, 2024).

Drawing on studies conducted outside the UK, including research from Asia (e.g. Kalim et al., 2025), therefore enhances rather than limits the relevance of this review when examining barriers to AI doctoral studies. At the same time, while socio-economic barriers may share common structural roots across countries, context-specific evidence remains essential, particularly when informing policy recommendations within the UK higher education system (Vargas-Solar, 2025).

## 2.2. AI Literacy

AI literacy is defined by Long and Magerko (2020) as:

> "a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace".

This is currently the most widely accepted definition in academic literature (Zhang, Prasad and Schroeder, 2024). It emerged from the concepts of digital literacy[1], information literacy and media literacy[2], emphasising ethical knowledge, critical thinking and interpretive skills (Chen et al., 2022; Lu and Lin, 2025). AI literacy has a proven positive effect on career intention and interest (D. Zhang et al., 2025), indicating that it could be used to examine prospective candidates' interest and aptness for doctoral studies in AI.

Although the link between AI literacy and career intention is not always made explicit, the importance of AI literacy for education and society as a whole is mentioned in grey literature globally (Elhussein et al., 2024; Mendigutxia, 2024; The Digital Cooperation Organization, 2025; New Jersey Department of Education, 2025; DSIT, 2025). Academic literature also considers the links between other recent constructs like AI self-efficacy[3] (Chen et al., 2024; Duong, 2025; Lu and Lin, 2025; D. Zhang et al., 2025) and AI capital[4] (Drydakis, 2024; 2025). This is useful for exploring the literature since, due to the recency of AI literacy's development, it may not have reached ubiquity – closely related concepts have been explored as well.

Most conceptualisations arise from the field of computer science or digital education (Zhang, Prasad and Schroeder, 2024). Long and Magerko (2020) presented a list of 17 competencies for AI literacy, which will be discussed at greater length in Section 3, but it is worth noting that among these are knowledge of "Programmability", "Sensors" and "Data Literacy", similarly suggesting a computing focus. Due to the emphasis on AI literacy for computing, the following sections will also explore the relevance of their popular definition in assessing the suitability of candidates for doctoral level AI programmes in the UK.

### 2.2.1. AI Literacy Development

AI literacy does not develop automatically. A series of institutional and psychological filters govern its trajectory; these filters determine who persists or leaves the learning pipeline.

---

[1] Defined at length by Bawden (2008) as a set of understandings, skills, and attitudes enabling individuals to locate, understand, create, evaluate, and communicate information using digital technologies.

[2] UNESCO (*Information Literacy,* n.d.) states that "Media and Information Literacy provides a set of essential skills to address the challenges of the 21st century including the proliferation of mis- and disinformation and hate speech, the decline of trust in media and digital innovations notably (AI)".

[3] AI self-efficacy is "Learners' competence to complete AI practices independently" (Chen et al., 2024).

Institutional structures limit who can progress from basic AI exposure to advanced AI research through technical entry requirements and rigid prerequisites. Among those who enter AI learning pathways, learners decide whether to continue based on two considerations. They judge whether AI is worth pursuing in relation to their goals, and they assess whether they believe they can succeed. These judgments of value and self-efficacy shape decisions to persist or withdraw. Understanding these interactions reveals the specific points where the journey toward expertise stops.

### 2.2.2 Institutional Access and Technical Gatekeeping

Although definitions of AI literacy do not require learners to possess programming skills or a prior computing background, doctoral study in AI globally continues to carry implicit expectations of technical preparedness.

In the UK, EPSRC-funded Centres for Doctoral Training (CDTs) have broadened access to AI doctoral study to some extent by recruiting applicants without formal backgrounds in AI or computer science. However, research shows that students from non-computing backgrounds and other underrepresented groups begin to experience sustained attrition at much earlier stages of the educational trajectory, giving rise to a pronounced "leaky pipeline"[5] effect well before doctoral training begins (Sarraju et al., 2023). At the undergraduate level, beyond the formal admission requirements, AI courses often require extended sequences of technical and mathematical training, which delay students' initial exposure to AI and reinforce entry barriers for non-technical learners (Niousha et al., 2024; Xia et al., 2024).

In sum, institutional access and technical prerequisites act as the initial gatekeepers that define who can begin the journey toward advanced AI research.

### 2.2.3. Value Perception as a Motivational Filter

Once learners pass the initial institutional gate, value perception operates as a motivational filter that shapes whether learners remain engaged long enough for AI literacy to develop.

Research based on Expectancy–Value Theory and the Theory of Planned Behaviour shows that learners' engagement with AI is driven more by perceived utility and relevance than by technical exposure alone (Yin and Goh, 2024). When AI is presented primarily as a mathematical or engineering discipline, this framing can weaken the formation of professional identity for students seeking socially oriented roles (Chun and Elkins, 2023). For example, students from underrepresented groups (URGs) often seek socially oriented applications of AI, such as social justice, healthcare, or law. When AI is framed mainly as a technical field, computing can appear socially irrelevant, leading to disengagement (Barretto et al., 2021).

### 2.2.4. AI Self-Efficacy as a Psychological Boundary

In addition to institutional access and motivational factors, learners' confidence in their own ability plays an important role in shaping continued engagement with AI learning.

Many novice learners report early intimidation, as they think AI literacy is often perceived as closely associated with programming expertise. AI is frequently perceived through stereotypes that portray the field as highly technical, competitive, and socially isolated, and as accessible mainly to those with prior experience or particular identity characteristics (e.g., White, Asian, man, genius) (Barker et al., 2002; Beyer, 2014; Cheryan et al., 2013; Lewis, Anderson and Yasuhara, 2016). Cowit and Fiesler (2024) show that such perceptions are associated with persistent uncertainty and self-doubt. Lower self-efficacy is, in turn, linked to earlier disengagement, often before sustained engagement with AI learning occurs. These confidence-related challenges are particularly evident among URGs and women. For URGs, limited exposure to professional role models and a lack of diverse representation in mainstream media can reduce familiarity with the field and increase self-doubt at entry. In addition, the dominant image of the "tech geek" weakens female students' sense of belonging and interest in computing, contributing to higher attrition rates even among academically high-performing women (Cheryan et al., 2013).

Taken together, these institutional, motivational, and psychological barriers suggest that unequal progression toward advanced AI literacy is not a matter of individual ability, but of structured access and support.

### 2.2.5. Intervention Design

AI literacy has a positive effect on career intention and interest (D. Zhang et al., 2025); the positive influence of AI literacy on digital entrepreneurship (Duong, 2025) and software development (Chen et al., 2024) have been measured in exploratory studies of university students in Viet Nam and Taiwan respectively.

Evaluative literature in HE or pre-HE contexts tends to focus either on improving career interest, or on improving participants' AI literacy itself. To understand how AI literacy can be developed in participants and how interest in applying for doctoral studies in AI increases, interventions related to both parts are useful.

### 2.2.6. Enhancing AI Career Interest

In an umbrella review, Zhang, Prasad and Schroeder (2024) identified that only one of 17 reviews looked at HE and adult education. Compared with exploratory research, there is little documentation on evaluative studies with HE learners, which would drive a need to increase intakes of university students entering undergraduate AI courses rather than at postgraduate level, as indicated in global reports on AI education (Mendigutxia, 2024; Maslej, 2025).

In studies aiming to improve career prospects or uptake of AI-related study, AI literacy is rarely named as an outcome. Nevertheless, some such interventions implicitly target AI literacy competencies, so the intervention's relevance can be assessed by mapping the reported activities onto Long and Magerko's (2020) 17 AI literacy competencies and 14 design considerations (listed in Figures 3a and 3b):

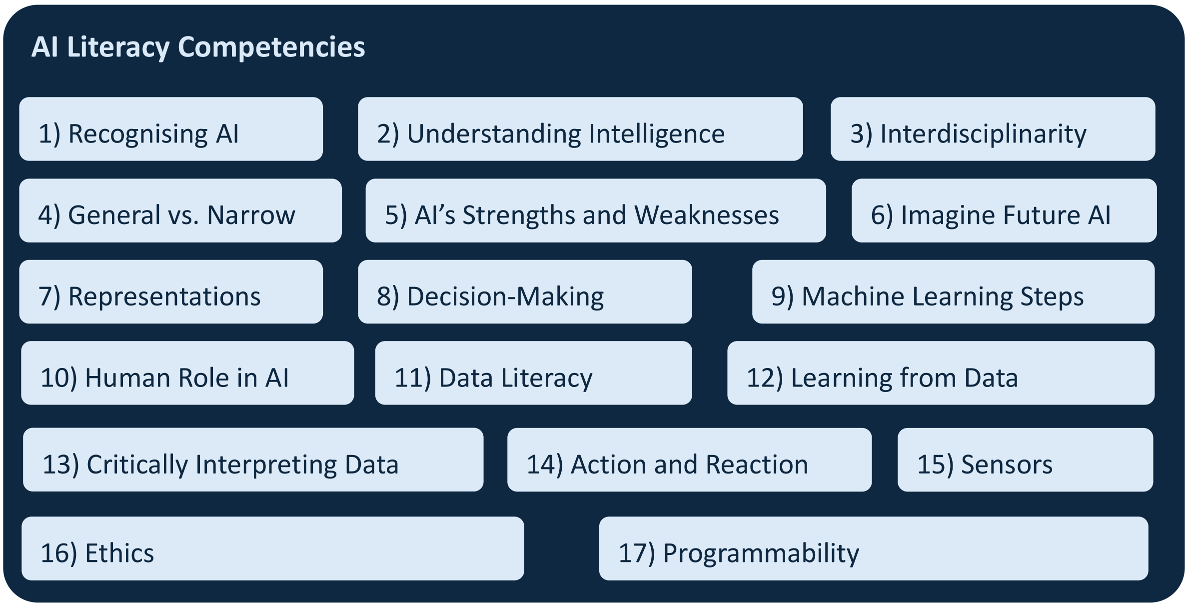


*Figure 3a. The 17 competencies of AI literacy as described by Long and Magerko (2020).*

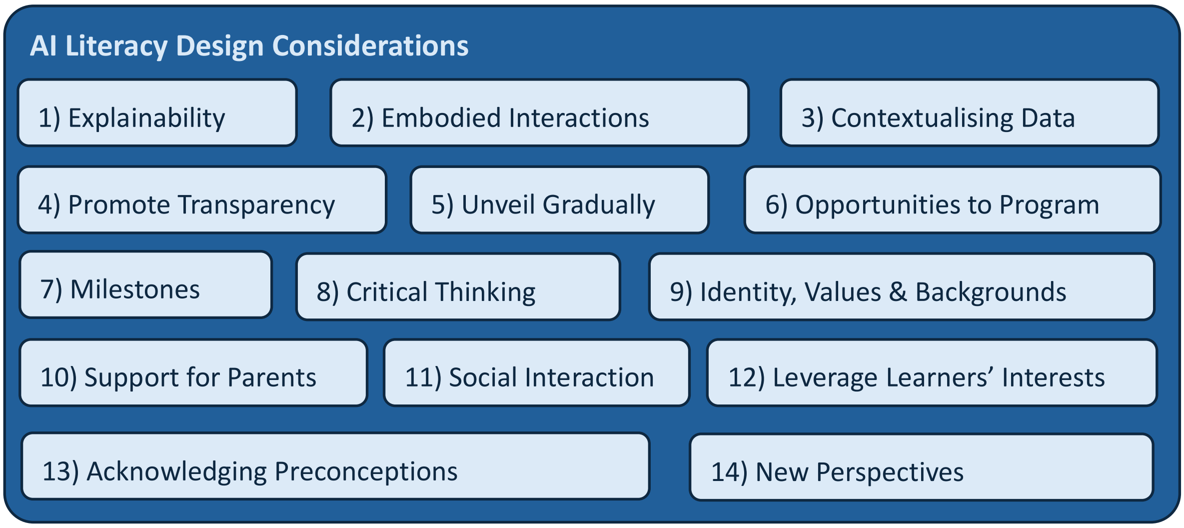


*Figure 3b. The 14 design considerations for activities to improve AI literacy as described by Long and Magerko (2020).*

Oskotsky et al. (2022) report on an introductory AI programme for female pre-biomedicine students to promote diversity in AI careers. It included discipline-relevant programming projects, lecture content, and community-building activities led by diverse faculty. Mapping the intervention components to AI literacy competencies (as seen in Figure 4) suggests that participants were afforded opportunities to develop this skill, though the absence of an explicit measurement of AI literacy limits conclusions about intervention efficacy. Further, while the authors report increases in students' AI knowledge and awareness of diversity issues, the mechanisms underlying these changes are not examined.

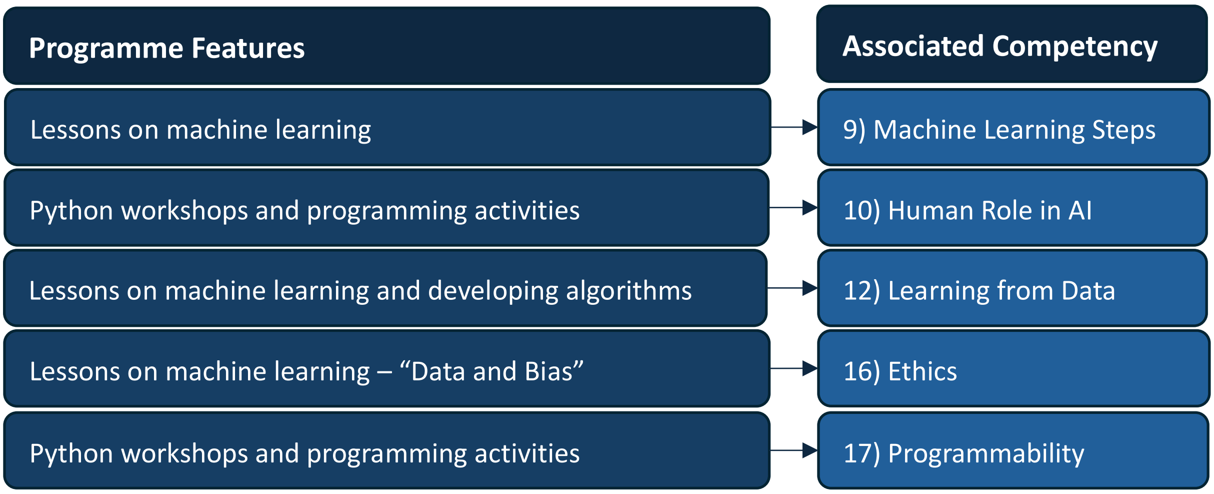


*Figure 4. The AI literacy competencies which relate to intended outcomes from Oskotsky et al.'s (2022) intervention with female Biomedical students from high schools in California.*

With the intention of increasing participants' AI literacy, Hopson et al. (2025) conducted a quasi-experimental evaluation of an extracurricular intervention with undergraduate pre-medical students. It followed a 4-week mini-curriculum which "integrated foundational AI concepts, ethical frameworks, hands-on engagement with (relevant medical image analysis techniques), and exposure to current AI research". These activities were successful according to the pre- and post-test results for their 16-item survey, though neither this instrument nor the programme design are directly derived from Long and Magerko's (2020) framework. Because of this, it is challenging to compare the AI literacy gains of students with other interventions – thus reinforcing the need for a standard tool as also highlighted by Zhang, Prasad and Schroeder (2024).

Together, these studies exemplify the gaps with previous interventions: firstly, that either AI literacy or career interest intend to be improved; secondly, that despite the popularity of Long and Magerko's (2020) definition, measures of AI literacy are still not consistent. Because of these factors, using an existing instrument to measure AI literacy which is aligned with the popular conceptualisation will create robust evaluative research whose results are transferrable. Laupichler et al.'s (2023) "Scale for the assessment of Non-experts' AI Literacy" (SNAIL) may be appropriate for evaluating an intervention group comprising students from outside of computer science. Though further content validity, reliability and cross-cultural validity testing would be required, according to Zhang, Prasad and Schroeder (2024), another measure aligned with the above framework is the 31-item scale by Hornberger, Bewersdorff and Nerdel (2023). A short version of this scale was more recently developed (Hornberger et al., 2025) which would be better suited to maintaining the brevity of a survey in which additional items (not just AI literacy) are to be measured.

### 2.2.7. Barriers to Attendance

With this project intending to remove barriers to PhD studies in AI, it is paramount that the intervention itself is accessible and equitable. The inability of members of the target group to access training will diminish the intervention's efficacy.

The digital divide, defined as "a division between people who have access and use of digital media and those who do not" (Dijk, 2020, p.1) is the most significant barrier to the AI literacy development of school age pupils in the UK (Department for Education, 2025). Despite this, the variability of internet access among UK university students is not well documented. So, it is advisable to implement hybrid engagement modes as a precaution, as is the case for UNESCO's Pre-AI initiative which comprises in-person and online training programmes as well as Massive Open Online Courses (UNESCO, 2025).This also facilitates the attendance of others who require this flexibility such as some disabled students or those with caring responsibilities (Y. Zhang et al., 2025).

### 2.2.8. Conclusion and Recommendations

Given that participants' engagement with AI learning is cyclically enabled by their AI literacy and self-efficacy, and that AI literacy increases career interest, this concept is a useful focus for the design of this training. Drawing on the review in this section, it is recommended to test participants' AI literacy as a result of attending the training programme.

The evaluative research in prior research is largely inconsistent with the prevailing definition for AI literacy (Long and Magerko, 2020), but following this comprehensive framework for both designing the programme and evaluating its efficacy would align with recent recommendations (Zhang, Prasad and Schroeder, 2024). The intervention itself needs to be suitable for its intended participants, and to this end the uncertainty of digital access among university students has been highlighted as an important design consideration.

Structural barriers presented by AI policy across UK universities are discussed in the following section (Section 4). The findings derive from an exploratory analysis into listings for PhD opportunities in the UK.

# 3. Structural Barriers

The motivation for undertaking the policy analysis stems from the increasing influence of AI in the HE sector, intertwining with broader socio-economic conditions that affect access and doctoral education more generally. First, we review the literature on AI use in the HE context, which we believe has an early influence on students' broader use of AI. Second, we outline key socioeconomic conditions that prevent students from pursuing PhD pathways. Finally, we produced an exploratory analysis of PhD job advertisements to further conceptualise the fragmentation of AI policy and expose some of the structural barriers to AI PhD pathways.

## 3.1. Shaping AI Use on Campus – Global and UK Perspective

The literature on policy and AI in the HE context focuses primarily on how universities govern the use of AI within their institutions. The available literature offers comparative insights into international and local policies.

Globally, a comparative analysis of 343 universities in Australia, Canada, China, the UK, and the US reveals a patchwork of AI policies ranging from permissive instructor discretion to outright bans, with academic misconduct cited as the primary rationale for opposing AI use. At the same time, universities often lack a clear explanation or support resources (Parker *et al.*, 2025). Similarly, a study of 10 AI policies across US institutions uncovers multilevel governance with varying levels of responsibility, supporting the aforementioned fragmentation of AI policy (Wu, Zhang and Carroll, 2024). The researchers highlight that the ambiguous AI policy and guidance at universities might, in fact, support the characteristics of academic enquiry, allowing students and staff to improve AI literacy and raise awareness of responsible AI use among communities, rather than policing its use.

The fragmentation of AI policy in HE is also highlighted in the UK Russell Group study (Atkinson-Toal & Guo, 2024). Researchers found that institutions employ a wide variety of approaches to determine what is permitted in the use of generative AI in learning and assessment. Although the unacknowledged use of AI typically leads to contract cheating and academic misconduct, the acceptance of AI use varies considerably across UK universities (Table 1). The provision of guidance materials also varies across institutions, with some offering no practical support. The authors Atkinson - Toal and Guo (2024) also emphasise that such a divide in the provision of AI skills may create risks of a digital divide in technology adoption and call for standardised core policy elements to ensure equal and supportive access to AI for students in the UK.

*Table 1. The Fragmentation of National Standards. This comparison illustrates how the broad mandates of the Russell Group are interpreted differently at the local level, resulting in distinct governance models ranging from "Co-Pilot" support (Bristol) to "Gatekeeper" restrictions (Oxford).*

| Policy Dimension | Russell Group Principles | University of Bristol (Our Focus) | University College London (The "Progressive" Model) | University of Oxford (The "Traditional" Model) |
|---|---|---|---|---|
| **AI Literacy** | **Mandate:** Universities *must* support students/ staff to become AI-literate, balancing opportunities with ethical risks. | **Approach:** "Foundational." Focuses on long-term skills but warns against "learning loss" and replacing intellectual rigour. | **Approach:** "Core Competency." explicitly treats AI proficiency as a required future workforce skill. Heavy emphasis on *how* to prompt and iterate. | **Approach:** "Cautious Engagement." Focuses on the limitations and risks of AI. Literacy is defined by understanding what AI *cannot* do (hallucinations/bias). |
| **Staff Guidance** | **Mandate:** Departments should apply institution-wide policies to their specific context. | **Approach:** "Co-Pilot Model." Staff use AI to sense-check and guide. Strong emphasis that AI cannot replace core disciplinary concepts. | **Approach:** "Designer Model." Staff are tasked with redesigning assessment specifically to either embrace or exclude AI. | **Approach:** "Gatekeeper Model." Tutors define strict boundaries. Emphasis on Socratic defence (Vivas) to verify student knowledge outside of AI. |
| **Assessment Categories** | **Mandate:** Assessment must ensure "equal access" and fairness regarding paid tools. | **Approach:** "Transparency & Trust." Students must disclose *why* and *when* AI was used. Links policy to **EDI/Decolonisation** goals regarding bias. | **Approach:** "The Three-Tier System." Assessments are explicitly labelled:<br>1. AI-Banned<br>2. AI-Assisted<br>3. AI-Integral (Required). | **Approach:** "Prohibited unless specified." The default stance is restriction. AI is generally treated similarly to third-party authorship unless explicitly allowed. |
| | | | | |

| | | | | |
|---|---|---|---|---|
| **Academic Integrity** | **Mandate:** Policies must be clear; students should not fear penalisation for discussing AI challenges. | **Approach:** "Student Agency." AI is permitted only when it does not diminish the student's capacity to derive meaning. It lacks "innate intelligence." | **Approach:** "Acknowledgement." Focus is on proper citation. If you use it, you must declare the prompt and the output trail. | **Approach:** "Authorship Definition." AI cannot be an author. submitting work generated by AI as one's own is treated strictly as plagiarism. |
| **Institutional Context** | **Context:** The 24 leading UK research universities. | **Context:** "AI University of the Year." Heavy investment in **Isambard-AI** creates pressure to lead technically while managing pedagogical risk. | **Context:** London-based, tech-forward. Positions itself as an innovator in assessment design. | **Context:** Historic prestige. Prioritizes the "human element" of tutorial-based teaching and original thought above tool usage. |

The influence of such fragmentation on institutional AI guidelines is reflected in studies that measure students' perceptions and awareness. In a survey by Jiang and colleagues, approximately 79% of students (124 undergraduates) reported the presence of AI policies and restrictions in their classes, but only 48% found these policies clear. Institutional-level guidance was even less recognised (48%), indicating ambiguity and inconsistency (Jiang et al., 2025). They also reported that students who used AI tools for writing or exam support reported greater awareness of AI's limitations, suggesting that active engagement with AI promotes critical understanding. In another study measuring perceptions of AI governance, students prioritised curriculum integration, ethical guidance, misuse detection and fair academic sanctions for unacknowledged use of AI in assessment (Barus et al., 2025).

## 3.2. Conditions Beyond AI Policies

While the AI use policies in HE shape the academic practices and those of doctoral researchers, they do not operate in isolation. As the literature review section on socio-economic barriers explains, there are wider structural barriers stemming from the broader STEM disciplines. The research of Russell Group Universities found minimal support for EDI and other accessibility factors within policy provisions (Atkinson-Toal and Guo, 2024). Most universities acknowledge AI's inclusive potential but focus instead on copyright and intellectual property risks. Notable exceptions included the University of Oxford's explicit recognition of AI's benefits for users with disabilities, the University of Sheffield's attention to environmental impacts and a RAISE framework, from Queen's University Belfast, which explicitly foregrounded equitable access to AI.

At the University of Bristol, the Centre for Doctoral Training in Interactive AI explicitly embeds EDI discussions in its curriculum, including content on demographic bias in AI systems, democratisation of AI technologies, and inclusive design of AI tools. However, this is specific to that programme rather than the institutional EDI policy itself (University of Bristol, no date).

Other factors relevant to this policy context include finance and the national migration and work visa rules. The financial cost of completing a PhD is one of the most significant barriers to postgraduate research (Lindner, 2020), with typical fees for home students at the University of Bristol ranging from at least £20,000 for a four-year PhD programme. Fees for overseas students are, of course, substantially higher. A study of postgraduate students at the University of Bristol in the 2022/23 academic year found that 79% of PGRs reported a need for supplementary income, with 82% citing the need to cover essential living costs (Smith et al., 2023)

Finally, the same report found that 86.5% of the international PGRs experience some form of hidden costs, in the form of visa costs, National Health Service (NHS), access to educational resources, health insurance, or other miscellaneous costs.

### 3.2.1. The University of Bristol's AI Position

University of Bristol has been named the UK's AI University of the Year in 2024. Such a position necessitates an encompassing strategy for inclusion of various expertise, schools and faculties and, as a research-intensive institution, a growing number of research staff, including PhD programmes, to drive its position forward. In this section, we provide an overview of Bristol University's AI policy for students, following the Atkinson-Toal and Guo (2024) study of the Russel Group Universities. We used the same approach to analyse the content of the digital document: "Chan (2023) proposed that effective AI policy requires three key considerations for effective management and practical integration: (1) pedagogical to improve learning and teaching; (2) governance to address privacy and security concerns; and (3) operational considerations regarding training, accessibility, and integration."

*Table 2. Summary of the University of Bristol AI policy for students.*

| Area | Key Points | Details |
|---|---|---|
| **Pedagogical (Improving Learning & Teaching)** | Integration of AI for learning support | AI allowed for *general revision*, planning, and understanding unit content. Students can use LLMs to generate study plans using Unit Aims, ILOs, covered topics, and timelines. |
| | Developing critical AI literacy | Policy emphasises fact checking, awareness of hallucinations, and critical engagement with AI-generated outputs. Students are encouraged to use AI tools to enhance understanding, not to replace academic |

| | | |
|---|---|---|
| | | thinking. |
| | Maintaining academic integrity and learning outcomes | AI should not undermine ILOs or degree classification standards. Writing or paraphrasing by AI for assessments is considered contract cheating. Grammar and spelling checks are permitted. |
| | Support for better AI use | Examples of recommended prompts are provided to help students generate more informed, critical results from LLMs. |
| **Governance (Privacy, Security & Academic Integrity)** | Consistent but flexible institutional policy | University-wide policy focused on LLMs, but module leaders may define their own assessment-specific AI guidance. |
| | Academic integrity and misconduct | Students must *acknowledge AI use* in assessments. AI-generated text without disclosure can constitute academic misconduct. |
| | AI errors and risks | The policy highlights hallucinations, errors, and the need for verification. Promotes responsible and informed use of AI. |
| | AI detection limitations | Turnitin's AI detection is allowed but considered unreliable. Staff are cautioned about false positives and accuracy issues. |
| | Guidance for ethical use in research | University neither bans nor promotes unrestricted use; encourages critical engagement while maintaining ethical standards. |
| **Operational (Training, Accessibility & Integration)** | Guidance and training resources | Direct links to citation support (e.g., Cite Them Right, librarians). Helps students and staff navigate referencing AI correctly. |
| | Staff support | Staff receive their own AI guidance to manage detection tools, assessment design, and marking support. |
| | Integration into learning workflows | AI tools framed as assistive companions for studying, revision, and learning processes, rather than substitutes for academic work. |
| | Accessibility and evolving technology | Policy acknowledges the rapid evolution and availability of AI, promoting engagement while preparing students for a digital future. |

## 3.3. Data Collection Pipeline

To complement the preceding literature review, this section provides a summary of the empirical landscape of AI doctoral studies in the UK. Whilst the previous sections identified the socio-economic, digital literacy and structural barriers that shape access to postgraduate studies in AI, the analysis that follows in this section examines how these barriers are articulated, or obscured in the live PhD marketplace.

We used a custom web scraping tool to collect a dataset of 2,643 unique listings from www.findaphd.com in December 2025, employing a broad, keyword query ("Artificial+Intelligence+OR+Machine+Learning+OR+Deep+Learning+OR+Data+Science+OR+Robotics+OR+Computer+Vision+OR+Natural+Language+Processing). This strategy ensured a coverage of core AI research projects as well as interdisciplinary applications, where AI methods are embedded within broader research domains.

After data scraping, we extracted structured information to use for analysis. Since the listings are varied and consist only of unstructured listing text and a title, a pipeline was developed to extract structured information on data relevant to our research questions. To extract information we developed an LLM JSON-schema extraction pipeline (Agrawal et al., 2022). Rather than relying on pre-defined coding categories, the prompt was designed by iterative work between the literature review and exploratory analysis of the PhD listings. We tested and updated the prompt separately for each dimension extracted by manually labelling randomly selected PhD listings and comparing the results with the LLM output. Where discrepancies emerged, the prompt was revised and re-evaluated until it achieved satisfactory alignment with the underlying text. Once we had satisfactory coverage of the key facts, we applied the entire prompt to all listings. The model Gemini-3-Flash was used for the extraction. The detailed prompt structure is available in Appendix A.

To assess the quality of the extracted data, we have added fields for the Large Language Model (LLM) to evaluate the sufficiency of information for each listing. These include True and False statements to the following question: "Are entry requirements missing?" “Is the supervisor missing?” “Is the university missing?”. Additionally, given the hallucination risks associated with short or vague PhD listings, we added the following True or False statement: True if the text is very short/vague but the schema is fully populated (indicating a likely hallucination). Taken together, these measures will help us identify potential data flaws and exclude non-significant results from the final analysis, leaving us with a final count of 2,382 listings.

### 3.3.1. Limitations

Before reading the findings, it is important to acknowledge several limitations of our methodology. Firstly, analysing the listing itself is misleading if information is present on other websites, for example a university website linked to by the listing might give additional context on entry requirements, funding, or EDI initiatives, which would be missed if not present on the findaphd.com listing. Secondly, while LLMs enable classification and extraction of text at scale, their outputs are probabilistic and thus depend on the training data. Rather than treating these findings as institutional truths or as recruitment strategies across universities, the output offers insights into patterns of information represented in the online text. Thirdly, such a classification approach is likely to be biased by context-specific cues. Although we manually reviewed the listings and compared outputs on approximately 5% of the data before employing the finalised prompt, certain attributes, such as gender, inclusivity, or funding certainty, cannot be directly observed and must be inferred from language, introducing the risk of oversimplification. These risks are particularly salient for non-Anglophone names, non-binary identities, and listings that rely on ambiguous or formulaic phrasing.

## 3.4 Findings

### 3.4.1. Financial Accessibility

The data analysis shows that PhD funding is uneven and uncertain. While almost half of the PhD listings (1,146) include a funding opportunity, the remaining listings either require self-funding (154) or require prospective students to apply via a university competition (507), where funding is allocated at the faculty or institutional level rather than tied to a specific project (Figure 5). Although such competitions are common and often provide full fee coverage and stipends, they introduce an additional layer of uncertainty for prospective candidates.

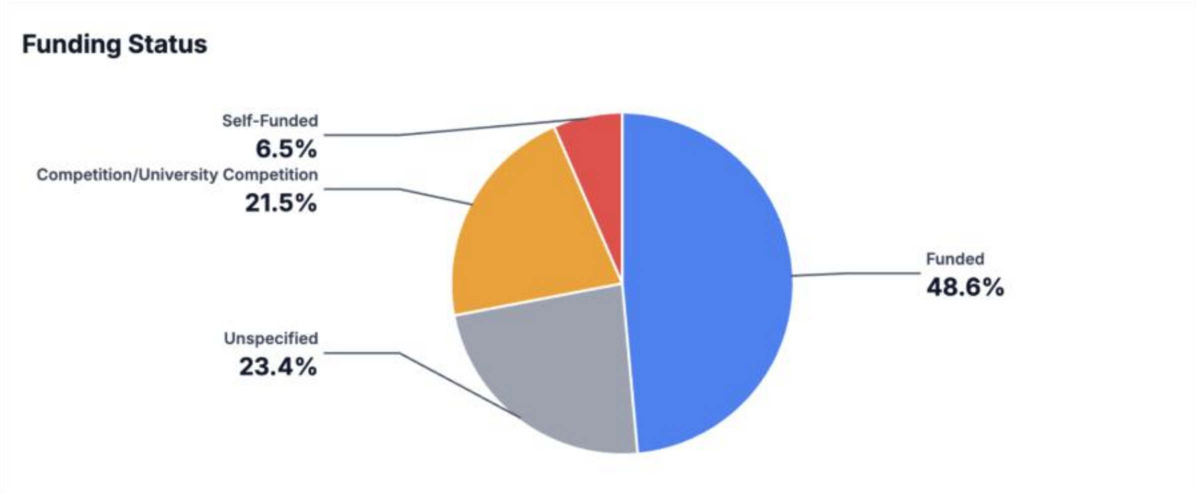


*Figure 5. Pie chart displaying funding status proportions among PhD listings.*

The financial burden of completing a PhD is a substantial barrier to entry into postgraduate research programs (Lindner, 2020). This is further reinforced by the finding that 23.4% of listings do not provide explicit information on funding arrangements, leaving applicants unable to assess the financial feasibility of pursuing the position. In a highly competitive and resource-intensive field such as AI, this lack of funding transparency may deter prospective candidates without independent financial support. We also recognise that, because of the LLM extraction method, not all funding details were captured. Some listings may include funding information on their institutional pages, which prospective students need to access when applying.

Therefore, we conclude that the number of unspecified listings might be lower, but our data collection and analysis methods were not able to capture this.

*Figure 6. Advertised annual PhD stipends. A distribution of explicit financial offers across the dataset. The data reveals a massive concentration around the national median of £20,780, confirming that when funding is*

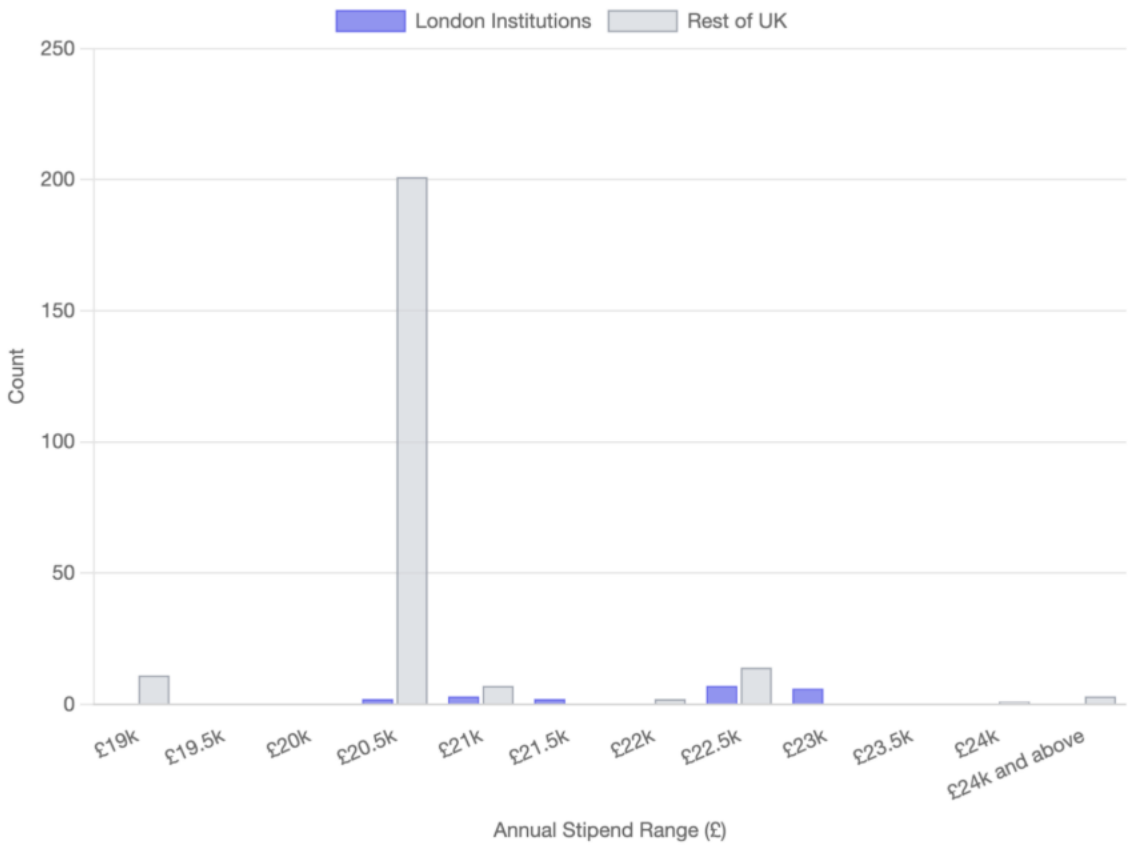


*openly advertised, it heavily aligns with standard UKRI minimum guidelines, alongside expected adjustments for London-based institutions.*

Although the financial security of prospective postgraduate researchers is often downplayed by unclear funding statements, we find that those that include financial details align with current UK Research and Innovation (UKRI) guidelines (Figure 6). It is also important to note that UKRI announced an increase of almost 5% in PhD stipends for students starting their studies in the academic year of 2026/27. The announcement was released in early 2026 ('Support for UKRI-funded students', no date) just after our data collection, therefore, our dataset does not capture this.

Beyond core funding, additional financial resources dedicated to training, conferences, and research dissemination serve as vital components of funding security. Such resources are essential for fostering skills development and enhancing academic visibility throughout doctoral training (Figure 7).


*Figure 7. Doughnut chart showing the presence of text mentioning explicit financial top-ups in UK AI PhD listings. Only 30.3% of adverts make explicit mention of additional funding beyond the standard stipend, while a majority of 69.6% contain no explicit information.*

Only 715 listings included information on the financial top-up, suggesting that opportunities for professional development support are often implicit. The analysis of available top-up examples indicates that their availability varies substantially. For instance, projects in collaboration with industry partners offer a financial stipend top-up of £10,000:

> "*This project benefits from the support of a large industrial partner and this covers the full cost of UK tuition fees and a significantly enhanced tax free stipend of £31,383, which includes the standard UKRI studentship of £21,383 for 26/27 and a £10,000 industrial top-up".* (Quote from the University of Southampton PhD listing: "Large-scale active noise and vibration control system design").

Other projects include a top-up of £1000 per annum to support research costs:

> *"Each GTA Scholarship is funded for 4 years and covers research degree tuition fees at the Home/EU rate, £1000 per year in research support costs, as well as a Doctoral research stipend to cover living costs, and a part-time teaching salary."*

Meanwhile, some listings include a total amount of additional funding across the duration of a PhD programme:

> *"Full scholarships include tuition fees, a stipend at the UKRI rate plus 10% ORC enhancement tax-free per annum for up to 3.5 years (totalling £22,858 for 2025/26, rising annually) and a budget of £4200 for things like conference travel."*

While the industrial partnership top-up is understandable given their commercial orientation, the other examples of financial support for conferences, training, and workshops remain concerning, given the increasing costs of participation in international events. For instance, the leading machine learning and artificial intelligence conferences, such as NeurIPS or ICLR, charge over US$400 (~£300) for full-time students.

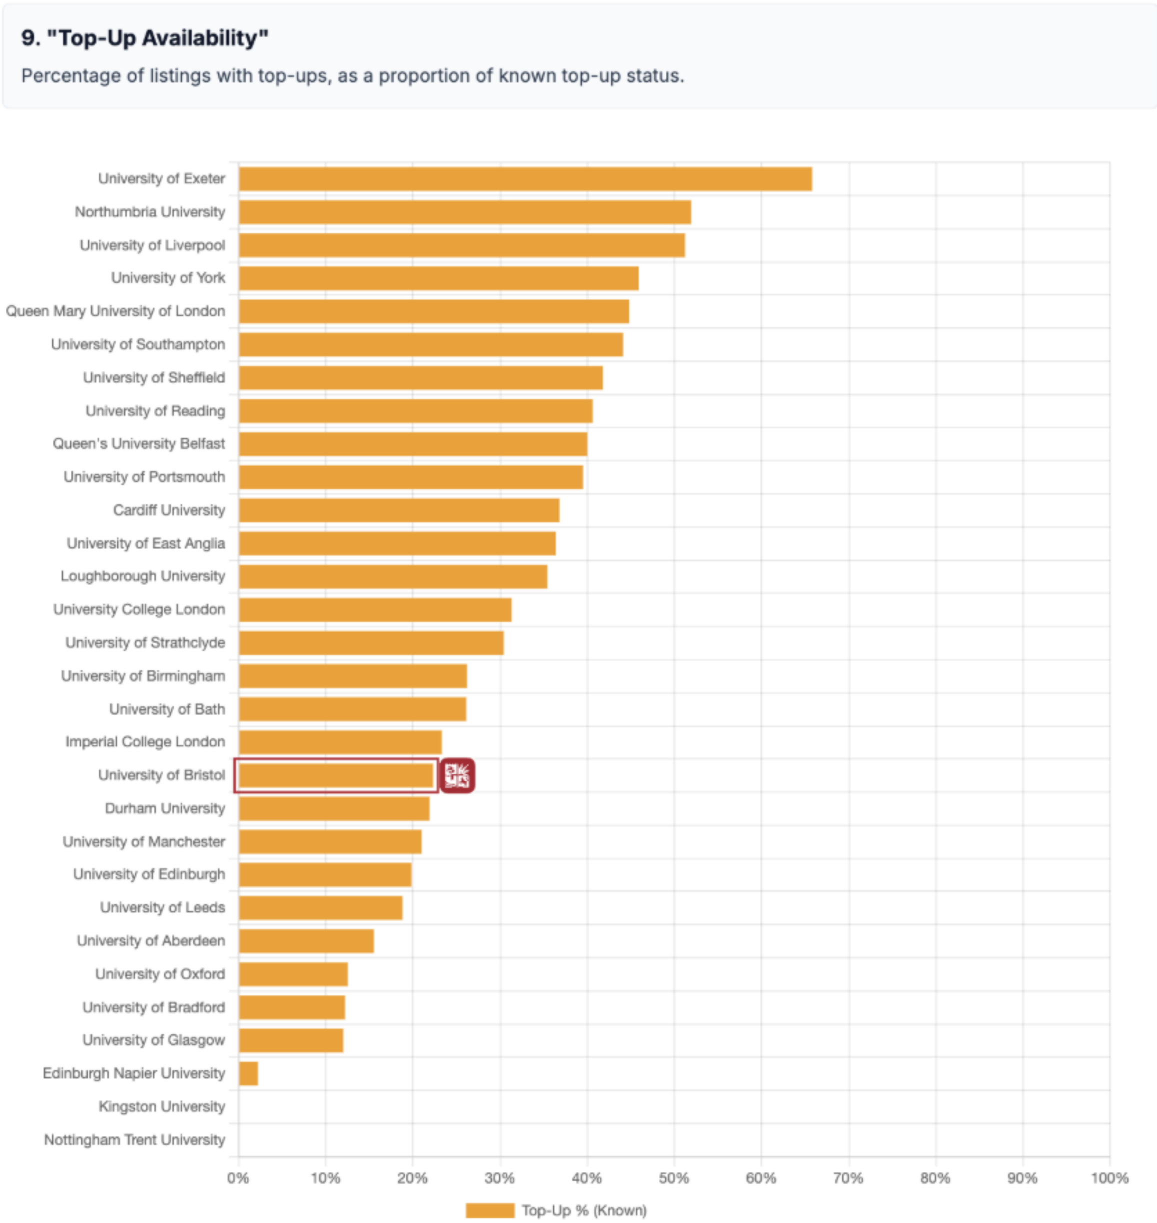


*Figure 8. Availability of financial top-ups by institution. Bar chart comparing the percentage of AI PhD listings that offer additional funding across UK universities. The University of Bristol (highlighted) sits near the middle of the distribution at approximately 22%.*

During our analysis, we found that some PhD listings that appear to be funded still require applicants to secure funding through an internal competition, typically at the faculty or university level. While our initial analysis, shown in Figure 8, distinguishes between directly funded and competition-based funding, some listings with funding are ambiguously worded. To better identify this, we employed large language models to evaluate the clarity of funding statements rather than relying only on the provided funding labels.

Using a keyword-based approach, we classified statements as either secured or competitive/conditional. Among funded listings (Figure 9), 962 (84,2%) include explicit statements of secured funding, indicating that funding is already secured for the advertised position. In contrast, only 9 listings are subject to conditional funding or an additional selection process, whilst 172 listings are unclear. This highlights a source of ambiguity in the PhD funding landscape: some positions that appear funded may entail uncertainty, particularly for those without institutional or insider knowledge of university funding mechanisms.

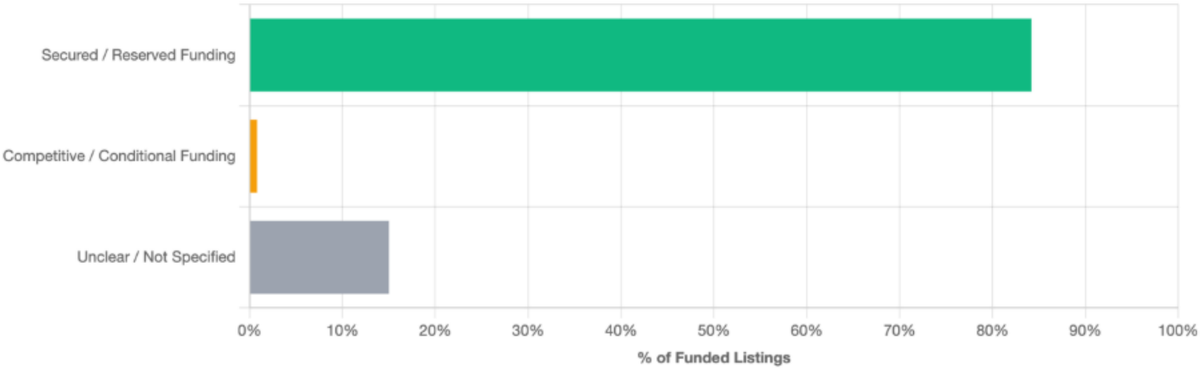


*Figure 9. Clarity of funding within funded PhD listings, showing that the vast majority (84.2%) of explicitly funded positions offer secured financial support, while roughly 15% remain vaguely defined or highly conditional.*

As we outlined above, university competition funding is relatively popular across PhD listings; other funding sources are also available to prospective postgraduate students. The most common contributors are universities and UK Research and Innovation (UKRI) Research Councils, with 630 and 604 listings, respectively. The Engineering and Physical Sciences Research Council (EPSRC) accounts for approximately 9.6% of the analysed PhD listings.

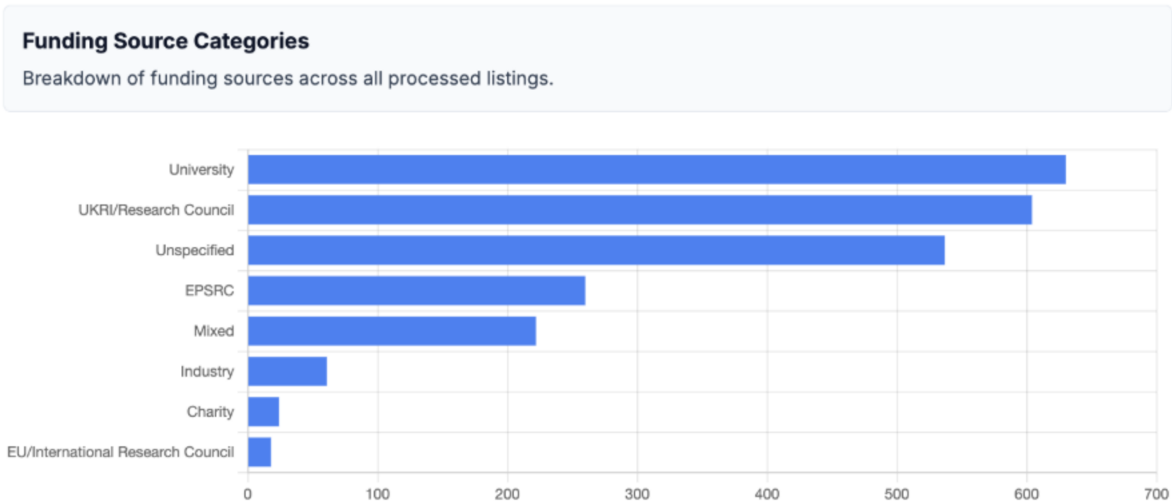


*Figure 10. Primary funding sources for AI doctoral programs. Horizontal bar chart illustrating that host Universities and UKRI/Research Councils are the dominant financial sponsors in the dataset, while direct industry and charity contributions remain comparatively low.*

### 3.4.2. Global Ambition, National Restrictions

When analysing the dataset (excluding listings with missing or corrupted text), fee coverage for funded roles is split almost evenly between Home-only (829) and International (821) students. However, the fragmentation between institutional ambition and national regulation creates a severe structural barrier for overseas talent.

The data indicates a high degree of advertised openness: 2,262 listings do not explicitly restrict applications to UK nationals in their advert text, implying international students are welcome to apply. However, when cross-referencing this advertised openness with actual financial support, a systemic financial trap emerges. There are 706 projects that allow international students to apply but strictly cap their funding at the "Home Only" fee level.

This discrepancy is largely driven by national UK Research and Innovation (UKRI) policies. While UKRI rules permit up to 30% of a cohort to be international students, the financial award is capped at the domestic fee rate. Unless the host university agrees to absorb the difference, international applicants are legally permitted to apply but are systematically priced out, forced to self-fund a tuition shortfall of roughly £15,000 per year. Ultimately, this relies on omission in academic advertising, leaving applicants to navigate complex funding caveats entirely on their own.

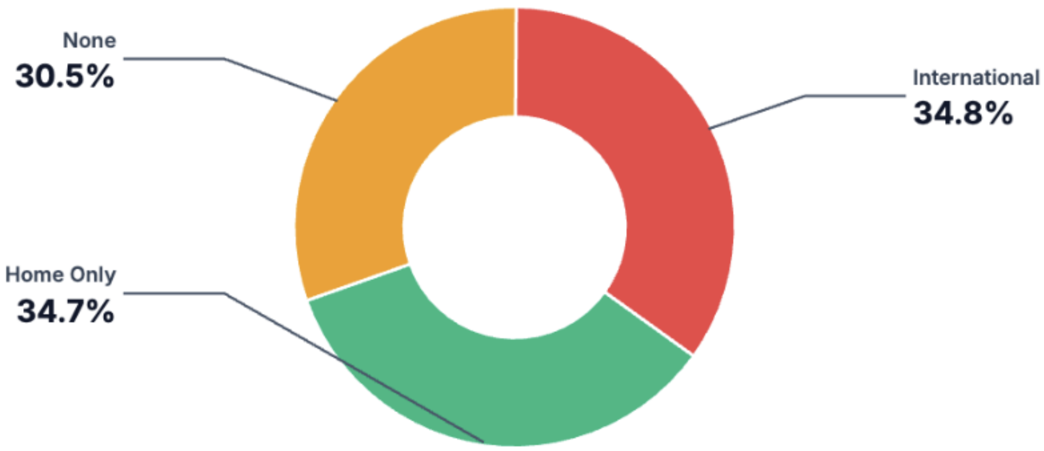


*Figure 11. Breakdown of PhD Listing Eligibility: A comparison of advertised openness versus funding restrictions, highlighting that over a third of listings (34.7%) restrict support to "Home Only" levels despite a similar volume of "International" eligibility.*

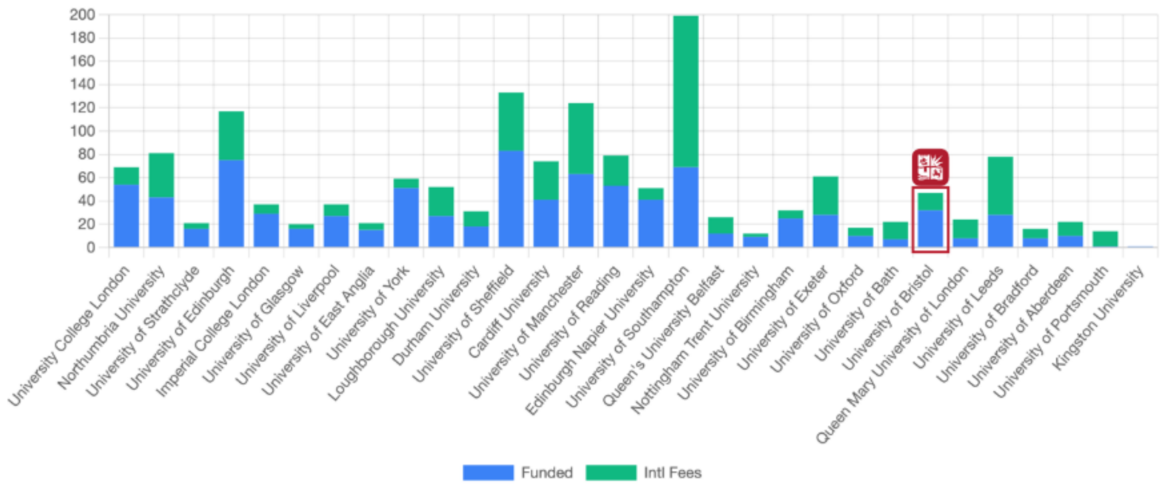


*Figure 12. The 'International Paywall' across UK universities. Stacked bar chart illustrating the "false hope" funding trap. It compares fully funded positions (blue) against the volume of projects that invite international applicants but leave them to cover significant fee shortfalls (green). The University of Bristol is highlighted for case-study context.*

### 3.4.3. Summary and Conclusions

This section examines the financial accessibility of AI PhD programmes in the UK. Our analysis reveals substantial variation in funding structures across universities, not only in terms of stipend availability but also in the mechanisms through which funding is allocated. While many listings present positions as funded, closer inspection shows that some require applicants to compete in additional institutional funding rounds, introducing an element of uncertainty into what appears to be secured support.

The stipends in our dataset align with the UKRI guidelines. Although the median stipend aligns with UKRI guidance, the scarcity of additional support may constrain doctoral candidates' ability to fully participate in professional development activities that are central to academic progression in AI. Securing a PhD admission is not only related to AI preparedness or digital skills literacy. As we outlined in Section 2 of the Literature Review, both AI literacy skills and self-efficacy or in other words, the psychological factors influencing personal confidence in the ability to engage with AI learning, act as gatekeepers to pursuing studies in AI. Therefore, the uneven financial access to developing core AI skills and connecting with like-minded peers might shape who can access the field and become visible.

Finally, as we outlined in Section 1 of the Literature Review, UK institutions are becoming increasingly international, with doctoral students and staff coming from diverse geographical and social backgrounds (Universities UK International, 2024). However, our findings do not support this ambition. Despite 95% of listings in our data being open to international students, only a fraction include a clear funding statement covering international fees.

## 3.5. Formal Requirements Show Lesser Grade Pressure

It is often presumed that pursuing postgraduate studies requires exceptional academic achievement and a First-Class degree. However, this notion appears to be evolving, as most universities now accept candidates with a commendable 2:1 degree, or, as described more flexibly by universities as a “strong degree,” suggesting a broad institutional shift away from rigid academic exceptionalism at the point of entry. This can be observed in the Figure 6 below, which shows the overwhelmingly high number of unspecified degree requirement (grey bar). However, me must be cautious with the formal degree requirements, as this does not necessarily reflect informal shortlisting practices at the institutional level, including supervisory preferences, or competitive dynamics among applicants.

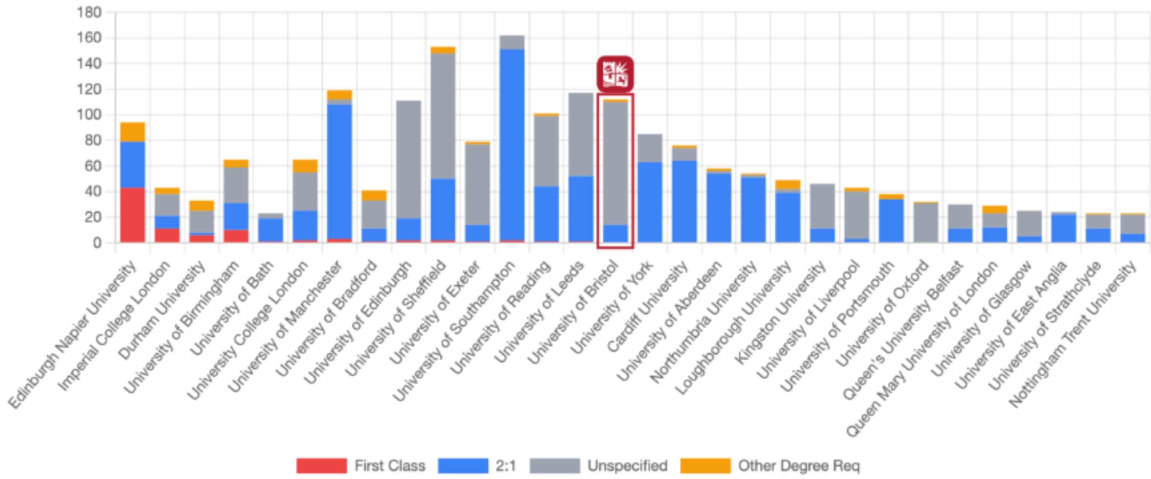


*Figure 13. Academic grade requirements across AI PhD listings. A comparison of formal degree expectations across universities, illustrating a broad institutional shift toward accepting 2:1 ("strong") degrees over rigid First-Class requirements.*

Beyond degree classifications, alternative pathways into doctoral study remain limited. Only around 35% of listings explicitly recognise equivalent practical or professional experience as a valid basis for assessing research suitability. This indicates that, despite some movement toward flexibility in formal requirements, academic credentials continue to dominate as the primary mechanism for evaluating candidates.

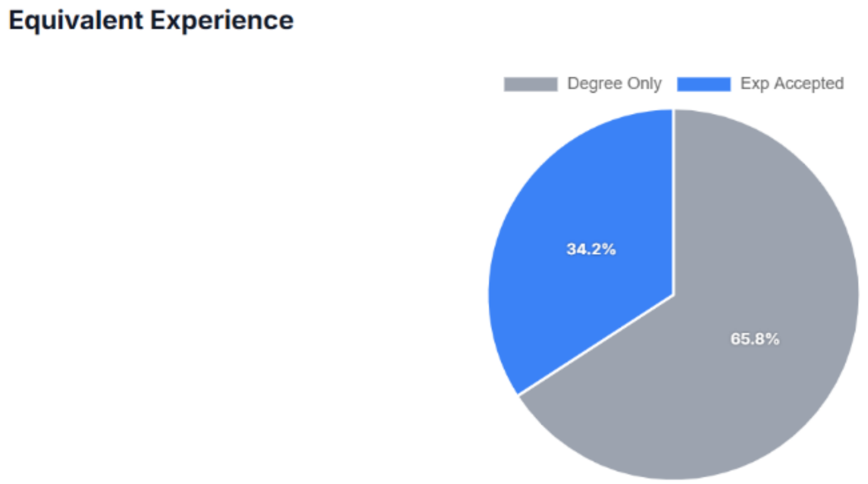


*Figure 14. Acceptance of equivalent professional experience. An illustration showing that nearly two-thirds of listings strictly require formal academic degrees, indicating that alternative pathways into doctoral study remain limited.*

## 3.5.1. Application Processes and Diversity Initiatives

In addition to formal entry requirements, we also analysed the more subtle indicators of potential exclusion in PhD listings. Elitist keywords (Figure 15), such as references to “world-class” institutions or “leading” research environments, are relatively uncommon overall. This suggests that symbolic signalling of prestige is not a dominant feature of AI PhD listings.

On the other hand, a small subset of universities requires students to contact supervisors before applying. While it is often understood as a means of early engagement, such a requirement can create a barrier for candidates with limited prior academic networks or confidence in approaching senior academics. For applicants without such qualities, this expectation might function as an additional gatekeeping mechanism, rather than an inclusive practice.

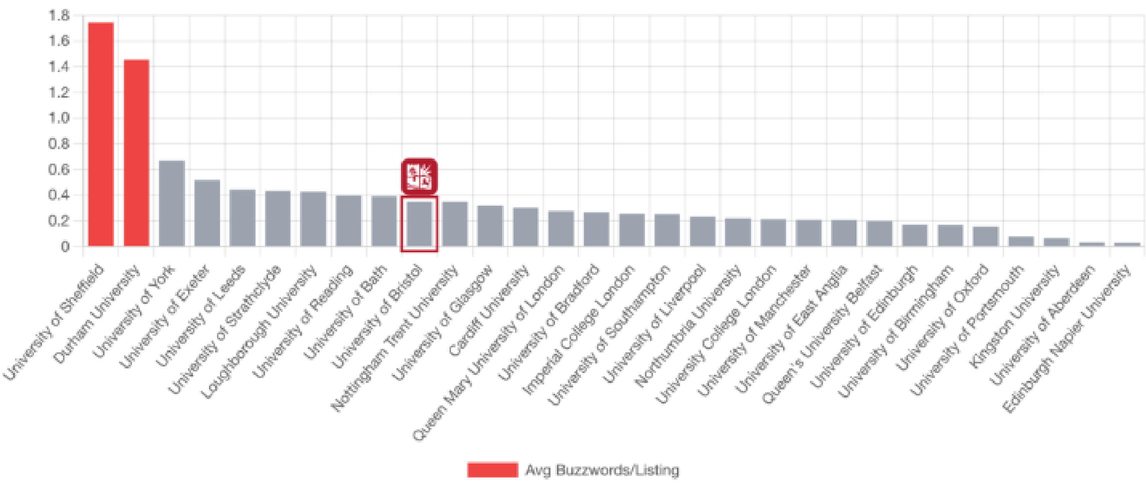


*Figure 15. Prevalence of elitist language by institution. An analysis of explicit prestige buzzwords (e.g., "world-class", "elite") per listing, demonstrating that symbolic signalling of prestige is relatively uncommon overall.*

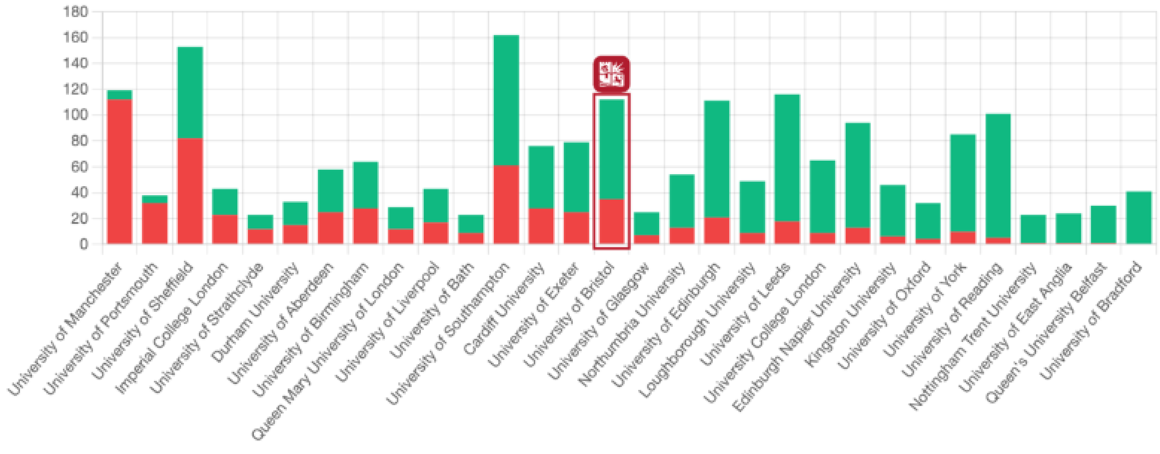


*Figure 16. Pre-application networking requirements. A breakdown of universities mandating applicants to contact a supervisor before applying, highlighting a potential social capital barrier for candidates lacking established academic networks.*

We also examined the presence and framing of diversity initiatives within PhD listings (Figure 16), distinguishing between two forms of inclusion signalling. The first category captures *active* diversity language, where listings explicitly reference concrete measures such as guaranteed interviews, targeted bursaries, or legal frameworks (e.g. the Equality Act). The second category captures *passive* diversity language, typically consisting of generic encouragements to apply without specifying any actionable support.

This absence is particularly striking given the openness of UK Universities and the growing institutional emphasis on equality, diversity and inclusion within the UK higher education. While the PhD advertisements are often driven by the institution-wide policies and regulations, the absence of such statements might act as an additional hidden barrier, particularly for candidates who require reasonable adjustments.

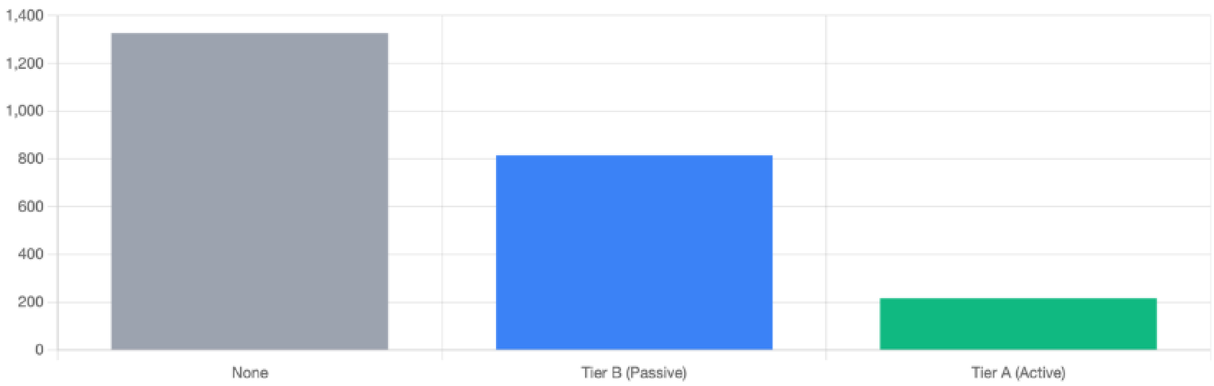


*Figure 17. Presence and framing of diversity initiatives. A comparison of listings containing no diversity statements versus those with passive encouragements (Tier B) and active, actionable support frameworks (Tier A).*

Notably, across the entire dataset, we found no explicit references to neurodiversity (Figure 17). This silence suggests that neurodivergent applicants remain largely invisible within the recruitment discourse of AI PhD programmes, despite increasing recognition of neurodiversity within both academic and technology sectors.

### 3.5.2. Gender Representation

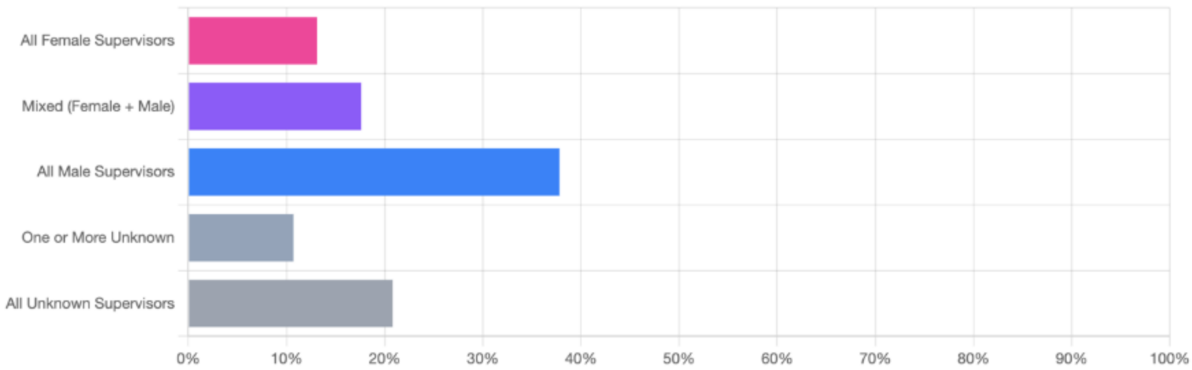


*Figure 18. Inferred gender composition of AI PhD supervisory teams. A breakdown of supervisor configurations across listings, illustrating that the market is heavily dominated by male-only supervisory teams, reflecting a persistent "gender void" in visible AI leadership.*

The data reveals a stark "Gender Void" in AI supervision (Figure 18). As shown in the figure above, the market is dominated by presumed male-only supervisory teams (blue bar). Research confirms that the lack of visible female role models significantly lowers the AI self-efficacy of female applicants, potentially causing them to self-select out of the field before they even apply.

While we consider supervisory gender composition an important dimension of access and symbolic inclusion, its measurement via LLM-based extraction raises significant methodological challenges. First, models frequently struggle to infer gender from non-Anglophone or gender-ambiguous names, introducing uncertainty that may disproportionately affect international or minority supervisors. Second, we were unable to infer race and ethnicity from supervisor names, which is an intentional omission due to data limitations. Third, gender inference relies on probabilistic associations rather than self-identification, risking misclassification. Therefore, as outlined in the introductory section, the results should be treated as indicative patterns rather than definitive measures of supervisory team composition.

### 3.5.3. Summary

This section covered various dimensions of the PhD application process, illustrating how the socio-economic barriers identified in the literature review are partially reproduced within the UK AI PhD market. In particular, the persistent exclusionary stereotype of AI, as a highly technical and male-dominated field, is reflected in supervisory structures in the listings. Consistent with prior research, our analysis shows that AI PhD supervision in the

UK remains overwhelmingly male-dominated, mirroring broader gender imbalances in STEM subjects (e.g. Frachtenberg and Kaner, 2022; Kalim et al., 2025; MIOIR, 2025; Stathoulopoulos et al., 2019).

Despite the increasing prominence of diversity discourse at the institutional level, explicit diversity initiatives are scarce. Where such statements appear, they are predominantly passive rather than action-oriented, suggesting that responsibility of inclusion is often deferred to broader institutional frameworks and policies. This finding aligns with the discussions related to underrepresented groups, who have historically faced limited access to STEM and AI subjects (Roberts et al., 2021; Ferguson et al., 2023). When coupled with the limited international funding availability, it reinforces the geographical barriers of the broader AI ecosystem.

At the same time, the data reveal some indications of reduced formal academic gatekeeping. PhD listings increasingly avoid rigid emphasis on First-class degree requirements, instead favouring broader criteria such as a “strong” or “good” honours degree. While this shift does not eliminate informal selection mechanisms or disciplinary hierarchies, it suggests a softening of barriers at the point of entry. The following section offers a deeper analysis and overview of this trend.

## 3.6. Discipline Requirement

Contrary to the common assumptions that AI PhD programmes are restricted to candidates with formal training in computer science or mathematics, the analysis suggests that PhD openings are relatively open in disciplinary framing. Only 96 openings were classified as exhibiting what we term “*disciplinary tunnel vision*” (Figure 19), where eligibility was explicitly restricted to a narrow set of technical degrees (see the Appendix for details).

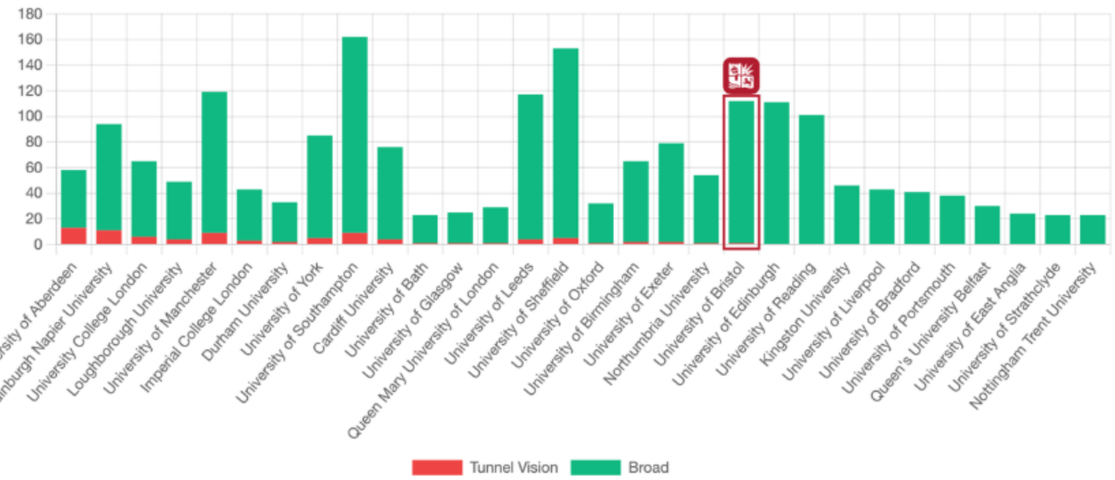


*Figure 19. "Tunnel vision" in academic degree requirements. A comparison of eligibility criteria across universities, demonstrating that explicit restrictions to narrow technical disciplines (like Computer Science or Mathematics) are relatively rare compared to broader academic framing.*

This openness, however, should be interpreted cautiously. Whilst only a minority of listings expect a science background, technical skills are often embedded within project descriptions. Programming languages (Figure 20), most notably Python, are the most commonly required alongside references to machine learning frameworks and computational methodologies. Overall (Table 3), 22.4% of listings explicitly required a programming language, and 25.3% included at least one hard technical requirement (either programming or a technical degree). A further 22.3% referenced specific computer science skills or frameworks, even where these were not framed as strict eligibility criteria.

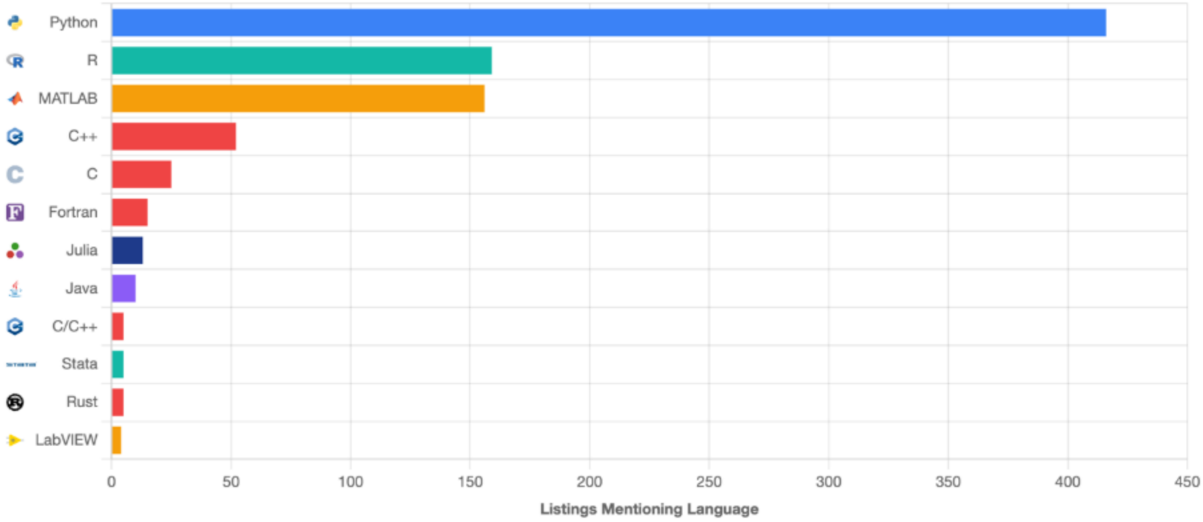


*Figure 20. Most frequently required programming languages. A ranking of the technical prerequisites explicitly embedded within project descriptions, highlighting Python's overwhelming dominance and illustrating the market's reliance on implicit, skill-based filtering.*

*Table 3. Technical requirements mentioned in the data.*

| Requirement Type | Count | Percentage |
|---|---|---|
| Programming language required | 536 | 22.4% |
| Technical degree required (tunnel vision) | 96 | 4.0% |
| CS skills/frameworks mentioned | 535 | 22.3% |
| Any hard requirement (prog OR degree) | 606 | 25.3% |
| No hard CS requirements | 1,756 | 73.2% |

This suggests a shift from formal credential-based gatekeeping (as discussed in the section above) to skill-based filtering. Rather than excluding applicants through narrow degree requirements, listings signal technical competencies. Such signals may function as implicit barriers, particularly for applicants from non-technical or interdisciplinary backgrounds who may meet formal eligibility criteria but lack confidence in applying without demonstrable coding experience.

This pattern is reflected in the predominance of Applied AI projects in the data (Table 4), where interdisciplinary collaboration is encouraged. Projects classified as Core AI show the highest concentration of hard technical requirements, with more than one-third of listings requiring either programming skills or a technical degree. By contrast, Applied AI, the largest category in the dataset, exhibits a more inclusive disciplinary framing, with interdisciplinary collaboration frequently encouraged.

*Table 4. Comparison of technical skills and educational background between areas of AI study.*

| Application Area | Total | Prog Req | Tech Degree | Any Hard Req | No Hard Req |
|---|---|---|---|---|---|
| Applied AI | 2066 | 497 (24.1%) | 81 (3.9%) | 559 (27.1%) | 1474 (71.3%) |
| Not AI | 138 | 10 (7.2%) | 2 (1.4%) | 11 (8.0%) | 127 (92.0%) |
| Tangential / Non-Core | 122 | 14 (11.5%) | 4 (3.3%) | 16 (13.1%) | 106 (86.9%) |
| Core AI | 52 | 15 (28.8%) | 9 (17.3%) | 20 (38.5%) | 29 (55.8%) |
| Policy & Ethics | 5 | 0 (0.0%) | 0 (0.0%) | 0 (0.0%) | 5 (100.0%) |

### 3.6.1. Summary

Taken together, this section demonstrates that AI PhD listings in the UK are, on paper, relatively open in their disciplinary framing. Explicit restrictions on computer science or mathematics degrees are rare, and most projects do not impose hard technical requirements. This challenges the perception of AI as an institutionally closed domain accessible only to narrowly trained specialists.

However, this formal inclusivity coexists with implicit skill-based filtering. Programming competence and familiarity with computational frameworks are frequently embedded within project descriptions. As a result, access may be shaped less by degree classification and more by prior exposure to technical practices.

Overall, the findings show a visible transition from overt disciplinary exclusion to a system structured around assumed technical fluency, which may differentially advantage applicants with prior computational capital. This reinforces the core idea behind the project's intervention, which aims to remove barriers to AI studies by providing training that enhances their skills and self-efficacy.

# 4. Method

## 4.1. Methodological Approach

This study followed a pragmatic paradigm (Islam, 2022), allowing for the exploration of the social issue of EDI (Equality, Diversity, Inclusion) in AI postgraduate studies. The study used a mixed methods approach (Gillespie et al., 2024) to offer an understanding of socioeconomic and individual barriers to AI studies. Descriptive statistics were used to evaluate the intervention's impact through a numerical pre- and post-test comparison. To explore the effect of the training on student attendees and the factors influencing their decisions to pursue an academic career in AI, qualitative data was gathered using free-text responses in the surveys and subsequent structured interviews.

## 4.2. Data Collection

This study was approved by the University of Bristol Faculty of Science and Engineering's ethics committee (approval code: 29619). Pre- and post- testing was completed using two respective online surveys: the first was administered in the days before the training began, and the second was offered after attending the first of two training sessions for completion within a two-week window. Following attendance of the second training session, survey participants were then invited to attend a structured online interview.

This mixed-methods approach was used to answer the following hypotheses with quantitative survey responses and research questions with interview responses, based on sections of the literature review:

H1) Does the AI training programme increase attendees' interest in AI study? (Section 3)

H2) Does the AI training programme increase attendees' AI literacy? (Section 3)

H3) Is there any difference of interest in AI study based on SES? (Section 2)

H4) Do STEM students show higher AI literacy than non-STEM students before and after the training? (Section 2 and 3)

RQ1) What do participants perceive as barriers and enablers for them to study AI following attendance of the training programme? (Sections 2 and 3)

RQ2) How do career expectations influence their choices around postgraduate study after attending the training programme? (Section 3)

RQ3) How does institutional policy impact interest in AI studies after attending the training programme?

Participants who completed both surveys were provided a "thank you" voucher of £10. For participating in the interviews, they were provided an additional voucher of £20. The average completion time for the pre-training survey was approximately 25 minutes, the post-training survey took a mean time of minutes, and the interviews lasted *up to one hour*.

### 4.2.1. Sample

Data was collected from self-selecting participants from the group of students attending the "Getting Started in AI" training programme run by the Jean Golding Institute at the University of Bristol. The eligibility criteria for the surveys and interviews were as follows:

- Aged 18 or over.
- Undergraduate students.
- Studying at the University of Bristol.
- Attendees of the "Getting Started in AI" training.

To maximise the number of participants, all attendees whose place was confirmed on the training programme were contacted by email about optionally taking part in the survey. In addition, time was allocated within the training sessions to complete the pre- and post-training surveys. They were also reminded of the option to take part in an interview at the end of the second training session.

The two training sessions (1 and 2) were each delivered across two groups (cohorts A and B) due to room size constraining the maximum group size to 30 attendees in each session. The research team supervised each session and were satisfied that there were no significances in the delivery of the two sessions between cohorts A and B.

### 4.2.2. Survey Design

The pre-training survey comprised items regarding demographic data, AI literacy, student attitudes towards AI, and their expectations for the training. Appendix A contains the full pre-training survey for reference. The post-training survey differed by exclusion of the demographic questions and measuring experiences of the training post-hoc.

For participants' SES status and demographics, standard SES indicators were taken from the Diversity and Inclusion Survey (DAISY) guidance (Molyneaux and Hunt, 2022) for ethnic background, gender, parents' educational background, and student's current occupation. The ethnic group multiple-choice and write-in options were similar to those used in the Census 2021 (ONS, 2020). Questions in the latter two items were structured similarly to those which the research team have received from universities or students' unions in research that analyses socioeconomic indicators. Furthermore, the research team agreed that the selected items reflected the inclusive and equitable foundations of the project in their wording.

The AI Literacy Scale (Short Version) or "AI-LITS" (Hornberger et al., 2025) was used as-is in the pre-training survey to measure students' AI Literacy. To reduce the likelihood that students would recognise these test questions from the pre-training survey and answer without consideration, the question order was altered for the post-training survey.

Items created for this survey – not derived from existing sources - were reviewed by two members of the team for item validity. These included five-point agreement scales and open questions with free-text response fields to measure student attitudes towards AI, their motivation for/experience of attending the training, and interest in studying PhDs.

### 4.2.3. Interview Design

The design of these interviews was informed by the three research questions which provided a qualitative insight into student participants' attitudes towards studying and using AI. Appendix C contains the interview questions used in this study.

Structured, qualitative interviews were conducted to gather data. Structuring the interviews served to improve consistency across separate instances, which were conducted by different members of the research team, and enable more direct comparisons of participants' responses. Each interviewer was accompanied in a recorded Microsoft Teams meeting by an observing member of the research team. These observers took field notes so that the interviewer could focus on inquiry; their notes were later used to compare the similarity of interviews among different interviewers.

Each interview was automatically transcribed in Microsoft Teams at the time of recording. The transcripts were reviewed and corrected by hand by the respective interviewers to ensure accuracy before analysing the text.

## 4.3. Data Analysis

### 4.3.1. Quantitative Survey Analysis

Questions on SES and demographic were only asked in the pre-training survey because they could not change (ethnic background, caregivers' educational level) or were not expected to change (gender, students' occupation) in the one or two weeks between surveys. Other than this, the quantitative survey data was compared between the pre-training and post-training datasets.

Before conducting paired analyses, the initial dataset of 14 matched pre- and post-training surveys was cleaned to ensure data integrity. One participant was excluded for completing both surveys consecutively (two minutes apart), indicating they did not attend the training between test administrations. For a participant with multiple post-survey submissions, only the first complete submission was retained to prevent practice effects from skewing the data. This resulted in a final matched sample of n=13 for the paired analyses.

To evaluate the impact of the training on AI knowledge (H2 and H4), two complementary scoring methods were applied to the 10-item AI-LITS test:

1. **Raw Percentage:** A standard calculation of total correct answers out of 10, providing a straightforward measure of overall knowledge acquisition.
2. **Item Response Theory (IRT) Ability Score (θ):** To provide a more robust evaluation of underlying conceptual mastery, responses were also scored using a 3-Parameter Logistic (3PL) IRT model.

Using the fixed item parameters validated by Hornberger et al. (2025), the IRT model estimates a "true ability" score (θ) by weighting each question based on its specific difficulty, how sharply it discriminates between ability levels, and adjusting for a 25% guessing probability on multiple-choice items. Because the IRT model weights hard questions differently than easy questions, it provides a stricter evaluation than the raw percentages.

### 4.3.2. Qualitative Interview Analysis

For the research team to compare students' interview data, a deductive approach to thematic analysis was followed due to time constraints (Clarke and Braun, 2017). Codes were defined according to previous literature on the topic (Fife and Gossner, 2024) regarding AI policies and admission policies in universities, AI literacy, and socioeconomic barriers to AI studies as defined in Sections 2 and 3. Some additional codes were added based on responses in pilot interviews conducted with each member of the research team. This allows for the testing of previous claims in a newly emerging social phenomenon and to expand the knowledge on this topic according to new themes emerging from the empirical data (Fife and Gossner, 2024). Researchers conducted the analysis in a shared codebook (Appendix D) document where codes were generated and reviewed both asynchronously and in dedicated meetings.

It is important to note that the AI training delivered by the JGI focused on AI literacy rather than other real or perceived barriers to AI doctoral studies as identified in the literature review (above). The partial lack of alignment between the literature review and the AI training had an impact on data collection and analysis, as some previously identified codes were not addressed as part of the post-training interviews (e.g., gender inequality in AI studies and confidentiality concerns). Additionally, some aspects of the literature review could not be addressed because of the sample of undergraduate students; for instance, admission policies related to doctoral studies in AI could not be investigated. Further studies could be developed to provide a more in-depth understanding of the different, intersecting dimensions that may impact doctoral studies in AI.

### 4.3.3. Qualitative Survey Analysis

The free-text responses from the surveys were analysed using abductive content analysis (based on iterative steps described by Erlingsson and Brysiewicz, 2017). Due to the brevity of most free-text responses, deep immersion and understanding of the data would have been difficult to achieve, so content analysis enabled meaning to be efficiently categorised across the dataset (Humble and Mozelius, 2022). The predefined codes generated for thematic analysis of the interviews were used in an initial deductive process, but due to the variety of responses, additional codes were added in successive passes through the data.

These qualitative approaches provided descriptive insights into how the training programme may have altered attendees' attitudes towards AI and further study. These complemented the quantitative survey data to understand a broader range of rationale for students' responses, rather than provided explanations for the impacts of the training on individuals.

# 5. Results

## 5.1. Quantitative Survey Sections and Interviews

For clarity and ease of interpretation, both the Raw Percentage and IRR Ability Score (θ) mentioned in Section 5.4.1 are presented in this analysis. Before conducting the paired analyses, the initial dataset of 14 matched pre- and post-training surveys was cleaned to ensure data integrity. One participant was excluded for completing both surveys consecutively (two minutes apart), indicating they did not attend the training between test administrations. For a participant with multiple post-survey submissions, only the first complete submission was retained to prevent practice effects from skewing the data. This resulted in a final matched sample of n=13 for the paired analyses.

Data from the interviews are provided alongside the quantitative analyses to provide descriptive context for the results. Results of the content analysis are presented in the following section as they reveal a greater variety of factors which do not all directly relate to the hypotheses here.

### 5.1.1. Pre- and Post- Training Interest in AI Study

**H1: Does the AI training programme increase attendees' interest in AI study?**

To evaluate shifts in student interest regarding future AI studies, pre- and post-training survey responses were analyzed using the non-parametric Wilcoxon Signed-Rank Test, which is appropriate for ordinal Likert-scale data. Two proxy measures were evaluated: student confidence in being offered a PhD place, and direct willingness to apply for an AI PhD.

When asked to rate their agreement with the statement, *"I think I would be offered a place to study a PhD in AI if I applied"* (1 = Strongly Disagree to 5 = Strongly Agree), the median score remained static at 2.0 (Disagree) before and after training. The mean score shifted positively from 2.15 to 2.54, which a Wilcoxon test indicated is a weak trend approaching, but not reaching, statistical significance ($W = 18.0$, $p = 0.094$). A similar pattern emerged regarding the direct question, *"Would you consider applying for a PhD in AI?"* (0 = No, 1 = Maybe, 2 = Yes). The median response remained firmly at 1.0 (Maybe). While the mean score rose slightly from 0.77 to 0.92, this change was not statistically significant ($W = 10.0$, $p = 0.375$). Collectively, the quantitative data suggests that while the training successfully increased objective AI literacy (as seen in H2), a short technical intervention alone is insufficient to significantly alter deeply held career intentions or overcome confidence barriers regarding PhD admissions.

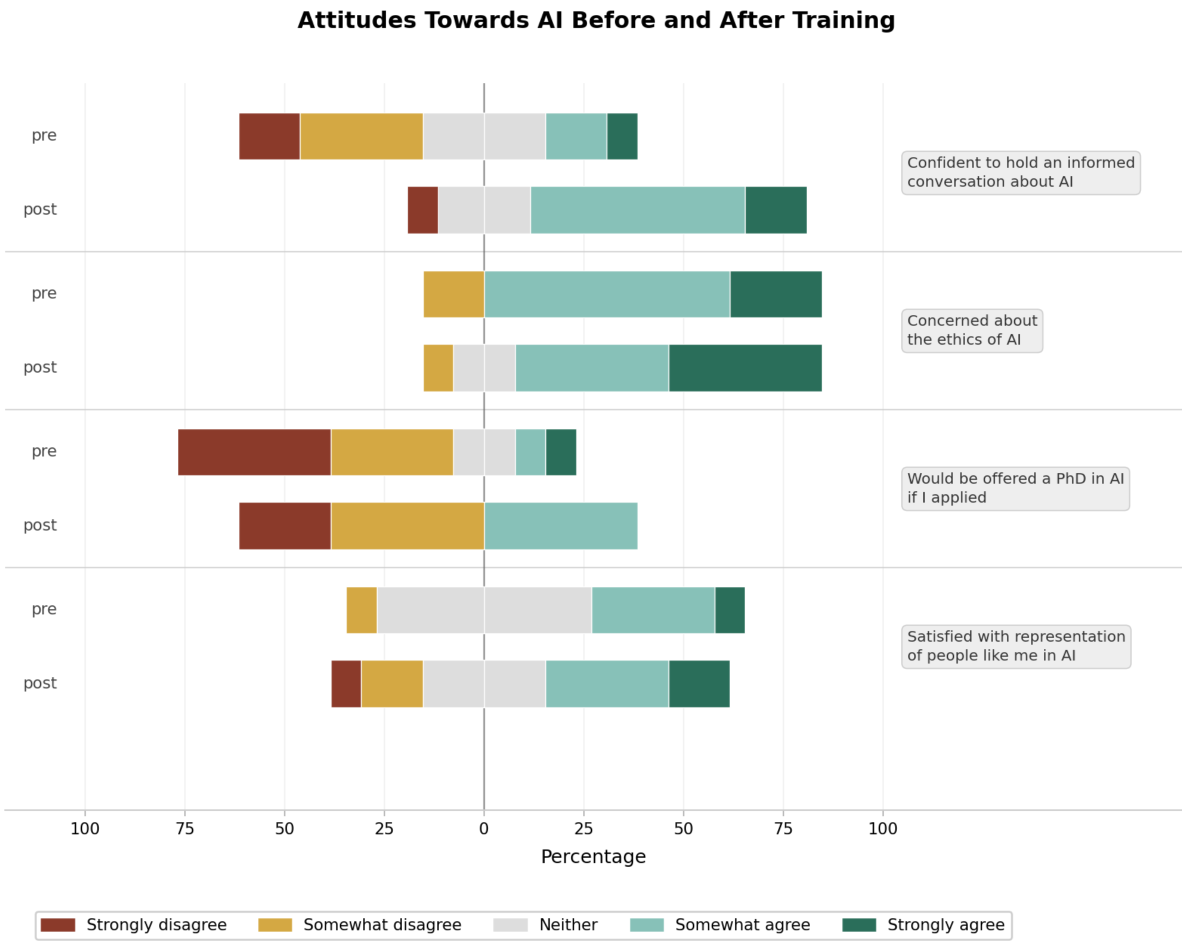


*Figure 21. Pre- and post-training shifts in AI attitudes (n=13). Diverging stacked bar chart showing a marked increase in confidence to discuss AI, alongside positive directional trends in perceived PhD eligibility and representation.*

These trends are further explored during interviews with participants, who outline explanations related to their current field of study, the limitations and opportunities of the job market, and personal career decisions.

For some interviewed participants, engaging with AI as a user of these systems was more important than studying AI. The application of AI was not clear to them in their fields, e.g., social sciences and marketing. This is also connected to the skills they have obtained in this pathway, which are less hard-science oriented. For these participants, the concern was to use AI ethically, understand its many applications in their work, and have the knowledge to make full and optimal use of it.

> “I don't know if taking a career, but learning more about it, getting more into more courses like this one”

While future careers may not be explicitly linked to barriers to doctoral AI studies, all participants in the present research project talked about their career expectations as part of their chosen degree and potential post-graduate studies. The ambiguity about future careers was one of the factors influencing the decision to pursue AI postgraduate studies. Concerns varied between the competitiveness of the field, uncertainty regarding the needed expertise within it, and the lack of understanding of how it connects to areas such as the social sciences. For participants with an engineering background, the concern was about the limited number of job opportunities in AI in comparison to a more general engineering specification:

> "I chose [this degree] because...I thought it's very broad here in England and I can go into any industry that I want. They will always need a mechanical engineer for something"

Participants clearly addressed that the application of AI is not well-known in most fields. More clarity about careers in AI was a type of support they perceived as needed to make this decision.

One interviewed participant shed light on more competitiveness in the job markets for migrants, especially in highly paid jobs.

> "There's always the risk of something being true right now, but then completely changing in a year's time because of the nature of the field where everything is changing rapidly"

However, for some participants, exposure to AI technical skills and concepts was important for discussing future careers. For others, it was eye-opening to the potential within a pathway of their interest.

> "I feel like it would depend on my future career choices, because right now I'm in between two options".

> "Yeah, I'm interested in financial, the financial side of marketing. So that's pretty much related to AI in terms of data analysis, so if I were to choose that, I think I would get more into AI"

> "So I felt like I didn't know enough about this whereas I should, especially because I want to pursue this in the postgraduate level […] I think there was a lot of space in robotics to incorporate artificial Intelligence and make autonomous robotic systems"

An important theme that emerged was the willingness to engage with AI in fear of future career competitiveness. Interviewed participants were cautious about the need for skills in AI for their future employment. Despite not knowing what the job market would look like, they believe that using AI systems is now essential for the future. This raises the question of whether the absence of this scenario would impact their willingness to enhance their knowledge and skills related to AI?

> "Because it would be like a major part of my future, like where I would be in my career in the future"

> "AI is very, very important nowadays and it is going to get more and more important"

It is also important to note that one participant implied a perception of an AI career in the UK being mostly open to UK population only:

> "And actually, one of the guys [delivering the training] was [from abroad], so we had a chat in [their home language] as well...I'm not from England, so that's really good to know that anyone can get into it"

The above quote may imply a degree of intersectionality between race/ethnicity/migration and chosen career path, hence higher education studies. However, the small number of participants and the narrow focus of the study (due to time constraints) mean that more studies should be conducted to further understand the above intersections.

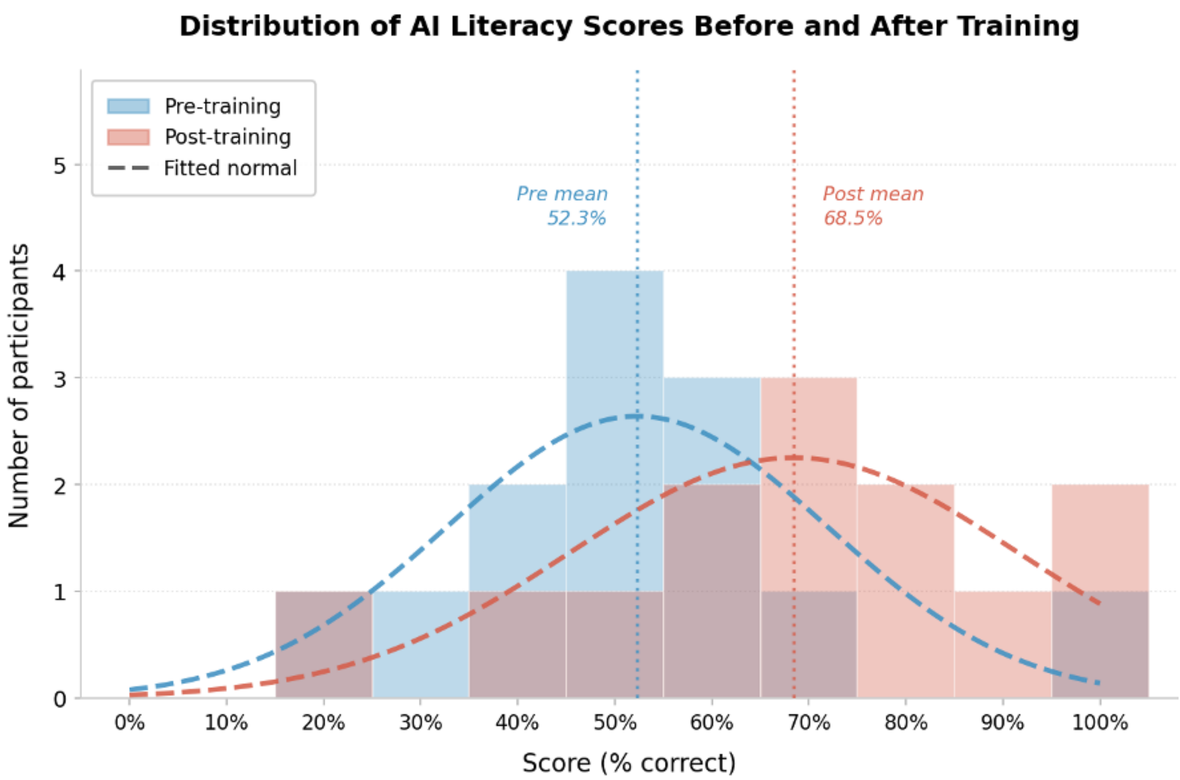


*Figure 22. Raw AI literacy questionnaire scores before and after training. This histogram illustrates a significant 16.2% increase in the mean score (from 52.3% to 68.5%) for the 13 participants, indicating clear knowledge acquisition.*

### 5.1.2. Pre- and Post- Training AI Literacy

**H2: Does the AI training programme increase attendees' AI literacy?**

Analysis of the 13 matched pairs revealed a strong positive trend following the training intervention. When evaluating the simple raw percentage scores, participants improved by an average of 16.2 percentage points—shifting from a baseline mean of 52.3% to a post-training mean of 68.5% (see Figure 22). A paired-samples t-test indicated that this raw improvement was statistically significant ($p = 0.024$), confirming the training successfully imparted measurable knowledge of core AI concepts.

When evaluating the IRT model ($\theta$), the mean true ability estimate demonstrated a substantial positive shift, increasing from 0.094 to 0.586, representing nearly half a standard deviation (see Figure 23). Because the 3PL model is highly sensitive to individual response patterns, it introduces higher variance than raw scoring, rendering the analysis statistically underpowered at n=13. Nevertheless, the strong directional alignment between the significant raw percentage increase and the positive $\theta$ shift supports the overall efficacy of the intervention.

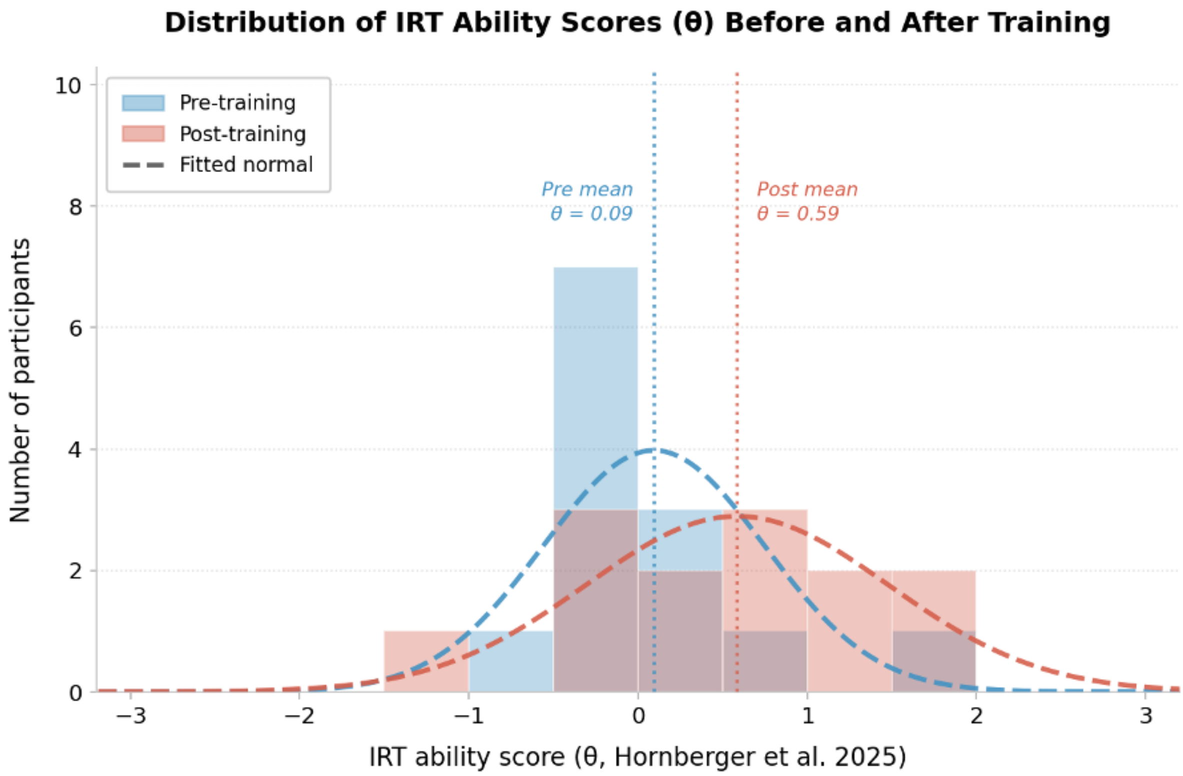


*Figure 23. IRT ability scores from the AI literacy questionnaire. Based on the 3PL IRT model, this histogram shows the shift in mean ability from 0.09 to 0.59 (n=13), representing growth of nearly half a standard deviation.*

The interviewed sample emphasised the importance of technical skills as a factor in their evaluation of what encourages or discourages them from engaging with AI studies.

Being knowledgeable of the general concepts, such as LLMs or neural networks in addition to coding skills, had a positive influence on participants' attitudes towards engaging more with AI as users in areas that might be related to their studies.

### 5.1.3. Interest in AI Study with Socioeconomic Factors

**H3: Is there any difference of interest in AI study based on SES?**
The self-selecting nature of the training attendees revealed a severe demographic skew: 12 of the 13 matched participants fell into the Higher SES category, leaving only 1 participant in the Lower SES category. Due to this extreme imbalance, formal comparative hypothesis testing is impossible.

While statistical conclusions cannot be drawn from a sample size of one, the fact that only a single first-generation student completed the AI training pipeline highlights a critical systemic issue regarding accessibility and early-stage engagement in AI education. As a key lesson learned, future research should employ a purposive sampling strategy to actively recruit students from diverse socioeconomic backgrounds, which would yield much stronger and more robust findings on this front.

*Table 5. Pre- and post-training AI literacy score changes by subject area. This table displays the mean raw percentage scores and the average percentage point (pp) change for STEM and Non-STEM cohorts. The STEM cohort started with a higher baseline score and gained more improvement from the training.*

| Group | n | Mean pre | Mean post | Mean change |
|---|---|---|---|---|
| **STEM** | 6 | 56.7% | 78.3% | +21.7pp |
| **Non-STEM** | 7 | 48.6% | 60.0% | +11.4pp |
| **All** | 13 | 52.3% | 68.5% | +16.2pp |

The conducted interviews were able to discuss only the aspect of accessibility to capacity building opportunities in terms of SES-related factors and in a specific group of undergraduate students studying at the University of Bristol.

In reflecting on the training the participants attended, they highlighted some aspects of accessibility to these opportunities. They believe that being part of the university allowed them to access a few trainings that have better introduced them to AI. While one of the training sessions some of the students attended focused on AI ethics, the "Getting Started in AI" training was more technical. Interviewed participants value this exposure as they identify its impact on their awareness of potential applications of AI within their current studies and the future. In general, the interviewed participants appreciated the technical development the training offered. Even though that was not a breaking factor in their decision to pursue AI studies, it offered awareness of the field and what it might include for students who were considering a related postgraduate degree.

"I also took the course to kind of help me decide whether I'd like to pursue a Masters of Science in AI rather than Masters of Science and Engineering, which was my original path".

indicated that being invited through email and posters, promotion of a free training, and being in a friendly environment encourages people to engage with AI more. This will allow people to be better informed and technically equipped, increasing confidence in pursuing a degree in AI.

Interestingly, the reflection on accessibility did not expand to minorities and disadvantaged people. This can be contextualised within the biases of the sample this study engaged with.

### 5.1.4. AI Literacy with Course Type

**H4: Do STEM students show higher AI literacy than non-STEM students before and after the training?**

To understand if prior academic background influenced AI literacy, the matched sample (n=13) was divided into STEM (n=6) and Non-STEM (n=7) cohorts based on their enrolled degree programs.

Prior to the training, STEM students exhibited a higher baseline AI literacy than their Non-STEM peers. The STEM cohort achieved a mean raw score of 56.7% (mean $\theta$ = 0.251), compared to the Non-STEM cohort's mean raw score of 48.6% (mean $\theta$ = -0.041). This confirms that STEM students enter with a slight measurable advantage in AI literacy.

While both groups benefitted from the intervention, STEM students experienced substantially higher growth. As visualized in Figure 3, the STEM cohort's raw scores surged by +21.7 percentage points to a final mean of 78.3% ($\theta$ increase of +0.816). In contrast, the Non-STEM cohort saw a more modest increase of +11.4 percentage points, finishing with a mean of 60.0% ($\theta$ increase of +0.215). Due to the small sub-sample sizes, within-group paired t-tests did not reach statistical significance. However, the descriptive data clearly indicates that while the training successfully elevates literacy across all academic backgrounds, it appears to act as a stronger catalyst for students already embedded in STEM disciplines, effectively widening the literacy gap between the two groups rather than closing it.

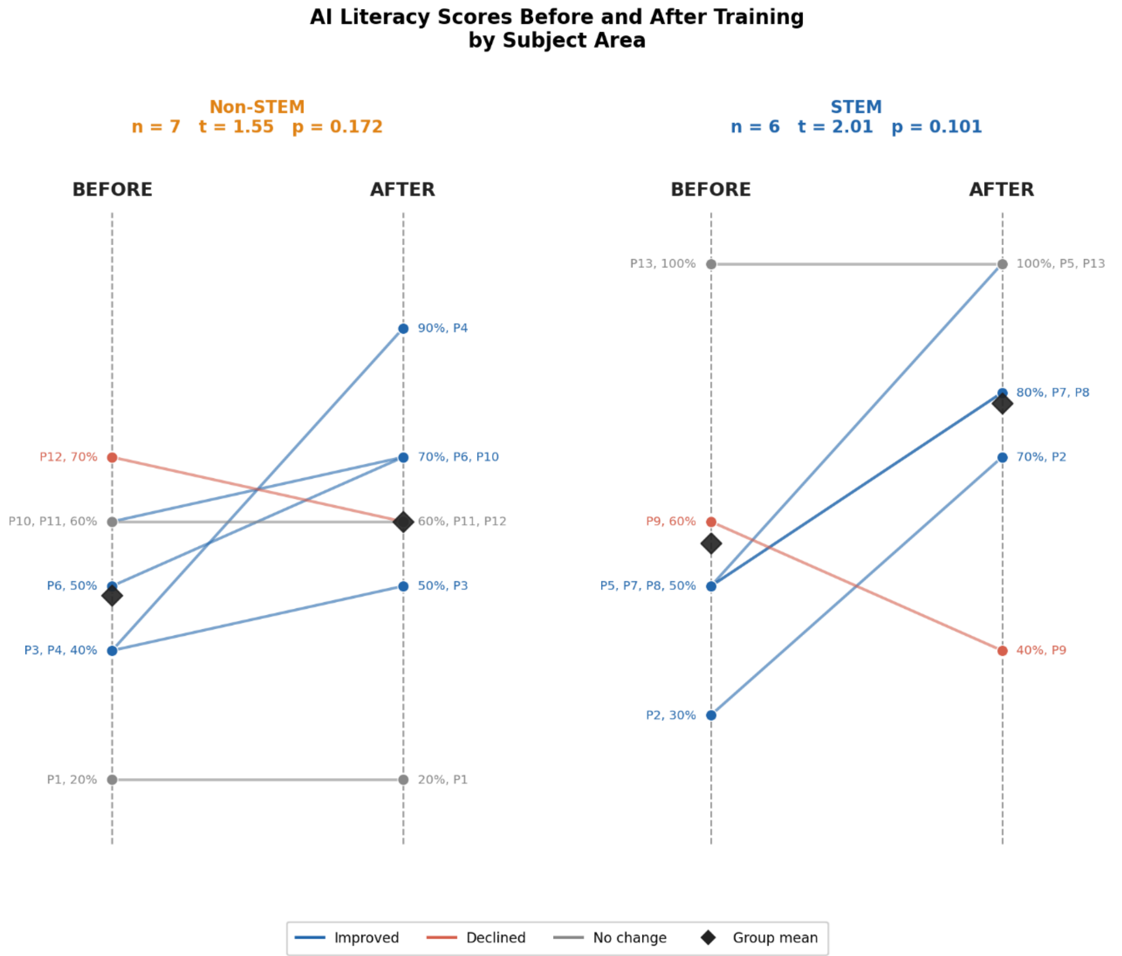


*Figure 24. Individual and group shifts in AI literacy by subject area. This slope chart compares the pre- and post-training scores for Non-STEM (n=,7 left) and STEM (n=7, right) students. The STEM group began with a higher mean baseline (56.7%) than the Non-STEM group (48.6%). The blue lines indicate improvement, showing that while both groups gained knowledge, the STEM cohort experienced a more pronounced average increase.*

The qualitative data explains these patterns in more detail, expanding on the interrelated themes of technical skills, training evaluation, and attitude towards AI. In addition, having previous coding skills and experience helped students feel more confident in expanding their knowledge about AI and considering further studies in it. Some participants expressed feeling less inhibited by the idea after taking the hands-on session that introduced them to some practical aspects of coding. Also, they emphasised the importance of understanding the mechanism behind AI and how it works, which helped them make better connections between it and the work that they do. Additionally, participants in the training particularly appreciated having AI substreams broken down for an easier learning experience.

> “So the way I use LLMs was also affected in terms of when realising and understanding how something works. You know how to better use it based on its underlying mechanism”

The change in perspective included the knowledge and skills needed to study AI. Participants declared that they were under the impression of a deep knowledge of maths and coding, which inhibited their engagement with AI. This perception changed following the training.

> "I come from an engineering background, so I've done some Python, and I know like some coding. If I read some code, I would probably understand what it's doing"

On the other hand, coding skills and limited knowledge of AI were factors that made participants less willing to further explore the training, let alone develop further interest in postgraduate studies in related topics.

> "Maybe it's because of like I don't have that technical background of it. So like I got the base of it, but still like I'm not- probably I won't go for an AI postgrad"

It is important to consider that these skills and knowledge are developed over the years through previous decisions in schools and undergraduate programmes and enhanced by personal interest and environmental influence.

> "Back in my high school, like uh, it's quite similar to the A-levels in the UK. So I took like Politics, Geography, History and Economics. I never got any chance to like study like Python and like stuff like that because it seems so complicated to like a person from social science major to study"

## 5.2. Free-Text Survey Responses

The content analysis of the free-text responses in the surveys revealed the frequency of different topics mentioned by the participants pre- and post-training, according to the following prompts:

A) Please list up to three things that interest you in studying AI.
B) Please list up to three things that interest you in undertaking a PhD.
C) Please write up to 100 words about [what has motivated you to take / your experience of taking] part in the AI training offered by the Jean Golding Institute.

The responses provide an insight into factors which contribute to participants' interest in studying AI (supplementing H1 further) and their motivations in AI study which may be useful in overcoming barriers (addressing motivational aspects of RQ1).

The frequency of prominent topics mentioned in response to these questions in the pre- and post- training surveys are displayed in Figures 25, 26 and 27. The results are presented in order of the highest to lowest percentage of responses to the pre-training survey containing each topic. It is worth noting that the titles used to represent the different ideas that interest students in studying AI, undertaking a PhD and taking part in the training programme are named “topics” here to distinguish them from the “codes” derived from the interviews.

### 5.2.1. Interest in Studying AI

All the topics featured in responses from the pre- and post-training survey question on participants’ interest in studying AI are listed below in Figure 25.

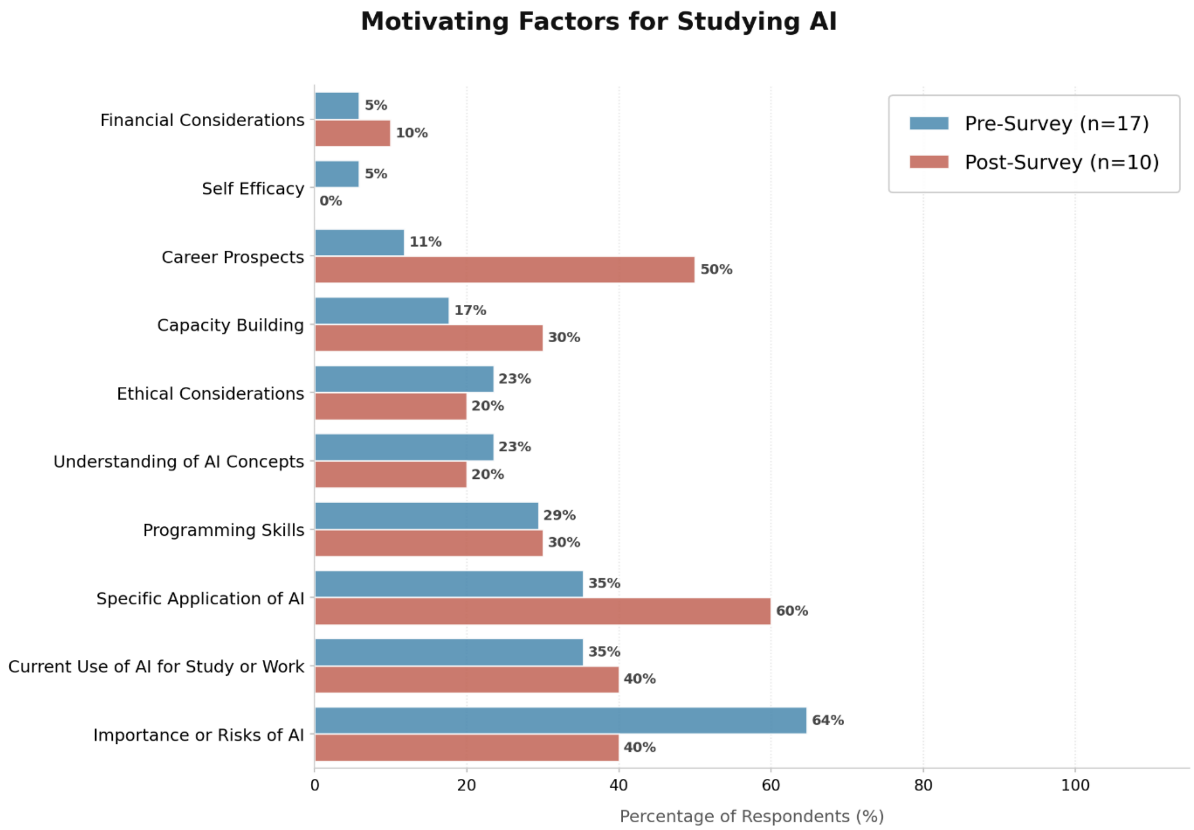


*Figure 25. Bar chart displaying the percentage of responses which feature the various factors that interest participants in studying AI.*

The “Importance or Risks of AI” related to respondents’ concern that learning AI is necessary for preparing for future careers or otherwise “futureproofing” as well as the inverse: that not understanding AI risks negative impacts on one’s aptitude for future careers. This was the most frequently stated motivator for participants before they underwent training (included in 64% of responses), though the proportion of respondents mentioning the “Importance or Risks of AI” in the post-training surveys reduced to 40%.

Participants expressed the belief that AI study would further their ability to use it in study or work in the near future (“Current Use of AI for Study or Work”) in 35% of responses prior to the training, rising to 40% of responses afterwards. They also frequently described “Specific Applications of AI”, especially in data science, as motivating their interest to study AI. Prior to the training, 35% of respondents mentioned this factor while after the training it featured in 60% of responses.

"Programming" was also commonly mentioned by respondents - 29% pre-training and, similarly, 30% post-training. Both before and after the training, this mostly related to respondents' expressions of interest in improving coding skills or understanding how to create AI tools. These were often connected with needing to learn to programme AI in order to use it better, and so "Current Use of AI for Study or Work" and "Programming" were often paired in responses.

"Ethical Considerations" and "Understanding of AI Concepts" were both cited in 23% of responses prior to the training and featured in 20% of responses afterwards. Before the training, participants who were concerned with the ethical – including "societal" - impacts of using AI tools often expressed these as the focus of their motivation for studying AI. They mostly provided explicit recognition of a need to understand the facets of how AI tools operate in order to understand these ethical impacts comprehensively. There were also responses where the importance of understanding how AI works was related more closely with the development of AI or applying it in specific areas; it was not always described in relation to ethical concerns.

Fewer respondents referred to a need for "Capacity Building" (17%), learning AI to improve "Career Prospects" (11%), wanting to improve confidence in using AI either now or in future ("Self Efficacy", 5%), and that studying AI may relate to finances in some way ("Financial Considerations", also 5%) prior to the training. Notably, "Capacity Building" and "Career Prospects" rose significantly after respondents' attendance to 30% and 50% respectively, further suggesting that as a result of the training, their interest in AI became more orientated towards skills acquisition and preparing for future careers than concerns around its use, which was more dominant in the pre-training surveys.

### 5.2.2. Interest in Studying a PhD

The topics arising from respondents' interest in studying a PhD before and after attending "Getting Started in AI" are summarised in Figure 26 below.

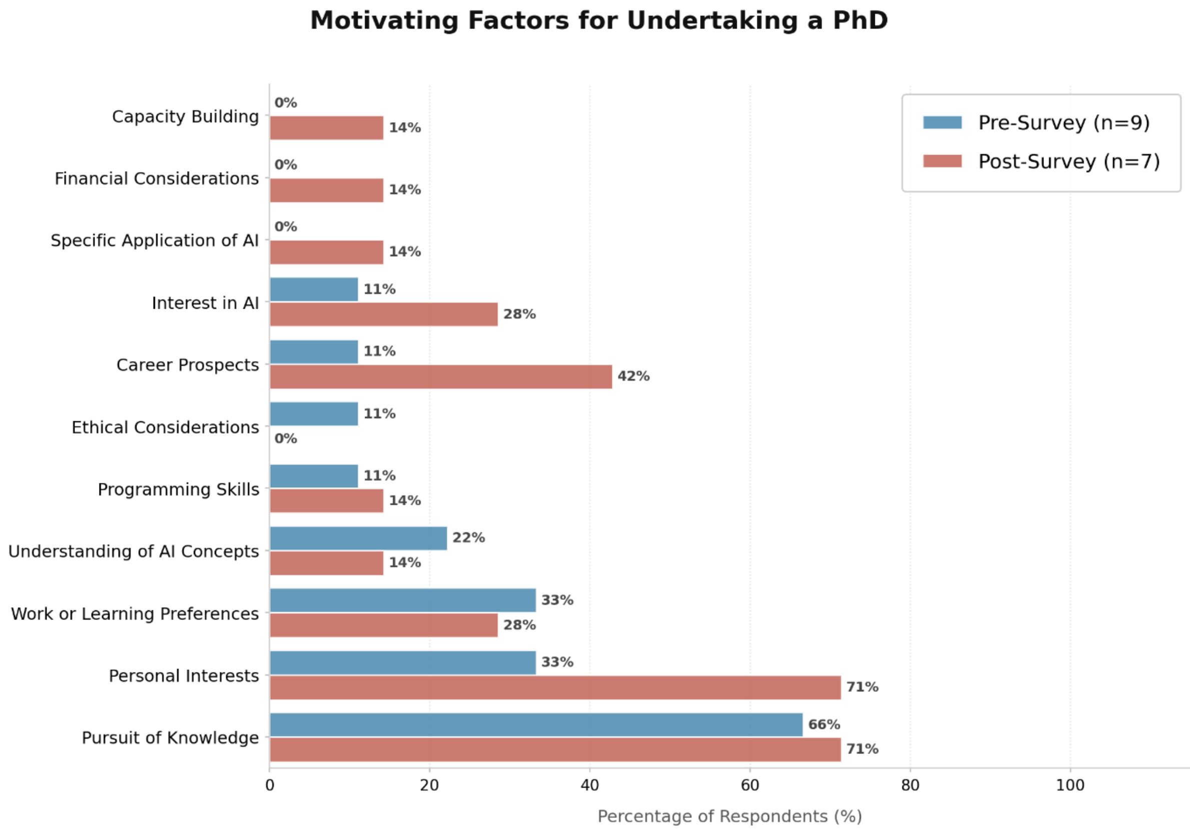


*Figure 26. Bar chart displaying topics related to participants' interest in undertaking a PhD by the percentage of responses in which they featured.*

Motivational factors for undertaking a PhD were more closely aligned with knowledge acquisition, personal interests and workstyle preferences than expressions of interest on what encourages them to study AI.

The "Pursuit of Knowledge" was mentioned frequently both before and after the training sessions (mentioned by six respondents, 66%, before and six, 71%, after the training). This topic was characterised by respondents' interest in developing knowledge, which was distinct from wanting to acquire functional knowledge that would improve their ability to use AI (which would relate to Using AI for Study or Work). Participants described wishing to study a PhD in order to be able to engage in prolonged, deep investigations of interesting topics, mostly without mentioning a specific field they wished to pursue. This was closely connected in the post-training responses with the idea that pursuing knowledge in this way would enable them to follow personal interests (hence they both featured in 71% of post-training responses), though only half the participants mentioning the "Pursuit of Knowledge" linked this explicitly with their "Personal Interests".

Respondents communicated an awareness that PhD study would enable them to work on long-term projects at a very high level, at rates of 33% pre-training and 28% post-training, which related to them displaying "Work or Learning Preferences".

Two students (22%) mentioned wanting to understand AI through studying a PhD before the training, one of which (11% for "Programming") described wanting to develop programming skills as part of this and having specified AI as a topic of interest (11% for "Interest in AI"). These topics were maintained by one student following the training as well (14% for both "Understanding of AI Concepts" and "Programming").

One participant (11%) related PhD study to "Career Prospects" before the training but three (42%) mentioned "Career Prospects" in their post-training response. This distinction was marked by descriptions that implied a recognition of undertaking a PhD as an alternative to entering other work after graduating from their course. Unlike perspectives on studying AI, "Career Prospects" were considered more in terms of what students desired to pursue rather than what preparations they must make for future work.

### 5.2.3. Attendance of the Training Sessions

Figure 27 below shows the frequency of various topics featured in responses to what participants anticipated before the training and in their reflections on the training.

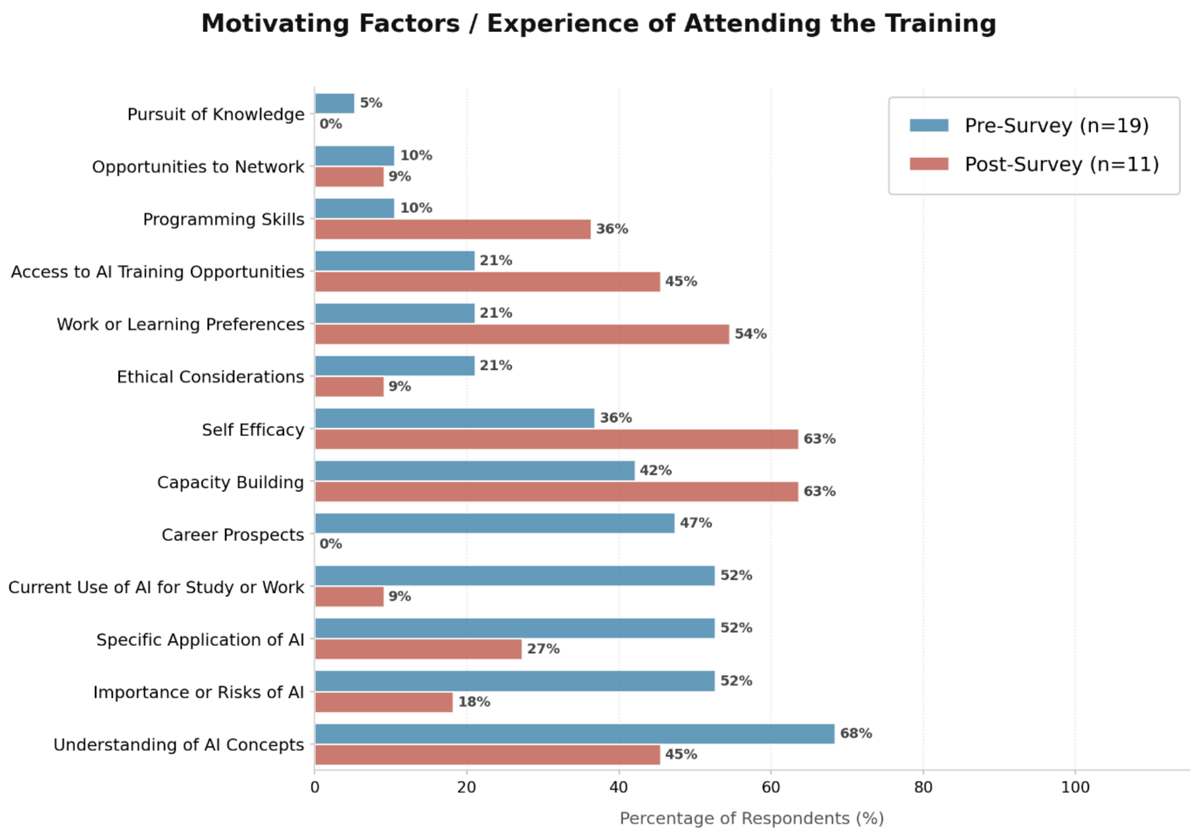


*Figure 27. Bar chart displaying the frequency of topics in responses to motivations for (pre-) and experiences of (post-) taking part in the Getting Started in AI training.*

The initial motivations for participating in the training were primarily driven by desires to develop their technical understanding and use of AI, as well as personal growth.

“Understanding of AI Concepts” (68%), alongside “Importance or Risks of AI” and “Specific Applications of AI” were the most commonly mentioned motivators for undertaking the training. They primarily relate to the sense communicated by respondents that AI is necessary for “future-proofing” or preparing for future careers. By reflecting on their potential future career paths, students referred to their respective fields in which they see increased AI use. This is also supported by the “Career Prospects” theme that appeared only in the pre-survey responses (47%), suggesting that despite the personal motivations, the training functioned as a bridge to support students in developing the technical skills necessary for their future success.

This effect can be observed in the reported increase in “Self-Efficacy” between the (pre-) and (post-) training surveys, from 36% to 63%. Here, students reflected on the training as a starting point for increasing their belief in their skills, supporting their AI learning journey, rather than as a final destination. “Capacity Building” was frequently mentioned in the post-training reflections, with students emphasising the training's role in providing valuable learning opportunities and overall accessibility to help them navigate the AI field. There was an increase in the proportion of respondents referred to the accessibility of the training as “Access to AI Training Opportunities” (including the value of incentives in their decision to attend): 21% before and 45% after. Moreover, many respondents communicated their enjoyment or appreciation of the applied, “practical” learning of AI concepts which were not as commonly considered before the sessions. This is reflected in the rise of “Work or Learning Preferences” from 21% to 54%.

Other than this, “Ethical Considerations” were more common in the pre-training surveys than post (in 21% versus 9% of responses) but other factors saw an increase in the post-training survey. “Programming Skills” and “Opportunities to Network” both saw mentions by 10% of respondents before the sessions with “Programming” increasing to 36% of responses post-training (which is as expected given the technical focus of the syllabus), while references to networking or learning with others reduced to 9% post-training. One participant (5%) expressed being motivated by “Knowledge Pursuit”: “I would love the opportunity to (learn about AI capabilities and ethics) to boost my knowledge” but no responses mentioned “Knowledge Pursuit” after the training in line with the definition provided in Section 5.2.2.

### 5.2.4. Discussion

***Participants’ AI Literacy Awareness***

The appreciation for developing AI skills found in the free-text responses supports that participants perceive skills and knowledge gaps relating to Long and Magerko’s (2020) AI literacy competencies. Some examples from the survey responses include “coding” relating to 17) Programmability, “ethics” clearly reflecting competency 16) Ethics, or a more complex example being “understanding... if (AI’s) objectivity ultimately depends on human design and development” comprising elements of 8) Decision-Making, 10) Human Role in AI, 12) Learning From Data, 13) Critically Interpreting Data and 16) Ethics.

***Self-Efficacy Trends***

The apparent increase in self-efficacy in 36% of pre-training survey and 63% of post-training survey responses for the third free-text question relates to findings from previous literature in Section 3.1.3 (Cowit and Fiesler, 2024). As such, the improved confidence in AI skills could aid participants in pursuing AI study if they so wished by reducing the likelihood of disengagement from learning AI earlier in the pipeline.

***Motivation to Improve Personal AI Usage***

The relevance of studying AI to survey respondents' immediate needs was a primary driver for attendance of the training sessions, as indicated by the significant percentage of citations (35% of pre-training respondents and 40% of post-training respondents). Respondents' interest in AI revolving around using it for current study or future work is consistent with Yin and Goh's (2024) observation that a key driver for learning AI due to its utility rather than pure technical interest in AI concepts.

BELOW ARE SECTIONS ON POLICY TO BE DISCUSSED.

1. **Perceived Policy:**

This study cannot claim a thorough and structured policy evaluation. However, it aimed to understand how the UoB AI policy was perceived by students.

Most interviewed participants think that the policy is clear. The examples given in the policy addressing the limits of using AI offered guidance for students. However, for others, they indicated that they would appreciate more clarity. These perspectives were related to exposure to the policy. Participants who were familiar with the guidebook were more positive about the clarity of the policy.

"Seeing that it had been considered by the university as well and it wasn't that black or white, yes or no type of situation, encouraged me. I was happy to see that"

"Because I've always been used to not using AI as much for essays or exams
So yeah, it basically was the same thing for me"

In addition, the exposure to the policy through lecturers, emails, and personal curiosity influenced the understanding of it. The answers fell on a spectrum of being too strict to a fair policy that respects academic integrity without inhibiting using AI. This will be further discussed in the next paragraph.

**4.1. The impact of the policy on using AI:**

For most students, their usage of AI was not affected by the policy. This was especially evident with students who used AI less than what the policy allows for. However, they indicated that they are aware of other students who are using AI less because of the policy. One of the interviewees highlighted how more flexibility with the policy can help international students with language barriers

A rare input on that matter mentioned that the restrictive use of AI in the university raises suspicions about the ethical considerations in using AI outside the university for work or personal matters.

"Sometimes I feel a bit limited because, for coursework obviously I'm not going to use it like to copy and paste or anything, but just reading this policy […] So even if I want to use it for personal use or to do something else, I'll always be thinking of, oh, what if there's a policy against what I'm doing?...I don't feel comfortable in the future using AI because I know there are all of these limitations"

However, there were no indications that the policy discourages students from pursuing a postgraduate degree in AI.

It is important to note that admission policies were not covered in our question due to sampling constraints. However, the literature on this topic addresses that thoroughly.

# 6. Conclusion

For convenience, the hypotheses and research questions are restated here:

H1) Does the AI training programme increase attendees' interest in AI study?

H2) Does the AI training programme increase attendees' AI literacy?

H3) Is there any difference of interest in AI study based on SES?

H4) Do STEM students show higher AI literacy than non-STEM students before and after the training?

RQ1) What do participants perceive as barriers and enablers for them to study AI following attendance of the training programme?

RQ2) How do career expectations influence their choices around postgraduate study after attending the training programme?

RQ3) How does institutional policy impact interest in AI studies after attending the training programme?

This study provides a mixed-methods evaluation of a short AI training programme delivered by the Jean Golding Institute (University of Bristol, UK). It offers insight into how AI training may shapes students' AI literacy, perceptions, and educational aspirations. The study aimed to understand how AI training may impact real or perceived barriers to AI doctoral studies. By combining mixed-methods surveys with qualitative interviews, the findings highlight a consistent pattern: training is effective in increasing AI knowledge and literacy but is not alone sufficient to transform longer-term educational pathways, particularly for postgraduate study.

The quantitative findings demonstrate a clear improvement in AI literacy following the training (supported by both raw scores and IRT-based measures). These insights are reinforced by qualitative insights which show that increased familiarity with key concepts and technical skills enhances participants' confidence in engaging with AI, although primarily as users within their existing discipline, rather than towards a potential PhD. The aforementioned AI literacy enhancement does not, therefore, necessarily translate into an increased interest in pursuing AI at the postgraduate level. Participants frame AI engagement as a means of potentially improving employability, rather than as an academic pathway.

A key contribution of this study is that it helped identify the gap between knowledge acquisition and postgraduate educational aspirations. Across the qualitative data, three interrelated factors emerge: the perceived relevance of AI to one's current field of study, confidence in technical skills (particularly coding), and uncertainty about career pathways in AI. For non-STEM students in particular, AI continues to be perceived as implying a shift in competencies, limiting, therefore, willingness to pursue further AI study, despite learning outcomes. In contrast, STEM students not only begin with higher baseline literacy but may also benefit more from the training, suggesting that AI training may (inadvertently) reinforce existing inequalities, rather than reducing them.

The findings also reveal important issues of accessibility and participation. The highly skewed sample, with minimal representation from lower socioeconomic backgrounds, highlights a critical limitation, which is certainly related to the time constraint and small sample, but also an important insight: engagement with AI training opportunities may already be unequal at the point of access. However, while participants seemed to value the relatively quick and easy availability and delivery of the training, their reflections rarely, if ever, extended to broader structural inequalities, underscoring the need for more inclusive recruitment strategies and targeted outreach in future research and practice.

This study also shows the impact of contextualisation. Participants frequently expressed uncertainty about how AI is linked to their discipline and career aspirations, suggesting that (technical) training alone may not be sufficient without clear links to the ‘real-world’ (i.e., career and/or disciplinary relevance). By addressing ‘real-world applications’, deeper and more sustained engagement with AI training and studies may, therefore, be possible.

Analytical limitations must be acknowledged: the small sample size limits the generalisability of the above findings and constrains statistical analysis, particularly when it comes to socioeconomic differences and determining effects of the training on AI interest. The short timeframe between pre- and post-training measures means that only immediate effects could be measured. There is no consideration for longer-term changes in AI-related perceptions, and/or skills, and/or behaviour. Additionally, the focus on undergraduate students limits the ability to explore questions related to AI doctoral study, particularly in relation to admissions.

Future research could adopt a longitudinal design and a more inclusive, purposive sample to better understand how AI engagement evolves over time and across diverse populations. Purposive sampling with underrepresented groups, and current or prospective PhD candidates who aspired to but could not pursue AI studies, would enable a more robust and comprehensive understanding of ensuring equitable access to AI doctoral studies, especially if paired with training that clearly and explicitly highlights career and disciplinary application.

To summarise, this study demonstrates that AI training may be an effective tool to improve AI literacy, but, in isolation does not address the broader barriers that shape students’ engagement with AI as a field of study. Addressing real and/or perceived structural and institutional barriers to AI doctoral studies may require a more holistic approach, accounting for both literacy development and contextualisation, while also considering inclusivity and support for diverse learners.

# Disclaimer: Use of AI

**(ChatGPT, Undermind, Google's AI-powered literature search tools)**

Authors used AI tool to summarise their spoken reflections; form the structure of the second section which was adapted through discussion; identify some relevant literature and pioneers in the field; locate relevant publications, obtain Harvard referencing guidance, and polish the final version to address different writing styles.

Each of these stages incorporated critical questioning and review by at least one author to ensure that the text accurately reflected the authors' knowledge, as well as to manually audit prose to be in the authors' own words.

# References


Agrawal, M. *et al.* (2022) 'Large language models are few-shot clinical information extractors', in Y. Goldberg, Z. Kozareva, and Y. Zhang (eds) *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. *EMNLP 2022*, Abu Dhabi, United Arab Emirates: Association for Computational Linguistics, pp. 1998–2022. Available at: https://doi.org/10.18653/v1/2022.emnlp-main.130.

Atkinson-Toal, A. and Guo, C. (2024) 'Generative Artificial Intelligence (AI) Education Policies of UK Universities', *Enhancing Teaching and Learning in Higher Education*, 2, pp. 70–94. Available at: https://doi.org/10.62512/etlhe.20.

Banh, L. and Strobel, G. (2023) 'Generative artificial intelligence', *Electronic Markets*, 33. Available at: https://doi.org/10.1007/s12525-023-00680-1.

Barretto, D. *et al.* (2021) 'Exploring Why Underrepresented Students Are Less Likely to Study Machine Learning and Artificial Intelligence', *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education V. 1*. New York, NY, USA: Association for Computing Machinery (ITiCSE '21), pp. 457–463. Available at: https://doi.org/10.1145/3430665.3456332.

Barus, O.P. *et al.* (2025) 'Shaping generative AI governance in higher education: Insights from student perception', *International Journal of Educational Research Open*, 8, p. 100452. Available at: https://doi.org/10.1016/j.ijedro.2025.100452.

Beyer, S. (2014) 'Why are women underrepresented in Computer Science? Gender differences in stereotypes, self-efficacy, values, and interests and predictors of future CS course-taking and grades', *Computer Science Education*, 24(2–3), pp. 153–192. Available at: https://doi.org/10.1080/08993408.2014.963363.

Chen, S.-Y. *et al.* (2022) 'Exploring the factors of students' intention to participate in AI software development', *Library Hi Tech*, 42(2), pp. 392–408. Available at: https://doi.org/10.1108/LHT-12-2021-0480.

Cheryan, S. *et al.* (2013) 'The Stereotypical Computer Scientist: Gendered Media Representations as a Barrier to Inclusion for Women', *Sex Roles*, 69(1), pp. 58–71. Available at: https://doi.org/10.1007/s11199-013-0296-x.

Chun, J. and Elkins, K. (2023) 'The Crisis of Artificial Intelligence: A New Digital Humanities Curriculum for Human-Centred AI', *International Journal of Humanities and Arts Computing*, 17(2), pp. 147–167. Available at: https://doi.org/10.3366/ijhac.2023.0310.

Clarke, V. and Braun, V. (2017) 'Thematic analysis', *The Journal of Positive Psychology*, 12(3), pp. 297–298. Available at: https://doi.org/10.1080/17439760.2016.1262613.

Cowit, N.Q. and Fiesler, C. (2024) 'Student Preconceptions of Artificial Intelligence: Results from Single Institution Survey', *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 2*. New York, NY, USA: Association for Computing Machinery (SIGCSE 2024), pp. 1610–1611. Available at: https://doi.org/10.1145/3626253.3635484.

Department for Education (2025) *AI in schools and colleges: what you need to know*, *The Education Hub*, 10 June. Available at: https://educationhub.blog.gov.uk/2025/06/artificial-intelligence-in-schools-everything-you-need-to-know/ (Accessed: 4 February 2026).

Dijk, J. van (2020) *The digital divide*. Cambridge, UK: Polity Press.

Ding, Y. *et al.* (2025) 'Voices of Her: Analyzing Gender Differences in the AI Publication World', in K. Atwell et al. (eds) *Proceedings of the Fourth Workshop on NLP for Positive Impact (NLP4PI)*. Vienna, Austria: Association for Computational Linguistics, pp. 196–214. Available at: https://doi.org/10.18653/v1/2025.nlp4pi-1.17.

Drydakis, N. (2024) 'Artificial intelligence capital and employment prospects', *Oxford Economic Papers*, 76(4), pp. 901–919. Available at: https://doi.org/10.1093/oep/gpae005.

Drydakis, N. (2025) 'The formation of AI Capital in higher education: Enhancing students' academic performance and employment rates', *Computers and Education: Artificial Intelligence*, 9, p. 100476. Available at: https://doi.org/10.1016/j.caeai.2025.100476.

DSIT (2025) *AI Opportunities Action Plan.* Department for Science, Innovation and Technology. Available at: https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan (Accessed: 12 January 2026).

Duong, C.D. (2026) 'AI literacy and higher education students' digital entrepreneurial intention: A moderated mediation model of AI self-efficacy and digital entrepreneurial self-efficacy', *Industry and Higher Education*, 40(2), pp. 242–255. Available at: https://doi.org/10.1177/09504222251370089.

Elhussein, G. *et al.* (2024) *Shaping the Future of Learning: The Role of AI in Education 4.0*. Available at: https://www.weforum.org/publications/shaping-the-future-of-learning-the-role-of-ai-in-education-4-0/ (Accessed: 30 March 2026).

Ferguson, S. *et al.* (2023) *Advancing a Model of Students' Intentional Persistence in Machine Learning and Artificial Intelligence*, *arXiv.org*. Available at: https://arxiv.org/abs/2311.10744v1 (Accessed: 30 March 2026).

Fife, S.T. and Gossner, J.D. (2024) 'Deductive Qualitative Analysis: Evaluating, Expanding, and Refining Theory', *International Journal of Qualitative Methods*, 23, p. 16094069241244856. Available at: https://doi.org/10.1177/16094069241244856.

Frachtenberg, E. and Kaner, R.D. (2022) 'Underrepresentation of women in computer systems research', *PLOS ONE*, 17(4), p. e0266439. Available at: https://doi.org/10.1371/journal.pone.0266439.

Gillespie, A., Glăveanu, V. and De Saint Laurent, C. (2024) *Pragmatism and Methodology: Doing Research That Matters with Mixed Methods*. 1st edn. Cambridge University Press. Available at: https://doi.org/10.1017/9781009031066.

Harehdasht, Z.F. and Saki, R. (2025) 'Decoding the Gender Gap: Addressing Gender Stereotypes and Psychological Barriers to Empower Women in Technology'. Available at: https://doi.org/10.5281/zenodo.17226899.

Hopson, S. *et al.* (2025) 'Enhancing AI literacy in undergraduate pre-medical education through student associations: an educational intervention', *BMC medical education*, 25(1), p. 999. Available at: https://doi.org/10.1186/s12909-025-07556-2.

Hornberger, M., Bewersdorff, A., and Nerdel, C. (2023). What do university students know about Artificial Intelligence? Development and validation of an AI literacy test. *Computers and Education: Artificial Intelligence*, *5*, 100165. https://doi.org/10.1016/j.caeai.2023.100165

Hornberger, M. *et al.* (2025) 'Development and validation of a short AI literacy test (AILIT-S) for university students', *Computers in Human Behavior: Artificial Humans*, 5, p. 100176. Available at: https://doi.org/10.1016/j.chbah.2025.100176.

Humble, N. and Mozelius, P. (2022) 'Content analysis or thematic analysis: Similarities, differences and applications in qualitative research'. *European conference on research methodology for business and management studies*, pp. 76–81.

Islam, M.R., Khan, N.A. and Baikady, R. (eds) (2022) *Principles of Social Research Methodology*. Singapore: Springer Nature. Available at: https://doi.org/10.1007/978-981-19-5441-2.

Jiang, Y., Xie, L. and Cao, X. (2025) 'Exploring the Effectiveness of Institutional Policies and Regulations for Generative AI Usage in Higher Education', *Higher Education Quarterly*, 79(4), p.

e70054. Available at: https://doi.org/10.1111/hequ.70054.

Kalim, U. *et al.* (2025) 'Barriers to AI adoption for women in higher education: a systematic review of the Asian context', *Smart Learning Environments*, 12(1), p. 38. Available at: https://doi.org/10.1186/s40561-025-00390-5.

Kelly, S., Kaye, S.-A. and Oviedo-Trespalacios, O. (2023) 'What factors contribute to the acceptance of artificial intelligence? A systematic review', *Telematics and Informatics*, 77, p. 101925. Available at: https://doi.org/10.1016/j.tele.2022.101925.

Kuhlman, C., Jackson, L. and Chunara, R. (2020) 'No computation without representation: Avoiding data and algorithm biases through diversity'. arXiv. Available at: https://doi.org/10.48550/arXiv.2002.11836.

Laupichler, M.C. *et al.* (2023) 'Development of the "Scale for the assessment of non-experts' AI literacy" – An exploratory factor analysis', *Computers in Human Behavior Reports*, 12, p. 100338. Available at: https://doi.org/10.1016/j.chbr.2023.100338.

Lewis, C.M., Anderson, R.E. and Yasuhara, K. (2016) '"I Don't Code All Day": Fitting in Computer Science When the Stereotypes Don't Fit', *Proceedings of the 2016 ACM Conference on International Computing Education Research*. New York, NY, USA: Association for Computing Machinery (ICER '16), pp. 23–32. Available at: https://doi.org/10.1145/2960310.2960332.

Lindner, R. (2020) *Barriers to doctoral education: Equality, Diversity and Inclusion for Postgraduate Research Students at UCL*. UCL Doctoral Shool: UCL. Available at: https://www.ucl.ac.uk/doctoral-school/sites/doctoral_school/files/barriers-to-doctoral-education_0.pdf?utm_source=chatgpt.com (Accessed: 6 January 2026).

Long, D. and Magerko, B. (2020) 'What is AI Literacy? Competencies and Design Considerations', *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery (CHI '20), pp. 1–16. Available at: https://doi.org/10.1145/3313831.3376727.

Low, A.M. and Kalender, Z.Y. (2024) 'Predictors and Socio-Demographic Disparities in STEM Degree Outcomes: A UK Longitudinal Study using Hierarchical Logistic Regression'. arXiv. Available at: https://doi.org/10.48550/arXiv.2408.05853.

Lu, W. and Lin, C. (2025) 'How do artificial intelligence literacy constructs work—based on a survey of university non-expert students', *Education and Information Technologies*, 30(10), pp. 13779–13805. Available at: https://doi.org/10.1007/s10639-025-13323-z.

Maslej, N. (2025) 'Artificial Intelligence Index Report 2025', *Artificial Intelligence* [Preprint]. Available at: https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf (Accessed: 11 February 2026).

Mendigutxia, A. and Pedró, F. (2024) *The role of higher education in national artificial intelligence strategies: a comparative policy review - UNESCO Digital Library*. Available at: https://unesdoc.unesco.org/ark:/48223/pf0000392047_eng (Accessed: 30 March 2026).

MIOIR (2025) *The AI Thesis Network: A glimpse into gender dynamics of UK STEMM PhD theses*, *Manchester Institute of Innovation Research blog*. Available at: https://blogs.manchester.ac.uk/mioir/2025/09/19/the-ai-thesis-network-a-glimpse-into-gender-dynamics-of-uk-stemm-phd-theses/ (Accessed: 30 March 2026).

Molyneaux, D.E. and Hunt, D.L. (2022) 'DIVERSITY AND INCLUSION SURVEY (DAISY) QUESTION GUIDANCE'. Available at: https://edisgroup.org/wp-content/uploads/2022/05/DAISY-guidance-current-upated-May-2022-V2.pdf (Accessed: 3 March 2026).

New Jersey Department of Education (2025) 'New Jersey Department of Education Announces Grant Awards to Support Artificial Intelligence in Education'.

Niousha, R. *et al.* (2024) 'Mapping the Pathways: A Comparative Analysis of AI/ML/DS Prerequisite Structures in R1 Institutions in the United States', *2024 IEEE Frontiers in Education Conference (FIE)*. *2024 IEEE Frontiers in Education Conference (FIE)*, pp. 1–9. Available at: https://doi.org/10.1109/FIE61694.2024.10893290.

Oliveira, J. *et al.* (2024) 'Exploring the Adoption Phenomenon of Artificial Intelligence by Doctoral Students Within Doctoral Education', *New Horizons in Adult Education and Human Resource Development*, 36(4), pp. 248–262. Available at: https://doi.org/10.1177/19394225241287032.

ONS, O. for N.S. (2021) *Census 2021 paper questionnaires*. Available at: https://www.ons.gov.uk/census/censustransformationprogramme/questiondevelopment/census2021paperquestionnaires (Accessed: 9 March 2026).

Oskotsky, T. *et al.* (2022) 'Nurturing diversity and inclusion in AI in Biomedicine through a virtual summer program for high school students', *PLOS Computational Biology*, 18(1), p. e1009719. Available at: https://doi.org/10.1371/journal.pcbi.1009719.

Parker, L. *et al.* (2025) 'Comparative analysis of artificial intelligence policies in universities across five countries', *Discover Computing*, 28(1), p. 267. Available at: https://doi.org/10.1007/s10791-025-09745-5.

Patole, S., Carpenter, R. and McWhorter, R. (2024) 'Artificial Intelligence (AI) and Ethical Use of Data in HRD'.

Roberts, S.F. *et al.* (2021) 'Review of Racially Equitable Admissions Practices in STEM Doctoral Programs', *Education Sciences*, 11(6). Available at: https://doi.org/10.3390/educsci11060270.

Sarraju, A., Ngo, S. and Rodriguez, F. (2023) 'The leaky pipeline of diverse race and ethnicity representation in academic science and technology training in the United States, 2003–2019', *PLOS ONE*, 18(4), p. e0284945. Available at: https://doi.org/10.1371/journal.pone.0284945.

Shams, R.A., Zowghi, D. and Bano, M. (2025) 'AI and the quest for diversity and inclusion: a systematic literature review', *AI and Ethics*, 5(1), pp. 411–438. Available at: https://doi.org/10.1007/s43681-023-00362-w.

Smith, P. *et al.* (2023) *PGR Finance*. University of Bristol. Available at: https://drive.google.com/drive/u/0/home (Accessed: 6 January 2026).

Stanford University (2023) *Artificial Intelligence Index Report: Introduction to the AI Index Report 2023 GP-003*. Available at: https://hai.stanford.edu/ai-index/2023-ai-index-report (Accessed: 30 March 2026).

Stathoulopoulos, K. and Mateos-Garcia, J.C. (2019) 'Gender Diversity in AI Research', *SSRN Electronic Journal* [Preprint]. Available at: https://doi.org/10.2139/ssrn.3428240.

The Digital Cooperation Organization (2025) *AI-REAL TOOLKIT AI READINESS TO EMPOWERMENT, ADOPTION, AND LEADERSHIP*. Available at: https://dco.org/dco-ai-adoption-playbook/ (Accessed: 30 March 2026).

UKCGE (2024) *Equity, Diversity, and Inclusion in Postgraduate Study 2022/23*, *UK Council for Graduate Education*. Available at: https://ukcge.ac.uk/resources/resource-library/equity-diversity-and-inclusion-in-postgraduate-study-2022-23 (Accessed: 30 March 2026).

Ulnicane, I. (2025) 'Governance fix? Power and politics in controversies about governing generative AI', *Policy and Society*, 44(1), pp. 70–84. Available at: https://doi.org/10.1093/polsoc/puae022.

UNESCO (2025) *Prep-AI: Preparing policy-makers and educators for AI and inclusive*. Available at: https://www.unesco.org/en/digital-education/artificial-intelligence/prep-ai (Accessed: 31 March 2026).

Universities UK International (2024) *Annual report 2023–24*. London. Available at: https://www.universitiesuk.ac.uk/sites/default/files/field/downloads/2024-10/Universities UK International Annual Report 23-24.pdf (Accessed: 17 December 2025).

University of Bristol (2024) *September: AI Awards | News and features | University of Bristol*. Available at: https://www.bristol.ac.uk/news/2024/september/ai-awards-2024.html (Accessed: 30 March 2026).

University of Bristol (no date) *Equality, Diversity and Inclusion | Centres for Doctoral Training | University of Bristol*, *UKRI Centre for Doctoral Training in Interactive Artificial Intelligence*. Available at: https://www.bristol.ac.uk/cdt/interactive-ai/about-the-centre/equality-diversity-and-inclusion/?utm_source=chatgpt.com (Accessed: 30 March 2026).

Vargas-Solar, G. (2025) ‘Bridging the Divide: Gender, Diversity, and Inclusion Gaps in Data Science and Artificial Intelligence Across Academia and Industry in the majority and minority worlds’. arXiv. Available at: https://doi.org/10.48550/arXiv.2511.18558.

van de Werfhorst, H.G. (2017) ‘Gender Segregation across Fields of Study in Post-Secondary Education: Trends and Social Differentials’, *European Sociological Review*, 33(3), pp. 449–464.

Williams, S. *et al.* (2025) ‘Women upskilling or reskilling to an ICT career: A systematic review of drivers and barriers’. arXiv. Available at: https://doi.org/10.48550/arXiv.2510.22508.

Wu, C., Zhang, H. and Carroll, J.M. (2024) ‘AI Governance in Higher Education: Case Studies of Guidance at Big Ten Universities’, *Future Internet*, 16(10), p. 354. Available at: https://doi.org/10.3390/fi16100354.

Xia, J., Li, M. and Li, J. (2024) ‘Comparative Analysis Vision of Worldwide AI Courses’. arXiv. Available at: https://doi.org/10.48550/arXiv.2407.16881.

Yin, S.X. and Goh, D. (2024) ‘Factors Influencing University Students’ AI Use and Knowledge Acquisition’, *Proceedings of the Association for Information Science and Technology*, pp. 1162–1164. Available at: https://doi.org/10.1002/pra2.1218.

Young, E., Wajcman, J. and Sprejer, L. (2023) ‘Mind the gender gap: Inequalities in the emergent professions of artificial intelligence (AI) and data science’, *New Technology, Work and Employment*, 38(3), pp. 391–414. Available at: https://doi.org/10.1111/ntwe.12278.

Zhang, D. *et al.* (2025) ‘Modeling the relationships between secondary school students’ AI learning attitude, AI literacy and AI career interest’, *Education and Information Technologies*, 30(17), pp. 25223–25250. Available at: https://doi.org/10.1007/s10639-025-13715-1.

Zhang, S., Prasad, G. P. and Schroeder, N.L. (2024) ‘Learning About AI: A Systematic Review of Reviews on AI Literacy’, *Journal of Educational Computing Research*, 63(5), pp. 1292–1322. Available at: https://doi.org/10.1177/07356331251342081.

Zhang, Y. *et al.* (2025) ‘Student Experiences of Hybrid and Online Engineering Labs in a Logic Control Course’, *Computer Applications in Engineering Education*, 33(3), p. e70032. Available at: https://doi.org/10.1002/cae.70032.

# Appendix A. LLM Prompt and Information Extraction Schema

Part 1: System Prompt and Instructions

The following prompt was wrapped around each individual PhD listing prior to processing:

```
You are an expert at analyzing PhD listings and extracting structured information about barriers to entry, funding, and diversity initiatives.

Please analyse the following PhD listing and extract all relevant information according to the provided schema. If a field is not explicitly mentioned in the text, set it to null. Do not infer or guess values.

Title: [Listing Title]

Full Description:

[Full Listing Text]

Extract all available information according to the schema. For fields where information is not available, use appropriate null/empty values. Be thorough and precise in your extraction.
```

Part 2: Extraction Schema Definition

The LLM was instructed to populate the following data points based on the text of the listing:

**Basic Info**

- **Id**: The Listing ID from the input file.
- **Title**: Title of the listing.

**University & Location**

- **Awarding University**: The institution granting the degree. If missing in text, use null.

**Supervisors** *Contains the following information for each supervisory team member:*

- **Name**: Text
- **Email**: Text
- **Profile Url**: Text
- **Gender Inference**: Value (Male, Female, Ambiguous/Unknown) and Confidence Score.

**Funding**

- **Status:** Funded, Self-Funded, Competition, Studentship, University studentship/University competition, Unspecified.
- **Funding Certainty:** Secured / Reserved Funding, Competitive / Conditional Funding, Unclear / Not Specified. (Includes exact evidence text).
- **Source Category:** UKRI/Research Council, EPSRC, EU/International Research Council, University, Industry, Charity, Mixed, Unspecified.
- **Financial Details:** Annual stipend in GBP.
- **Financial Top Up:** True if any additional funding beyond stipend/fees is mentioned (e.g., training funds, conference budget, consumables).
- **Fees Coverage:** Home Only, International, None, Unspecified.
- **Duration Years:** Number (decimal).

**Entry Requirements**

- **Degree Level Required:** First, 2:1, 2:2, Masters, Unspecified, Other. (Instruction: Do not infer 2:1 unless explicitly stated. Use Unspecified if no grade is listed).
- **Discipline Tunnel Vision:** True if the listing accepts ONLY Computer Science, Math, or Physics. False if it accepts Biology, Psychology, Humanities, etc.
- **Equivalent Experience:** True if accepted.
- **Ielts Requirement:** Extract specific English language requirement details if present (e.g., 'IELTS 6.5 overall').
- **References Count:** Number of academic references required.
- **Own Transport Required:** Does it mention needing a car or driving licence?
- **Requires Research Proposal:** Yes/No
- **Requires Cv:** Yes/No
- **Requires Cover Letter:** Yes/No
- **Contact Supervisor First:** Does it say 'Contact supervisor before applying'?

**Immigration Status** *Legal ability to study (Visa/ATAS), distinct from funding.*

- **Can Sponsor Visa:** Can the university sponsor a Student Visa? Default TRUE for UK universities unless 'UK Nationals Only' is stated.
- **Requires Atas:** Is ATAS security clearance mentioned?
- **Raw Evidence Text:** Exact text justifying the classification.

**Interview Details**

- **Mentions Technical Test**: Yes/No

**Study Mode**

- **Part Time Availability**: True if the text includes part-time, flexible arrangements, including part-time PhD, part-time study, flexible study mode, or similar.
- **Hybrid Availability**: True if hybrid/remote working is explicitly mentioned.

**Diversity Initiatives**

- **Tier**: Tier A (Active) or Tier B (Passive) or None. (Instruction: Tier A = Guaranteed interviews, citing Equality Act, specific bursaries. Tier B = Generic 'we encourage applications' statement).
- **Raw Text**: Raw text of the diversity initiative if present.

**Culture And Language**

- **Elitist Language Score**: Simple keyword-based score. Count the number of distinct explicit prestige terms (e.g., 'world-class', 'elite', 'prestigious', 'top-tier', 'leading university'). Score equals the count, capped at 5. Do not infer prestige beyond exact terms used.

**Topic Classification**

- **Category**: Core AI, Applied AI, Policy & Ethics, Tangential / Non-Core, Not AI, Unsure.
- **Confidence Score**: Number (decimal).
- **Keywords**: List of topics justifying the classification.

**Technical Skills**

- **Explicit Languages**: List any programming languages explicitly mentioned (e.g., Python, R, C++, Java, MATLAB).
- **Explicit Libraries Or Frameworks**: List any named libraries, frameworks, or platforms (e.g., PyTorch, TensorFlow, scikit-learn, NumPy, HuggingFace, ROS).

**Data Quality Flags** *Flags to indicate missing or poor quality source text.*

- **Is Entry Requirements Missing:** Yes/No
- **Is Supervisor Missing:** Yes/No
- **Is University Missing:** Yes/No
- **Is Critical Hallucination Risk:** True if the text is very short/vague but the schema is fully populated (likely hallucination).

# Appendix B. Pre-Training Survey

The pre- and post-training surveys contained virtually identical questions, hence why only the pre-training survey is provided here.

## Pre- AI Training Survey

In this section, you will be asked about your background.

19. Please write your course and year below:

For example, "Year 2 Politics and International Relations", "Year 3 Modern Languages (French and Portuguese)" or "Year 4 Mechanical Engineering with a Year in Industry"

Enter your answer

20. What is your age?

Enter your answer

21. What is your ethnic group? Please select all the options that best describe your ethnicity or background (e.g. you could select "Any other Black Background" and "English / Welsh / Scottish / Northern Irish / British" if this best reflects your identity.)

- [ ] Bangladeshi
- [ ] Chinese
- [ ] Indian
- [ ] Pakistani
- [ ] Another Asian background (please describe in the next question)
- [ ] African
- [ ] Caribbean
- [ ] Any other Black / African / Caribbean background (please describe in the next question)
- [ ] English / Welsh / Scottish / Northern Irish / British
- [ ] Gypsy or Irish Traveller
- [ ] Irish
- [ ] Roma
- [ ] Any other White background (please describe in the next question)
- [ ] Arab
- [ ] Hispanic
- [ ] Latina / Latino / Latinx
- [ ] Any other ethnic group (please describe in the next question)

22. If you selected any of the "Any other Asian/Black/African/Caribbean/White background" or "Any other ethnic group" options in the previous question, please write your ethnic background here:

Enter your answer

23. Which of the following best describes your gender?

- ○ Man
- ○ Non-binary
- ○ Woman
- ○ Prefer not to say
- ○ Prefer to self-describe

## AI Knowledge

The following questions have been adapted from Hornberger et al. (2025). They ask about your current knowledge of AI as we want to explore how this changes with you completing both training sessions. Because of this, it is essential that you do not use any additional materials to answer these questions. Also, you will not be shown the correct answers.

24. In which of these areas is AI typically applied?

- Detecting credit card fraud
- Cryptocurrency mining
- Web tracking
- Encryption for instant messaging services

25. Which of the following interdisciplinary research fields is also a subfield of AI?

- Blockchain
- Natural Language Processing
- Psychology of Learning
- Bioinformatics

26. For which task was AI first shown to be superior to human experts?

- Detecting tumours
- Conducting software projects
- Translating novels
- Designing cancer therapies

27. How does supervised learning differ from unsupervised learning?

- In supervised learning, the output values of the training data are known.
- In supervised learning, humans must supervise the AI during learning and intervene if necessary.
- In supervised learning, all computational steps are documented.
- In supervised learning, stricter legal regulations apply.

28. Sort the process steps in supervised learning into the correct order by dragging and dropping:

- Train model with training data.
- Predict test data with the model.
- Collect and prepare data.
- Divide data into training and test data.
- Calculate accuracy of prediction.

29. How do AI developers most typically shape the results of the machine learning process?

- ◯ Through calculation of the accuracy of the prediction.
- ◯ Through randomised division into test and training data.
- ◯ Through selection of the model.
- ◯ Through abstraction of the model.

30. What primarily determines the behaviour of AI systems?

- ◯ AI systems strive for autonomy.
- ◯ AI systems pursue a goal that has been given to them by humans.
- ◯ AI systems perform behaviours randomly.
- ◯ AI systems seek out goals independently and pursue them.

31. You are testing a machine learning model that is supposed to classify images of animals. You notice that the model is better at recognising cats than dogs. What could be the reason for this?

- ◯ Dogs are more difficult to recognise than cats since there are fewer images of dogs on the internet.
- ◯ Small objects (cats) are better recognised than large ones (dogs).
- ◯ Most models are generally better at recognising cats than dogs.
- ◯ The training data of the dogs were not representative of all dog breeds.

32. What is the black box problem?

- ◯ AI entails a residual risk that is hard to calculate.
- ◯ It is often difficult to determine how an AI system makes decisions.
- ◯ Users are often not informed that an AI system is being used.
- ◯ Many users have little knowledge about AI.

33. What is a central risk in using AI for predictive policing?

- ◯ Vulnerability to hacking
- ◯ Discrimination against suspects based on origin and status
- ◯ Lack of legal certainty in the event of an AI failure
- ◯ Undermining the authority of police officers

## This section will ask about your interest in and awareness of AI

34. Please indicate your agreement with the following statements.

| | Strongly disagree | Somewhat disagree | Neither agree nor disagree | Somewhat agree | Strongly agree |
|---|---|---|---|---|---|
| I am confident that I can hold an informed conversation about AI. | ◯ | ◯ | ◯ | ◯ | ◯ |
| I am concerned about the ethics of AI. | ◯ | ◯ | ◯ | ◯ | ◯ |
| I think I would be offered a place to study a PhD in AI if I applied. | ◯ | ◯ | ◯ | ◯ | ◯ |
| I am satisfied with the number of people within the field of AI who are like me. | ◯ | ◯ | ◯ | ◯ | ◯ |

35. At this stage, would you consider studying AI in a specialised master's course, short course or other professional course?

- ◯ Yes
- ◯ No
- ◯ Maybe
- ◯ Prefer not to say

36. At this stage, would you consider applying for a PhD in AI?

By this, we mean whether you would consider applying for study of a postgraduate research or postgraduate taught degree which focuses on the applications, development and understanding of AI.

- ◯ Yes
- ◯ No
- ◯ Maybe
- ◯ Prefer not to say

37. Please list up to three things that interest you in studying AI.

Enter your answer

38. Please list up to three things that interest you in undertaking a PhD.

Enter your answer

## Your motivation to participate in the training.

39. Please write up to 100 words about what has motivated you to take part in the AI training offered by the Jean Golding Institute.

Enter your answer

# Appendix C. Interview Schedule

| No. | Question | Probe / Follow-up | Research Questions |
|---|---|---|---|
| 1 | What were the main factors in you deciding to study your current course / choose your current job? | Has the AI training affected which careers now interest you for the future? | RQ2 (primary) RQ1 (secondary) |
| 2 | What's influencing your decision to pursue/not pursue studies in AI? | What are your thoughts about doing a postgraduate degree in AI? Have these changed as a result of the training? | RQ1 (primary) RQ2 (secondary) |
| 3 | What type of support do you think you would want or need for studying AI? | Do you think you are likely to receive that support easily? Why or why not? | RQ1 (primary) RQ2 (secondary) |
| 4 | From your point of view, do you think that studying AI (courses undergraduate, postgraduate, trainings) seems accessible? | Why or why not? | RQ1 (primary) RQ2 (secondary) |
| 5 | How clear are the university's AI policies to you? | What effect would you say the university's AI policies have had on your interest and use of it? What would you change about the policy, if anything? | RQ3 (primary) RQ1 (secondary) |
| 6 | How has the AI training programme changed the way you think about studying AI? | Is there anything you hoped to get out of it but didn't? | RQ1 (primary) RQ2 (secondary) |
| 7 | Is there anything else you would like to add? | | All RQs. |

# Appendix D. Codebook for Qualitative Interviews

The sub-themes, sub-sub-themes and sub-sub-sub-themes are indicated with a corresponding number of arrows (e.g: subthemes of the above theme are marked “> Title”).

| Name | Description | Files | References |
|---|---|---|---|
| Personal factors | Includes personal elements/circumstances that affect academic paths choices positively or negatively. | 0 | 0 |
| > Attitudes towards using-studying AI | This code includes participants' willingness to engage with AI whether to use it in their work or to take on a postgraduate degree in AI. It is connected to the other codes regarding technical skills and future careers. | 0 | 0 |
| >> Negative | | 1 | 3 |
| >> Not interested in a future career | | 4 | 5 |
| >> Positive | | 2 | 4 |
| >> Positive but sceptical | | 2 | 2 |
| >> Positive due its necessity for the future | | 4 | 6 |
| > Factors influencing self efficacy | Includes the factors behind student's perception of their ability to pursue a postgraduate study in an AI related field. | 3 | 4 |
| > Motivation influencing academic choices | Factors that encourage/discourage students to take a certain academic path. | 0 | 0 |
| >> Create positive impact | | 1 | 1 |
| >> Environmental influence (family, tutors) | | 1 | 3 |
| >> Passion and | | 3 | 7 |

| | | | |
|---|---|---|---|
| interests | | | |
| >>> Practical decision | | 2 | 4 |
| Skills and knowledge | Includes codes of knowledge and skills in AI and their positive and/or negative impact on students' attitudes to study AI or use it. | 0 | 0 |
| > Capacity building | This code describes the role of capacity building in influencing students' decisions to engage with AI postgraduate studies, their daily work and education tasks. It falls under the theme of AI literacy and relates to the technical skills aspect. It includes the children codes related to the "getting started in AI" training, self learning efforts, and increased awareness. | 5 | 11 |
| >> Increased awareness | | 3 | 10 |
| >> Self teaching | This code includes the personal effore made by interviewees to learn AI-related skills. | 3 | 3 |
| >> Training evaluation | This code includes general reflections on the quality of the training and specific reflections on its influence on attitudes towards using AI and\or sudying AI in the future. | 2 | 3 |
| >>> Accessibility | This codes includes evaluation of accessibility to AI capacity building opportunities. | 4 | 17 |
| >>> Content | This code includes the evaluation of the content of the training and the role it plays in influencing attitudes towards using and/or studying AI. | 5 | 14 |
| >>> Delivery | This codes includes evaluations regarding the way in which the training was delivered and its impact on attitudes towards using and/or studying AI. | 4 | 15 |
| > Technical Skills | Includes barriers and enablers related to technical skills perceived to be needed in AI related fields. | 0 | 0 |
| >> Coding | Positive and negative influences of the exposure/lack of exposure to coding-related skills and knowledge on students attitudes towards using it in their work or taking on a postgraduate degree in AI | 3 | 5 |
| >> Previous | Positive and negative influences of previous | 4 | 9 |

| | | | |
|---|---|---|---|
| knowledge-experience-degrees | exposure/ lack of exposure to AI related skills and knowledge on students attitudes towards using it in their work or taking on a postgraduate degree in AI | | |
| >> Understanding of concepts | Positive and negative influences of understanding general concepts and mechanisms of AI systems on students attitudes towards using it in their work or taking on a postgraduate degree in AI | 3 | 5 |