Educating the Public in Artificial Intelligence: Insights from a Large-Scale Course

Maria Kasinidou[1], **Styliani Kleanthous** [1], **Jahna Otterbacher**[1, 2], **Evgenia Christoforou** [1]

[1]Open University of Cyprus
[2]CYENS Center of Excellence, Nicosia, Cyprus
{maria.kasinidou, styliani.kleanthous, jahna.otterbacher, evgenia.christoforou}@ouc.ac.cy

Abstract

As AI technologies grow more influential in shaping modern life, there is an urgent need to make AI literacy accessible beyond academic and technical communities. This paper presents the design, delivery, and evaluation of an online AI course targeting the general public. The course combined asynchronous lectures, interactive live sessions, and reflective assignments. Of the 343 people who registered, 169 completed the program. Using validated instruments administered before and after the course, we measured changes in participants' attitudes toward AI and their AI literacy. Our findings revealed statistically significant changes in AI literacy, specifically in awareness, usage, and evaluation constructs, as well as a rise in positive attitudes toward AI. High satisfaction scores and qualitative feedback further support the course's effectiveness. These findings reinforce the importance of inclusive, scalable educational interventions for empowering the public to navigate AI technologies.

Introduction and Related Work

Artificial Intelligence (AI) technologies now permeate everyday life, from personalized recommendations in news feeds and streaming platforms, to chatbots like ChatGPT that support writing and information-seeking. Yet many people remain unaware of how often they interact with AI and continue to hold misconceptions about its nature and capabilities (Kelley et al. 2021). AI's capabilities continue to expand rapidly and be integrated in many aspects of our lives, leading to a pressing call for AI literacy across the population to equip individuals with the knowledge needed to recognize, use and evaluate AI considering ethical premises (Long and Magerko 2020; Ng et al. 2021). Research community highlight the urgency of upskilling and reskilling the public's AI literacy, noting that people often face challenges in acquiring AI-related competencies (Pham et al. 2024).

Research has consistently shown that educational interventions enhance understanding of AI (Lee et al. 2021a; Laupichler et al. 2022) and can shape acceptance and utilization of AI technologies (Schiavo, Businaro, and Zancanaro 2024; Hollands et al. 2024). Hence, organizations such as the Association for the Advancement of Artificial Intelligence (AAAI), OECD, and the European Commission, are working with educators and researchers on designing effective AI educational interventions (Touretzky et al. 2019; Almatrafi, Johri, and Lee 2024) and developing assessments on their impact on participants' perception of AI and AI literacy (Masla et al. 2025; Ng et al. 2021). These initiatives primarily target university students like computer science (CS) students (Wollowski et al. 2016; Kasinidou et al. 2021; Stadelmann et al. 2021) as well as students from non-CS disciplines (Wunderlich, Higgins, and Lichtenstein 2021; Kong, Cheung, and Zhang 2022, 2023; Kasinidou et al. 2024; Southworth et al. 2023; Biswas et al. 2025).

Courses aimed at improving AI literacy among students with no CS background show promising results, with participants demonstrating increased understanding of AI and greater awareness of related career paths (Lee et al. 2021b). For example, non-CS learners enrolled in AI courses demonstrated substantial improvements in their understanding of AI (Kong, Cheung, and Zhang 2022, 2021) and a greater awareness of AI's ethical and socio-technical challenges (Kong, Cheung, and Zhang 2023; Kasinidou et al. 2024). These courses also boosted their confidence in using and working with AI (Kong, Cheung, and Zhang 2021; Kupferschmidt et al. 2025). Early AI education has also demonstrated success in fostering nuanced and realistic perspectives of AI (Druga and Ko 2021; Arn and Huang 2024) while also enhancing learners' understanding of AI's ethical implications (Kong, Cheung, and Tsang 2023). Similarly, educational programs for professionals have been shown to help participants understand how AI can be effectively incorporated into their workplace operations (Johnson et al. 2022).

Beyond higher education, increasing efforts are being made to design AI education for K-12 students (Sabuncuoglu 2020; Long and Magerko 2020; Lee et al. 2021b). Although recent initiatives have made significant strides in reaching K-12 (Yue, Jong, and Dai 2022; Kasinidou, Kleanthous, and Otterbacher 2026b,a), and university students (Laupichler et al. 2022; Ng et al. 2022), efforts to engage the broader public remain limited. Yet, emerging studies show that such outreach can lead to meaningful gains in AI literacy among non-experts (Kasinidou, Kleanthous, and Otterbacher 2023; Stoyanovich et al. 2025). As AI becomes increasingly embedded in everyday life, educational efforts must extend to reach individuals from diverse backgrounds (Maher and Tadimalla 2024). This shift calls for inclusive,

accessible learning models that reflect the real-world needs and perspectives of the public (Long, Teachey, and Magerko 2022). Moreover, embedding ethical and socio-technical dimensions into such programs is key to fostering responsible and informed engagement with AI (Tadimalla and Maher 2025). While various approaches have explored how to introduce AI to K-12 and university students, questions remain about how to best design educational experiences that are both accessible and impactful for the general public. In this paper, we share insights from the design, implementation, and evaluation of a large-scale open course aimed at raising AI literacy among adults with no technical background. We focus in particular on how the course shaped participants' attitudes toward AI and enhanced their AI literacy.

'AI in Everyday Life' Course

AI in Everyday Life was developed as part of the PINNACLE project (Kasinidou et al. 2025) in the spirit of lifelong learning (Knapper and Cropley 2000), featuring both synchronous and asynchronous lectures to facilitate learning at participants' own pace. The course was open to up to 400 participants with no registration fees, as a service to the community, and it was offered in both English and Greek and took place from April to June of 2025. Those who completed all required assessments (i.e., quizzes and assignments) received a certificate of completion (CoC). The objectives of the course were to develop an understanding of the components of contemporary, data-driven AI and its development processes; to gain familiarity with foundational concepts of AI; to articulate the principles of Machine Learning (ML) and its relationship to Big Data; to critically analyze AI applications by identifying their core functionalities; to evaluate them concerning their potential social and ethical implications; and to acquire a comprehensive understanding of the role of personal data within the broader AI ecosystem.

Structure. The course comprised eight units and followed a blended learning approach, combining pre-recorded video lectures with weekly synchronous interactive sessions. Each unit included one or two short video lectures (10–15 minutes each), along with a self-assessment quiz and an assignment. Learners were expected to watch the videos and complete their assignments before the live sessions. The assignment for Unit 1 was not mandatory; it served as a practice activity to help learners become familiar with the assignment format.

The synchronous sessions were held once a week, lasted approximately one hour, and were recorded and archived on the course's Moodle page. Each session began with a brief 10-minute recap of the week's content, followed by a presentation of selected data from participants' assignments, offering preliminary insights into emerging trends and findings. Following each lecture, a self-administered online quiz was made available, assessing participants' understanding of key concepts and providing instant feedback. While the quizzes were mandatory only for those seeking a CoC, all were encouraged to complete them for their own self-assessment.

Learning Units

Introduction. The objectives of the first unit are threefold and were discussed during the live session of the course. First, it aimed to familiarize learners with the fundamental principles of online and hybrid learning. As the course was delivered in a blended format, it was essential for learners to understand how the asynchronous and synchronous components would interact. Second, the unit introduced the overall aims of the course and explained how pre-recorded video lectures, quizzes, assignments, and live sessions would contribute to those outcomes. Third, learners were introduced to the Moodle platform and the assignments that would support their learning. This unit included only one pre-recorded video, which provided an overview of how AI is increasingly present in everyday life. Key topics covered were the democratization of AI (Gilman 2023) and the concept of Participatory AI (Birhane et al. 2022), as well as how participatory AI would be applied during the course through the assignments and the evaluation of AI applications.

Fundamentals of AI Applications. The second unit focused on deepening the understanding of what makes AI applications "intelligent" and how the democratization of AI has influenced its development. The unit had four learning objectives: to explore the impact of democratization on AI development processes, to identify the common characteristics of intelligent AI applications, to describe and analyze the basic components of such systems, and to critically examine everyday AI applications in terms of their functionality, data sources, and their potential benefits and drawbacks for end-users. The unit included two video lectures. The first video introduced learners to a working definition of AI and broke down its core capabilities. It further examined what it means for a system to be rational and intelligent. In the second video, learners were introduced to how AI applications are built to reason and learn from data, and how they are applied to both well-defined and open-ended problems.

Trustworthy AI. The third unit introduced learners to the ethical dimensions of AI, with a focus on the concept of Trustworthy AI. The learning objectives of this unit were to describe the characteristics of trustworthy AI applications; explain the purpose and scope of the EU AI Act, particularly the distinction between high- and low-risk applications; identify key ethical and social issues that may arise from the use of AI; and understand why addressing social bias in AI applications presents persistent challenges. The unit also included two pre-recorded video lectures. The first addressed a range of ethical concerns surrounding AI, such as bias and discrimination. Learners explored multiple strategies for mitigating ethical risks, including technical safeguards, public awareness, and policy interventions. The second focused on the concept of Trustworthy AI, which outlines seven key requirements for trustworthy AI [1]. The video also introduced practical tools and reflection questions that can be used to evaluate AI applications for trustworthiness and alignment with ethical and human rights principles.

Machine Learning and Big Data. The fourth unit intro-

[1] https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai

duced learners to ML and its relationship with Big Data, as well as emerging developments in Generative AI (GenAI) and Foundation Models. The learning objectives of the unit were to enable learners to describe ML in their own words, understand how big data supports ML processes, explain the role of ML in contemporary AI applications, and develop a basic understanding of Generative Models and Foundation Models. The first video provided a historical overview of ML and explored the defining features of Big Data, focusing on how these two areas have evolved in tandem to drive modern AI. It introduced learners to key ML algorithms, their purposes, strengths, and limitations, and highlighted how ML techniques are applied in various domains—from business operations to everyday consumer technologies. The second video delved deeper into ML architectures and techniques such as supervised, unsupervised, and reinforcement learning. In this context, the concept of Foundation Models was revisited, framing Transformers as versatile systems trained on massive datasets for diverse tasks.

Vision Technologies. The fifth unit focused on Vision Technologies, exploring how AI applications interpret and analyze visual information. The learning objectives of the unit were to introduce learners to computer vision, familiarize them with its common tasks, and illustrate how these technologies are integrated into everyday applications. In addition, the unit aimed to help learners understand how advanced models, such as Convolutional Neural Networks and Generative Adversarial Networks, have driven progress in the field, along with newer architectures like Transformers. The first video introduced the concept of Computer Vision, explaining how machines are trained to process and interpret images and video. The lecture also addressed the growing concerns around Facial Recognition Technologies, emphasizing the importance of considering ethical and societal impacts when deploying such systems. The second video shifted focus to the role of datasets in training Computer Vision models. Participants learned about the importance of labeled datasets, how they support pattern recognition, and how datasets can influence model performance. Prominent public datasets were discussed, alongside the challenges of data bias, the risks of human labeling errors, and how stereotypes may be unintentionally perpetuated through AI applications trained on unrepresentative or flawed data.

Language Technologies. The sixth unit focused on the field of Language Technologies, explaining how they process, interpret, and generate human language. The learning objectives of this unit were to understand the concept of Natural Language Processing (NLP) and recognize its presence in everyday applications. Additionally, it aimed to raise awareness about ethical concerns in Language Technologies. It also aimed to introduce how Large Language Models (LLMs) are advancing the capabilities of modern NLP systems. A key focus was placed on understanding how transformers, like encoder-based models and generative decoder models, have reshaped the landscape of language-based AI. The first video provided an overview of NLP, highlighting its goals and key techniques for enabling machines to analyze, understand, and generate human language. It explored a variety of real-world applications, including machine translation, chatbots, and automated content moderation. The benefits and drawbacks of NLP were also discussed. The second video expanded on this by exploring the emergence and impact of LLMs. It examined how models like BERT and GPT are built using Transformer architectures. The video also addressed the limitations and challenges of NLP and LLMs.

Search Technologies. The seventh unit focused on AI-enabled Search Technologies and the field of Information Retrieval (IR). The learning objectives of the unit were to introduce learners to the primary goals of IR, familiarize them with key search-related tasks and their everyday applications, and encourage critical reflection on the advantages and limitations of modern search systems. Learners were also introduced to how GenAI is shaping the future of search by enabling more contextual, conversational, and user-adaptive experiences. This unit included a single video lecture that provided an overview of the fundamental concepts and processes involved in search technologies. Learners were introduced to the traditional mechanics of search engines, such as crawling, indexing, and ranking, as well as how ML and NLP have improved search relevance and user satisfaction. The lecture traced the evolution from keyword-based Retrieval Systems to more advanced, AI-powered models capable of understanding user intent. The video concluded with a critical examination of the ethical implications of AI in search, including concerns about algorithmic bias, misinformation, and filter bubbles in search systems.

Personalization Technologies. The eighth unit explores the role of personalization in AI applications, focusing on how digital systems tailor content and services to individual users. The learning objectives of the unit were to help learners understand the purpose of AI-driven personalization, recognize its presence in everyday applications, and critically examine both the benefits and potential dangers of its use. The first video introduced learners to personalization in AI applications, explaining how user data, collected explicitly and implicitly, is used to adapt interfaces and recommendations. Examples from platforms such as Netflix, Spotify, and YouTube were used to illustrate different types of personalization systems, including adaptable, adaptive, and hybrid approaches. Learners also explored the underlying AI techniques that power these systems, such as ML and Deep Learning. The second video delved deeper into the implications of AI-powered personalization. It examined how advanced algorithms process vast amounts of behavioral data to generate customized search results. The video also addressed critical ethical challenges, including filter bubbles, algorithmic bias, and privacy concerns.

Assignments

The assignments were designed to help participants reflect on the technologies discussed each week by connecting them to their own experiences with relevant applications in their daily lives. There were two types of tasks:

- Bottom-up: Participants independently identified AI-enabled applications they encountered in their everyday lives, classified the type of underlying technology, and recorded basic observations.

- Top-down: Participants were guided to analyze a specific type of AI-enabled applications, with step-by-step instructions tailored to the weekly topic.

We began the course using a bottom-up approach for the first four assignments. Participants explored and reported on AI-enabled applications they used. Starting from Assignment 5, we shifted to a top-down approach, where they were directed to find and analyze applications that aligned with the specific technology discussed each week. For example, in Assignment 5, they were asked to focus on applications that use Vision Technology. Each week, they completed a log in which they identified the applications or systems they used that were related to the unit's topic. They were asked to report which systems or applications they believed utilized the technology discussed in that week's unit. If they were unable to identify relevant applications, support was provided by suggesting example activities and corresponding applications. After identifying the applications they used, they were asked to select one application to evaluate in more detail. They described their experience using the application and assessed its trustworthiness. They were also asked to evaluate their trust in the application they selected. In the first three units, before the trustworthiness principles were introduced, they simply indicated whether they trusted the app and believed it was worthy of users' trust. Starting from the fifth assignment, once the principles had been introduced, they were asked to evaluate the chosen application using all seven trustworthiness principles and to provide an overall rating of whether they considered it trustworthy.

Methodology

For this study, we received ethical approval from the Cyprus National Bioethics Committee. Participants were asked to respond to the following scales before and after the course:

- General Attitude towards AI scale (GAAIS), which includes 20 items (adopted from (Schepman and Rodway 2020)). GAAIS has two subscales: 12 items assess positive attitudes and 8 items assess negative attitudes.
- AI Literacy scale (AILS) with 12 items (adopted from (Wang, Rau, and Yuan 2023)). AILS assesses proficiency in recognising, using, and evaluating AI under the premise of ethical standards and consists of four constructs: awareness, use, evaluation, and ethics.

Both scales (Likert-scale, 1–5) were validated in English and Greek (Kasinidou, Kleanthoys, and Otterbacher 2025). A course evaluation form was sent to all registered participants, including those who did not complete the course.

Participants. The course announcement was distributed on the social media and a total of 489 individuals expressed interest, with 343 registering. Of those, 169 completed the required assignments, quizzes, and the pre- and post-assessments and received their CoC. Table 1 presents the number of learners who attended the live sessions, watched the video lectures and recordings asynchronously, and those who submitted the assignments and quizzes for each unit. Out of the 343 registered participants, 305 completed the pre-course assessment, of which 169 completed the post-course assessment. Additionally, 103 participants completed the evaluation form. Of these, 90 had completed the course, while 13 respondents had dropped out before completion but still provided feedback. The demographics of the respondents are presented in Table 2.

Course Evaluation and Performance

Attitudes Towards AI and AI Literacy

Wilcoxon signed-rank tests were employed to identify differences in participants' attitudes toward AI and their AI literacy between pre- and post-course responses.

Positive Attitude. We examined changes in participants' positive attitudes toward AI before and after the course. The analysis revealed a statistically significant difference in overall positive attitudes (z=–2.155, p=.031), with the mean score rising from 3.793 to 3.871. Further analysis of individual statements showed statistically significant increases in several items. Participants reported a greater interest in using AI in their daily lives, as reflected in the item "I am interested in using AI applications in my daily life" (P1) (z=–2.918, p=.001). A significant increase was also found in the statement "There are many beneficial applications of AI" (P2) (z=–2.157, p=.016). Moreover, the statement "AI is exciting" (P3) showed a statistically significant increase in post-course responses (z=–3.060, p<.001), indicating heightened enthusiasm for the topic. Similarly, they expressed greater admiration for AI capabilities in the item "I am impressed by what AI can do" (P7) (z=–2.130, p=.017). While no other items showed statistically significant differences, descriptive statistics indicated small increases in mean scores across all twelve attitude statements from pre- to post-assessment, suggesting a general trend toward more positive attitudes following the course (see Table 3).

Negative Attitude. In contrast to the increase in positive attitudes, negative attitudes toward AI did not change significantly. Statistical analysis indicated no significant difference in negative attitude scores (z=–0.792, p=.426), with a slight, non-significant increase in the mean from 2.729 to 2.760. However, analysis of individual items revealed one statistically significant change. Agreement with the statement "Organizations use AI unethically" (N2) increased significantly after the course (z=–2.134, p=.019), suggesting increased awareness regarding ethical implications. Although no other items showed statistically significant changes, minor shifts in response patterns were observed. In particular, small in-

Unit	1	2	3	4	5	6	7	8
Live #	167	121	106	77	75	62	62	56
Video 1	285	231	192	180	172	171	164	162
Video 2	–	217	187	178	171	169	–	162
Live Views	234	188	148	136	135	129	127	33
Assignments	219	209	195	185	181	180	175	175
Quizzes	243	213	190	181	177	178	173	173

Table 1: Learner Engagement in Course Activities.

	Pre	Post	Evaluation
Respondents	306	169	103
Gender			
Male	39.87%	41.42%	45.63%
Female	60.13%	58.58%	54.37%
Age			
18–24	0.65%	1.18	0.97%
25–34	7.52%	4.73%	2.91%
35–44	34.97%	30.77%	33.01%
45–54	38.89%	40.24%	40.78%
55–64	16.67%	21.89%	20.39%
Over 65	1.31%	1.18%	1.94%
Country			
Cyprus	50.33%	50.89%	44.66%
Greece	44.12%	45.56%	48.54%
Other countries	5.56%	3.55%	6.80%
Education level			
High School	2.29%	2.96%	2.91%
Bachelor degree	19.61%	17.75%	19.42%
Master degree	67.32%	71.60%	67.96%
Doctoral degree	9.80%	6.51%	8.74%
Other Diploma	0.98%	1.18%	0.97%
Occupation			
Education	33.66%	32.54%	28.16%
Finance / accounting	9.80%	10.06%	11.65%
Medical	10.13%	10.06%	16.50%
Public administration	18.63%	20.71%	18.45%
Other (e.g., police)	24.51%	24.26%	21.36%
Not currently employed	3.27%	2.37%	3.88%

Table 2: Participants Demographics.

creases were noted in agreement with the statements "AI is used to spy on people," "I shiver with discomfort when I think about future uses of AI," and "I think artificially intelligent systems make many errors." These trends may reflect a growing recognition among participants of the potential risks and limitations associated with AI.

AI Literacy. AILS scores increased significantly (z=–5.821, p<.001) from M=3.782 to M=4.060. Significant differences were observed in *Awareness* (z=–5.626, p<.001.), which increased from M=3.791 to M=4.174. All three items showed significant improvements, including the ability to distinguish between smart and non-smart devices (AW1) (z=–2.798, p=.002), and the ability to identify AI technologies used in everyday applications (AW3) z=–4.515, p<.001). Additionally, there was a significant decrease in the item "I do not know how AI can help me" (AW2) (z=5.185, p<.001). Significant differences were also observed in *Usage* (z=–4.376, p<.001), which increased from M=3.856 to M=4.148. Two items showed significant improvement, showing better ability to skillfully use AI applications to assist with daily tasks (US1) (z=–4.312, p<.001), and use of AI applications to improve their work efficiency (US3) (z=–3.992, p<.001). The item "It is usually hard for me to learn to use a new AI application" (US2) showed a small decrease, suggesting a potential improvement (see Table 4. Significant differences were also found in *Evaluation* (z=–5.269, p<.001), with scores increasing from M=3.493 to M=3.890. Significant improvement were observed in the ability to evaluate capabilities and limitations of AI applications (EV1) (z=–3.552, p<.001), to select the most appropriate solution among several offered by a smart agent (EV2) (z=–4.688, p<.001), and to choose the most suitable AI application for a given task (EV3) (z=–4.919, p<.001). In contrast, we found no significant differences in *Ethics* construct (z=–0.728, p=.466) and items, with a slight increase from M=3.986 to M=4.028.Small increases were observed in the ability to comply with ethical principles when using AI applications (EV1), to remain alert to the potential abuse of AI (ET3), and a slight drop in the item "I am never alert to privacy and information security issues when using AI applications" (ET2).

	Pre		Post	
	Mean	SD	Mean	SD
Positive	**3.793**	**0.454**	**3.871**	**0.476**
P1	4.189	0.587	4.349	0.647
P2	4.266	0.540	4.391	0.618
P3	4.243	0.622	4.414	0.593
P4	4.024	0.763	4.101	0.704
P5	4.243	0.659	4.284	0.692
P6	3.497	0.795	3.515	0.920
P7	4.219	0.612	4.337	0.596
P8	3.953	0.644	3.959	0.649
P9	3.030	0.855	3.160	0.889
P10	3.195	0.796	3.243	0.835
P11	3.604	0.758	3.568	0.878
P12	3.059	0.956	3.130	1.044
Negative	**2.729**	**0.511**	**2.760**	**0.597**
N1	2.775	0.769	2.763	0.895
N2	3.663	0.723	3.805	0.718
N3	1.793	0.731	1.787	0.733
N4	2.751	0.924	2.864	1.006
N5	2.840	0.915	2.852	0.943
N6	2.757	0.849	2.728	0.898
N7	3.018	0.760	3.089	0.815
N8	2.237	0.766	2.195	0.888

Table 3: Pre- and Post-Course Attitudes towards AI.

Course Evaluation

Regarding participants' engagement with the course content, 40% of respondents reported that they typically attended the live sessions, while 34.4% indicated that they watched the recorded live sessions asynchronously. 24.4% reported using a combination of both live and recorded sessions. One participant indicated that he did not attend or watch the live sessions at all. When asked whether they watched the pre-recorded video lectures before attending/watching the live sessions, 51.1% stated that they did so consistently, 32.2% reported doing so occasionally, and 16.7% indicated that they did not watch the pre-recorded material beforehand.

The majority expressed high levels of satisfaction with various aspects of the course. Specifically, 94.4% of respon-

	Pre		Post	
	Mean	SD	Mean	SD
Awareness	**3.791**	**0.585**	**4.174**	**0.484**
AW1	4.036	0.672	4.243	0.583
$AW2^R$	2.527	0.913	1.970	0.767
AW3	3.864	0.690	4.249	0.634
Usage	**3.856**	**0.618**	**4.148**	**0.515**
US1	3.893	0.707	4.240	0.575
$US2^R$	2.148	0.784	1.988	0.794
US3	3.822	0.819	4.183	0.584
Evaluation	**3.493**	**0.681**	**3.890**	**0.514**
EV1	3.609	0.780	3.929	0.642
EV2	3.550	0.779	3.947	0.570
EV3	3.320	0.855	3.793	0.606
Ethics	**3.986**	**0.536**	**4.028**	**0.567**
ET1	4.112	0.727	4.172	0.794
$ET2^R$	2.189	0.852	2.213	0.946
ET3	4.036	0.830	4.124	0.757
AI Literacy	**3.782**	**0.431**	**4.060**	**0.368**

R - reversed items

Table 4: Pre- and Post-Course AI Literacy.

Unit	1	2	3	4	5	6	7	8
#	243	213	190	181	177	178	173	173
score (μ)	7.76	7.81	6.97	8.72	8.74	7.47	9.20	8.95
median	8	8	7	9	10	8	10	9

Table 5: Learner performance (# responses to the quizzes (Q), mean/median score (out of 10)).

dents selected the higher end of the Likert scale (4 or 5 on a five-point scale) to indicate that the pre-recorded video lectures were helpful, while 85.6% found the live sessions to be helpful as well. 92.2% agreed that the instructional approach used by the teaching team was effective. Similarly, 92.2% reported that the structure of the course was clear and effective, and 93.3% believed that the educational content and accompanying materials were both understandable and accessible. Additionally, 93.3% agreed that the course included sufficient examples and explanations to support their understanding of the material. With regard to assessments, 90% stated that the assignments helped them to better understand the course content, and the same proportion found the assignments to be interesting. Furthermore, 84.4% indicated that they found it engaging to see how the data from the assignments was analyzed and discussed during the live sessions. Overall, 85.6% reported that the course met their expectations. Finally, 90% agreed with the statement that the course is potentially "useful for everyone," highlighting its perceived value and broader applicability.

After completing the course, the vast majority reported improvements in their understanding of key AI-related concepts. Specifically, 96.7% indicated that they were able to better understand the behavior of AI applications, and 94.4% felt more confident in analyzing the functions of "everyday" AI applications. Additionally, 94.4% stated that they were able to explain the ethical and social issues that may arise from the use of AI technologies. A large proportion (91.1%) also reported being able to recognize that methods such as ML are foundational to the development of AI applications. Furthermore, most indicated that they were able to identify the use of different AI technologies in real-world contexts, including vision technologies (92.2%), language technologies (91.1%), search technologies (94.4%), and personalization technologies (91.1%). They reported a high ability to understand the potential risks associated with AI applications (93.3%). These self-reported learning gains were supported by the quiz results, which showed generally high average scores, indicating strong comprehension (Table 5).

When asked whether there were any topics they wished had been covered in more depth, 55 participants provided open-ended responses. Participants expressed a desire to learn more about *"the actual designing and functions of algorithms and the development process of AI-powered apps."* participant 5 – p5). Others suggested including more examples of real-world AI applications (p36), and placing greater emphasis on practical training in the use of AI tools (p39, p43, p65). For instance, one participant noted the value of including *"a topic on how to properly use the various AI tools such as ChatGPT that [they] use daily."* (p38). Some highlighted the importance of exploring further *"issues related to ethics"* (p83) and risks associated with AI. A few proposed the development of a more advanced program that delves deeper into technical implementation (p61, p81), including topics such as *"creating programs in Python utilizing AI to solve problems as well as automate processes"* (p62). In the final open-ended question about general comments, some emphasized the broader importance of AI education, stating *"that more people should be educated about AI in everyday life."* (p75), while another suggested that *"this particular course should be given regularly so that as many people as possible can attend it because it is a topic that most people come into contact with daily."* (p41).

Drop-out. We noticed that 176 out of the 343 registered for the course dropped out, and we invited them to complete the online evaluation, with 13 of them completing the form. They were asked to answer five 5-point Likert-scale items aimed at understanding the reasons they dropped out, and one free-text question asking them to specify why they did not complete the course. The majority (69%) indicated that they did not complete the course because they were very busy during the period it was offered and did not have enough time. Most (62%) also stated that the lack of time to complete the assignments was one of the main reasons for not finishing the course. However, 62% did not consider the course to be more time-consuming than expected. Additionally, most disagreed that dissatisfaction with the instructors (85%) or that the course did not meet their expectations (77%) played a role in their decision not to complete the course. Four participants further explained that the timing and day of the course were not convenient for them (p14), or that they had simply *"just ran out of time"* (p13).

Reflection & Concluding Remarks

Our first goal was to design and implement a large-scale course on AI, and secondly, to understand its impact on

participants' attitudes toward AI and their AI literacy. Our findings build on existing work (Long and Magerko 2020), which underscores the need to equip individuals with the skills to critically engage with AI, both by harnessing its benefits effectively and by assessing its potential risks.

Attitudes Towards AI. The analysis of participants' attitudes toward AI yields important insights into the ways educational interventions can influence perceptions of AI. Similar to (Zhang et al. 2022), our participants joined the course with a positive attitude towards AI, which was not surprising given that the course was a self-selected program. Notably, the course fostered an increase in positive attitudes, with participants expressing greater interest in integrating AI in their daily lives and showing enhanced enthusiasm for its potential. These findings are consistent with prior research indicating that AI education can promote more informed and favorable attitudes toward AI (Kwon et al. 2024; Kasinidou, Kleanthous, and Otterbacher 2023; Kong, Cheung, and Zhang 2022). Interestingly, while negative attitude did not shift significantly, the item, "Organizations use AI unethically" showed a significant increase. This points to a heightened awareness of ethical concerns, which was one of the main objectives of our course. Like in (Bewersdorff et al. 2023; Vandenberg and Mott 2023), this shift may suggest that even when participants retain their enthusiasm for AI, they can simultaneously develop a more critical understanding of its ethical implications.

AI Literacy. Our participants' AI literacy was also improved, demonstrating increased progress across multiple constructs of AI literacy, suggesting that non-technical courses can have a meaningful impact on learners' abilities to engage with AI. Awareness and Usage constructs increased significantly after the completion of the course, with participants reporting enhanced ability to recognize AI in everyday contexts and greater confidence in utilizing AI applications to assist with daily tasks and improve productivity. This reflects a shift away from abstract notions of AI toward a more practical, functional understanding, consistent with findings from previous studies (Kong, Cheung, and Zhang 2022, 2023). Similar to prior work, we found that AI education enables learners to become more critical consumers and fosters more realistic expectations of AI applications (Candon et al. 2025; Kong, Korte, and Cheung 2023). Participants also demonstrated a notable improvement in Evaluation skills, showing stronger abilities to assess the capabilities and limitations of AI, make informed decisions between different tools, and choose technologies that fit specific tasks. These results are based on previous work that highlights that AI education not only improves how participants use AI, but also strengthens critical thinking needed to evaluate it (Kasinidou et al. 2024, 2021; Lee et al. 2021a). In contrast, the Ethics component did not show significant changes, though slight improvements were observed. This may suggest that ethical reflection requires more time, depth, or targeted discussion than what a single introductory course can provide. Still, even the small changes in participants' self-reported ethical concerns are promising and align with findings from earlier studies showing that even brief exposure to AI ethics can begin to stimulate ethical awareness (Kasinidou, Kleanthous, and Otterbacher 2023; Gómez, Dabbah, and Benotti 2024).

Course Evaluation Most participants reported satisfaction with both the structure and delivery of the course, including the clarity of content, helpfulness of pre-recorded videos and live sessions, and the effectiveness of the teaching approach. This is in line with research highlighting the importance of accessible, well-supported learning environments in fostering meaningful learning outcomes in AI education (Lee et al. 2021a; Kasinidou et al. 2021). The flexibility of engaging with content either synchronously or asynchronously appears to have supported different learning needs, with many successfully using a mix of formats. In addition, the integration of real-life examples, the presentation of assignments' data analysis during the live session, and open-ended discussions helped participants not only absorb knowledge but also apply it practically (Kasinidou et al. 2024; Kong, Cheung, and Zhang 2023).

Overall, the course succeeded in fostering a more balanced and informed understanding of AI. Participants gained a clearer sense of both the opportunities and challenges posed by AI. Our findings reinforce that accessible, and non-technical educational interventions can meaningfully enhance AI literacy in the general public (Kwon et al. 2024; Kasinidou et al. 2024). Participants not only became more aware of how AI shows up in their everyday lives, but also built the skills to engage with it more critically echoed in earlier research on practical AI learning (Kong, Cheung, and Zhang 2023; Candon et al. 2025). As AI continues to permeate daily life, these results add to the growing recognition that AI literacy is not optional, but a foundational competency for navigating the digital world (Bewersdorff et al. 2023; Kasinidou, Kleanthous, and Otterbacher 2023).

Limitations and Future Work. We note several limitations that should be considered when interpreting this work. While 343 participants registered to the course, only 169 completed the required assessments and received a CoC. This indicates a notable dropout rate and partial engagement, highlighting the need for future work to better understand learners' motivations and the barriers to sustained participation. Furthermore, since participants self-selected the course and were likely already motivated to learn about AI, the findings may not be generalizable to the broader population. Third, we relied on self-reported data, which may not fully capture actual skill development or long-term impact.

We plan to continue offering this course through our university's platform on a semester-based schedule, keeping it open to the general public. Given the observed dropout rates and the self-selected nature of our current audience, our next steps will focus on attracting a more diverse and representative group across different age ranges and social backgrounds. This broader reach will help us identify barriers to sustained participation and advance our goal of making AI education accessible to all. It will also enable us to evaluate the course at a larger scale, refine its structure, and develop targeted strategies to improve engagement and retention.

Acknowledgments

This project is partially funded by the EU's Horizon Europe Programme, agreement No. 101214000 (TRUMAN), the Cyprus Research and Innovation Foundation Programme, agreement No. BRIDGE2HORIZON/0823E/0203 (PINNACLE) and the EU's Erasmus+ Programme, agreement No. 2025-1-CY01-KA220-HED-000364525 (AI4Everyone).

References

Almatrafi, O.; Johri, A.; and Lee, H. 2024. A systematic review of AI literacy conceptualization, constructs, and implementation and assessment efforts (2019-2023). *Computers and Education Open*, 100173.

Arn, L.; and Huang, E. M. 2024. "Robots Can Do Disgusting Things, but Also Good Things": Fostering Children's Understanding of AI through Storytelling. *ACM Transactions on Computing Education*, 24(3): 1–55.

Bewersdorff, A.; Zhai, X.; Roberts, J.; and Nerdel, C. 2023. Myths, mis-and preconceptions of artificial intelligence: A review of the literature. *Computers and Education: Artificial Intelligence*, 100143.

Birhane, A.; Isaac, W.; Prabhakaran, V.; Diaz, M.; Elish, M. C.; Gabriel, I.; and Mohamed, S. 2022. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–8.

Biswas, J.; Fussell, D.; Stone, P.; Patterson, K.; Procko, K.; Sabatini, L.; and Xu, Z. 2025. The Essentials of AI for Life and Society: An AI Literacy Course for the University Community. *Proceedings of the AAAI Conference on Artificial Intelligence*.

Candon, K.; Georgiou, N. C.; Ramnauth, R.; Cheung, J.; Fincke, E. C.; and Scassellati, B. 2025. Artificial Intelligence for Future Presidents: Teaching AI Literacy to Everyone. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 28988–28995.

Druga, S.; and Ko, A. J. 2021. How do children's perceptions of machine intelligence change when training and coding smart programs? In *Proceedings of the 20th Annual ACM Interaction Design and Children Conference*, 49–61.

Gilman, M. E. 2023. Democratizing AI: Principles for meaningful public participation. *Data & Society*.

Gómez, M. J.; Dabbah, J.; and Benotti, L. 2024. A workshop on artificial intelligence biases and its effect on high school students' perceptions. *International Journal of Child-Computer Interaction*, 100710.

Hollands, F. M.; DiPaola, D.; Breazeal, C.; and Ali, S. 2024. AI Mastery May Not Be For Everyone, But AI Literacy Should Be. In *Proceedings of the 2024 on ACM Virtual Global Computing Education Conference V. 1*, 60–66.

Johnson, D.; Alsharid, M.; El-Bouri, R.; Mehdi, N.; Shamout, F.; Szenicer, A.; Toman, D.; and Binghalib, S. 2022. An experience report of executive-level artificial intelligence education in the united arab emirates. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 12766–12773.

Kasinidou, M.; Kleanthous, S.; Busso, M.; Rodas, M.; Otterbacher, J.; and Giunchiglia, F. 2024. Artificial Intelligence in Everyday Life 2.0: Educating University Students from Different Majors. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1*.

Kasinidou, M.; Kleanthous, S.; Orphanou, K.; and Otterbacher, J. 2021. Educating Computer Science Students about Algorithmic Fairness, Accountability, Transparency and Ethics. In *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education*.

Kasinidou, M.; Kleanthous, S.; and Otterbacher, J. 2023. Artificial Intelligence in Everyday Life: Educating the Public Through an Open, Distance-learning Course. In *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1*.

Kasinidou, M.; Kleanthous, S.; and Otterbacher, J. 2026a. AI in the Eyes of Middle Schoolers: Perceptions, Attitudes, and Literacy. In *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V.1 (to appear)*.

Kasinidou, M.; Kleanthous, S.; and Otterbacher, J. 2026b. 'What Do Children Think About AI?': Insights and Educational Implications from Primary School Students' Perceptions of AI. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Kasinidou, M.; Kleanthous, S.; Otterbacher, J.; and Christoforou, E. 2025. PINNACLE: Public Engagement in AI Evaluation for Education and SystemBenchmarking. In *Proceedings of the 2nd International Conference on Digital Skills and IT Professionalism.*

Kasinidou, M.; Kleanthoys, S.; and Otterbacher, J. 2025. Cypriot teachers' digital skills and attitudes towards AI. *Discover Education*, 4(1): 1.

Kelley, P. G.; Yang, Y.; Heldreth, C.; Moessner, C.; Sedley, A.; Kramm, A.; Newman, D. T.; and Woodruff, A. 2021. Exciting, Useful, Worrying, Futuristic: Public Perception of Artificial Intelligence in 8 Countries. In *Proceedings of the 2021 aaai/acm conference on ai, ethics, and society*.

Knapper, C.; and Cropley, A. J. 2000. *Lifelong learning in higher education*. Psychology Press.

Kong, S.-C.; Cheung, W. M.-Y.; and Tsang, O. 2023. Evaluating an artificial intelligence literacy programme for empowering and developing concepts, literacy and ethical awareness in senior secondary students. *Education and Information Technologies*, 28(4): 4703–4724.

Kong, S.-C.; Cheung, W. M.-Y.; and Zhang, G. 2021. Evaluation of an artificial intelligence literacy course for university students with diverse study backgrounds. *Computers and Education: Artificial Intelligence*, 2: 100026.

Kong, S.-C.; Cheung, W. M.-Y.; and Zhang, G. 2022. Evaluating artificial intelligence literacy courses for fostering conceptual learning, literacy and empowerment in university students: Refocusing to conceptual building. *Computers in Human Behavior Reports*, 7: 100223.

Kong, S.-C.; Cheung, W. M.-Y.; and Zhang, G. 2023. Evaluating an artificial intelligence literacy programme for developing university students' conceptual understanding, lit-

eracy, empowerment and ethical awareness. *Educational Technology & Society*, 26(1): 16–30.

Kong, S. C.; Korte, S.-M.; and Cheung, W. M.-Y. 2023. Nurturing Artificial Intelligence Literacy in Students with Diverse Cultural Backgrounds. In *International Conference on Innovative Technologies and Learning*. Springer.

Kupferschmidt, K. L.; Wan, F.; Alvarez, J. C.; Castano, D. G.; Taylor, G. W.; and Kocak, S. A. 2025. Bridging the AI Gap: Evaluating the Impact of an AI Education Program for Caregivers on Parental Leave. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Kwon, K.; Kim, K.; Ottenbreit-Leftwich, A. T.; Glazewski, K. D.; Brown, M. L.; Bae, H.; and Closser, F. M. 2024. Exploring AIFORGOOD Summer Camp Curriculum to Foster Middle School Students' Understanding of Artificial Intelligence. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 2*, 1710–1711.

Laupichler, M. C.; Aster, A.; Schirch, J.; and Raupach, T. 2022. Artificial intelligence literacy in higher and adult education: A scoping literature review. *Computers and Education: Artificial Intelligence*, 100101.

Lee, I.; Ali, S.; Zhang, H.; DiPaola, D.; and Breazeal, C. 2021a. Developing Middle School Students' AI Literacy. In *Proceedings of the 52nd ACM technical symposium on computer science education*, 191–197.

Lee, I.; Ali, S.; Zhang, H.; DiPaola, D.; and Breazeal, C. 2021b. *Developing Middle School Students' AI Literacy*.

Long, D.; and Magerko, B. 2020. What is AI literacy? Competencies and design considerations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.

Long, D.; Teachey, A.; and Magerko, B. 2022. Family learning talk in AI literacy learning activities. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.

Maher, M. L.; and Tadimalla, S. Y. 2024. Increasing Diversity in Lifelong AI Education: Workshop Report. In *Proceedings of the AAAI Symposium Series*, volume 3, 493–500.

Masla, J.; Bosch, C.; Ravi, P.; Guterman, L.; Wharton, S.; Gustafson-Quiett, M. C.; Hegly, S. A.; Macatantan, C.; Klopfer, E.; Breazeal, C.; et al. 2025. Supporting AI Literacy Teaching Through the Development of Assessments for Classroom Use. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 29178–29185.

Ng, D. T. K.; Lee, M.; Tan, R. J. Y.; Hu, X.; Downie, J. S.; and Chu, S. K. W. 2022. A review of AI teaching and learning from 2000 to 2020. *Education and Information Technologies*, 1–57.

Ng, D. T. K.; Leung, J. K. L.; Chu, S. K. W.; and Qiao, M. S. 2021. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*.

Pham, L.; O'Sullivan, B.; Scantamburlo, T.; and Mai, T. 2024. Addressing Digital and AI Skills Gaps in European Living Areas: A Comparative Analysis of Small and Large Communities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 23119–23127.

Sabuncuoglu, A. 2020. Designing One Year Curriculum to Teach Artificial Intelligence for Middle School. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education*, ITiCSE '20.

Schepman, A.; and Rodway, P. 2020. Initial validation of the general attitudes towards Artificial Intelligence Scale. *Computers in human behavior reports*.

Schiavo, G.; Businaro, S.; and Zancanaro, M. 2024. Comprehension, apprehension, and acceptance: Understanding the influence of literacy and anxiety on acceptance of artificial Intelligence. *Technology in Society*, 77: 102537.

Southworth, J.; Migliaccio, K.; Glover, J.; Glover, J.; Reed, D.; McCarty, C.; Brendemuhl, J.; and Thomas, A. 2023. Developing a model for AI Across the curriculum: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence*.

Stadelmann, T.; Keuzenkamp, J.; Grabner, H.; and Würsch, C. 2021. The AI-atlas: didactics for teaching AI and machine learning on-site, online, and hybrid. *Education Sciences*, 11(7): 318.

Stoyanovich, J.; Lewis, A.; Corbett, E.; Bynum, L. E.; Rosenblatt, L.; and Khan, F. A. 2025. We Are AI: Taking Control of Technology. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 29070–29077.

Tadimalla, S. Y.; and Maher, M. L. 2025. Sociotechnical AI Education Course Design for CS Majors and Non-Majors. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 2*, 1631–1632.

Touretzky, D.; Gardner-McCune, C.; Martin, F.; and Seehorn, D. 2019. Envisioning AI for K-12: What should every child know about AI? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 9795–9799.

Vandenberg, J.; and Mott, B. 2023. " AI Teaches Itself": Exploring Young Learners' Perspectives on Artificial Intelligence for Instrument Development. In *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1*, 485–490.

Wang, B.; Rau, P.-L. P.; and Yuan, T. 2023. Measuring user competence in using artificial intelligence: validity and reliability of artificial intelligence literacy scale. *Behaviour & information technology*.

Wollowski, M.; Selkowitz, R.; Brown, L.; Goel, A.; Luger, G.; Marshall, J.; Neel, A.; Neller, T.; and Norvig, P. 2016. A survey of current practice and teaching of AI. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Wunderlich, L.; Higgins, A.; and Lichtenstein, Y. 2021. *Machine Learning for Business Students: An Experiential Learning Approach*.

Yue, M.; Jong, M. S.-Y.; and Dai, Y. 2022. Pedagogical design of K-12 artificial intelligence education: A systematic review. *Sustainability*, 14(23): 15620.

Zhang, H.; Lee, I.; Ali, S.; DiPaola, D.; Cheng, Y.; and Breazeal, C. 2022. Integrating Ethics and Career Futures with Technical Learning to Promote AI Literacy for Middle School Students: An Exploratory Study. *International Journal of Artificial Intelligence in Education*, 1–35.