Check for updates

CASE STUDY

How to train the upcoming generation of AI researchers? Design principles and results from the Ethos+Tekhne initiative as a case study

[version 1; peer review: 4 approved with reservations]

Stefano Vrizzi[1,2], Alberto Arturo Vergani[3,4], Laura State[5,6], Filippo Galli[5], Amric Trudel[2,7], Sara Mazzucato[4,8]

[1]UB School of Economics, Universitat de Barcelona, Barcelona, 08034, Spain
[2]Département d'Études Cognitives, École Normale Supérieure, PSL Research University, Paris, 75005, France
[3]Department of Brain and Behavioral Sciences, Università di Pavia, Pavia, 27100, Italy
[4]The BioRobotics Institute and Department of Excellence in Robotics and Artificial Intelligence, Scuola Superiore Sant'Anna, Pontedera, Pisa, 56025, Italy
[5]Scuola Normale Superiore, Pisa, 56126, Italy
[6]Università di Pisa, Pisa, 56127, Italy
[7]Centre de recherche en Épidémiologie et Santé des Populations, Université Paris-Saclay, UVSQ, INSERM, Paris, France
[8]Department of Medical Informatics, Erasmus Medical Center, Rotterdam, 3015, The Netherlands

v1 **First published:** 22 May 2026, **6**:174
https://doi.org/10.12688/openreseurope.20225.1
Latest published: 22 May 2026, **6**:174
https://doi.org/10.12688/openreseurope.20225.1

Abstract

A lasting challenge for the next generation of AI researchers and professionals will be to shape a global AI infrastructure capable of accommodating stakeholder heterogeneity, while addressing technical constraints. It becomes crucial to provide early-career researchers and professionals in AI with the tools and the knowledge to understand the complexity of this challenge, and with the technical skills to partake in it. Here, we present core design features and results from a novel training initiative combining AI and ethics. This initiative is called "Ethos+Tekhne" because it centers on the discussion of two main questions: what kind of AI we want (*ethos*), and how to achieve it concretely (*tekhne*). Its programme was designed around three key principles: pluralism, narrative, and balance. In terms of pluralism, the training combined several types of theoretical and practical learning activities, as well as speakers from a variety of fields, including non-academic sectors and citizens. With regard to narrative, the training followed a programme-wide thread of thematic days. Finally, with respect to balance, on-site participant group size was limited to 30–40 individuals, in order to offer each participant the space and time to actively engage in discussions, while supporting

Open Peer Review

Approval Status ? ? ? ?

	1	2	3	4
version 1 22 May 2026	? view	? view	? view	? view

1. **Febblina Daryanes**, Universitas Riau, Pekanbaru, Indonesia
2. **Rajiansyah Rajiansyah**, Universitas Mulawarman, Samarinda, Indonesia
3. **Louis Robert C. Sison**, Bulacan State University, Malolos, Philippines
4. **Fei Ping Por**, Tunku Abdul Rahman University of Management and Technology, Kuala Lumpur, Malaysia

Any reports and responses or comments on the article can be found at the end of the article.

online attendance to promote open access. Moreover, the programme schedule aimed to balance activity and break time to encourage participants to exchange their thoughts in informal settings. The hallmark of success of this initiative has been its regenerative nature: some first-edition participants led the organisation of the second edition, and some second-edition participants are currently organising the third edition. Content-wise, three axes of action in ethical AI development emerged from participants' discussion: *AI for Good*, *Good for AI*, and *Good AI.* We report our rationale, programme design, and results, in the hope that these could provide material for future training programmes in AI, and potentially other technologies.

Keywords
Artificial Intelligence, Education, Ethics, Human-centered AI, Interdisciplinarity, Pluralism, Society

This article is included in the Horizon 2020 gateway.

This article is included in the Horizon Europe gateway.

This article is included in the Marie-Sklodowska-Curie Actions (MSCA) gateway.

This article is included in the Erasmus+ gateway.

This article is included in the Artificial Intelligence and the Social Sciences and Humanities collection.

Corresponding author: Stefano Vrizzi (vrizzi@ub.edu)

Author roles: Vrizzi S: Conceptualization, Data Curation, Formal Analysis, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Software, Supervision, Validation, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Vergani AA**: Conceptualization, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Supervision, Writing – Original Draft Preparation, Writing – Review & Editing; **State L**: Conceptualization, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Supervision, Writing – Original Draft Preparation, Writing – Review & Editing; **Galli F**: Conceptualization, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Writing – Original Draft Preparation; **Trudel A**: Formal Analysis, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Writing – Original Draft Preparation, Writing – Review & Editing; **Mazzucato S**: Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Writing – Original Draft Preparation, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: The "Ethos+Tekhne: a new generation of AI researchers" spring school was funded by the first and the third EELISA Joint Calls for Inter-Institutional Activities under the European Union's Erasmus+ programme (Grant Agreements N. 101004081 and 101124676) and from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement N. 835294). Its last session at Cinema Arsenale in Pisa during the first edition was financed by Razionalisti Pisa. Stefano Vrizzi was funded by the AI4theSciences PhD programme (Grant Agreement N. 945304, Marie Skłodowska-Curie Actions-COFUND, Horizon 2020 programme) and by The ToBe project (Grant Agreement N. 101094211, Horizon Europe Research and Innovation Programme).
The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2026 Vrizzi S *et al*. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Vrizzi S, Vergani AA, State L *et al.* **How to train the upcoming generation of AI researchers? Design principles and results from the Ethos+Tekhne initiative as a case study [version 1; peer review: 4 approved with reservations]** Open Research Europe 2026, **6**:174 https://doi.org/10.12688/openreseurope.20225.1

First published: 22 May 2026, **6**:174 https://doi.org/10.12688/openreseurope.20225.1

1. Introduction

1.1 Rationale

Artificial Intelligence (AI) is projected to become a global infrastructure capable of making autonomous decisions and mimicking human cognitive skills. Its development inevitably concerns human rights. Firstly, becoming a global infrastructure means fundamentally shaping the lives of individuals and the evolution of society. Secondly, unlike previous industrial revolutions that primarily automated routine tasks, the AI innovation wave targets higher-order reasoning and autonomous decision-making skills. It is therefore crucial to prepare the upcoming generation of AI researchers and professionals to embrace responsibly the impact of their work.

Motivated by these challenges, we convened in 2022 as a group of early-career researchers and professionals from across Europe. The goal was to devise an opportunity for early-career researchers and professionals in AI to discuss two core questions: 1) what kind of AI to pursue and 2) how to achieve it, i.e. how to combine *ethos* and *tekhne* (Srikumar et al. 2022).[1] This is the rationale underlying the creation of the "Ethos+Tekhne"[2] school for early-career researchers and professionals in AI.

Our initiative contributes to the growing ecosystem of interdisciplinary study courses and programmes, research groups, and conferences around Europe and globally that combine the technical perspective on AI with legal, ethical and philosophical considerations. In support of this momentum, this paper offers key insights to devise an interdisciplinary seasonal school, hoping to provide useful elements that other training programmes on emerging technologies can adopt, leverage and improve.

1.2 Structure

In this paper, we present design principles and results from the first two editions of the Ethos+Tekhne school, held in Pisa, Italy, in 2023 and Paris, France, in 2025. We structure the paper in two main sections: one outlining the underlying design principles and differences of the two training programmes (Section 2), and the other one discussing the outcomes from this initiative, by highlighting the axes of action in ethical AI development that emerged from participants' discussions, their feedback, and the ripple effects at team and institutional level (Section 3). We conclude by summarising the innovative approaches taken and their significance in preparing future researchers for the ethical challenges in building the global infrastructure of AI.

2. Methods

We designed the programme based on three main principles: pluralism (Section 2.1), narrative (Section 2.2), and balance (section 2.3). Each principle underpinned several factors that contributed to the value and success of Ethos+Tekhne.

Ethics and Consent

This work reports participant data from both editions. All data were collected anonymously through online questionnaires, and with written informed consent from the respondents to participate in this research. In France, an ethical opinion is only mandatory under the Jardé law for scientific research that improves biological or medical knowledge and/or changes the care of subjects.

2.1 Pluralism

The most fundamental pillar to design our programme was pluralism: multidisciplinary perspectives, diverse learning activities, and heterogeneous backgrounds of participants.

2.1.1 Topics, speakers, and organising team

One key element was the originality and boldness in conceiving, proposing, and implementing a horizontal educational scenario within the traditionally vertical sci-tech pillar of AI. This approach enabled the integration of interdisciplinary perspectives and fostered dialogue between ethics, technology, and society (Sections 2.2.1–2, and Table 1A-B). Moreover, the diversity of speakers further enriched the experience. Speakers differed in terms of gender, age, sector of expertise, and career stage, ranging from AI researchers and big-tech engineers to philosophers, lawyers, economists,

[1]"*An analysis of preprints found that researchers assessed impact through the narrow lens of technical contributions.* […] *Work to better understand how AI researchers view their responsibility about the effects of their work in light of changing social contexts is also crucial.*" (Srikumar et al. 2022).

[2]More information available at the EELISA community webpage of "Ethos+Tekhne": https://digitalcampus.eelisa.eu/community/d5b455fd-4da7-45bf-9152-2730bd03d962/details.

policy advisors, journalists, and Members of the European Parliament (MEPs). This diversity was an outcome of the diversity in terms of network, resources, and experience within the organising team, which included early-career researchers in AI, both STEM and non-STEM, as well as a journalist and a young professional from the private sector. The heterogeneity of topics and speakers created a dynamic environment where multiple voices could intersect, contributing to a more inclusive and multifaceted understanding of the challenges and opportunities at the interface of AI and ethics.

2.1.2 Learning activities

We also diversified learning activities. We planned theoretical lectures in the morning, to provide working knowledge to all participants, followed by more practical sessions in the afternoon, to leverage the content learnt in the morning, with activities ranging from discussion panels to workshops, coding tutorials, group writing, and a moot court competition (Fig. 1). Moreover, we further broadened the discussion by organising a cinema night where participants, as future AI developers and professionals, met with the local community for an open conversation on AI development, fostering a mutual exchange of knowledge and opinions (Seismic Foundation 2025).

2.1.3 Target audience and participation

The training programme targeted early-career researchers (PhD students, post-docs, master's students, research assistants). In terms of participant profile (Fig. 2A), the first edition was composed mainly of PhD students. In contrast, the second edition featured a more diverse group, including a greater number of master's students and the presence of independent participants, who had not been represented previously. Regarding academic fields (Fig. 2B), the first edition targeted primarily participants from STEM (science, technology, engineering, and mathematics), while the second edition showed a more balanced distribution. Participation from philosophy, law, and medicine increased, and field representation expanded to include the social sciences and management. This shift reflects a deliberate effort to broaden the scope of the initiative, making it more inclusive across both career stages and disciplinary perspectives. In both editions, heterogeneity of backgrounds and educational levels played a key role in sparking debates on AI that combined both STEM and non-STEM perspectives.

2.2 Narrative

The narrative of the school programme was based on starting with *why* (Sinek 2009), before discussing the *how*. The aim was to first create the space for participants to form their own intrinsic motivation to engage with the impacts of AI, and then address the ethical and technical constraints in AI development. We placed the most technical details about AI algorithms mid-programme, following an hourglass model (Fig. 1). The programme's topics were inspired by the Ethics Guidelines for Trustworthy AI (High-Level Expert Group on Artificial Intelligence 2019), and designed to promote the Sustainable Development Goals (SDGs), focusing on SDG 5 (Gender equality), SDG 8 (Decent work and economic growth), SDG 9 (Industry, Innovation and Infrastructure), SDG 10 (Reduced inequalities), SDG 16 (Peace, Justice and Strong Institutions).

Table 1. The day-specific questions addressed by the first (A) and second (B) editions of Ethos+Tekhne, with a summary of the daily topics of the training programmes to address such questions. Daily questions, topics, and their order were chosen to create a programme-wide narrative based on these questions (Sections 2.2.1 and 2.2.2).

A		
First edition		
Day	**Question addressed**	**Topic**
1	Why should we care about AI?	The socio-economic impact of AI
2	What are the legal constraints within which we should operate in AI?	EU regulations for a fair and private AI
3	Once aware of the legal constraints, how can we gain a deep technical understanding of AI?	eXplainable AI (XAI)
4	Once we have gained a technical understanding of AI models, how can we improve their fairness?	AI bias identification and trustworthiness assessment
5	Once all the legal and technical precautionary measures have been adopted, what are the legal consequences in accidents involving AI?	Legal liability in AI

Table 1. *Continued*

B		
Second edition		
Day	**Question addressed**	**Topic**
1	Why should we care about AI?	The socio-economic impact of AI
2	How can we design and deploy Generative AI responsibly, ensuring technical soundness and ethical alignment?	The technical and ethical dimensions of Generative AI
3	What are the legal constraints and compliance challenges in AI?	The EU regulatory framework for AI systems, intellectual property and copyright

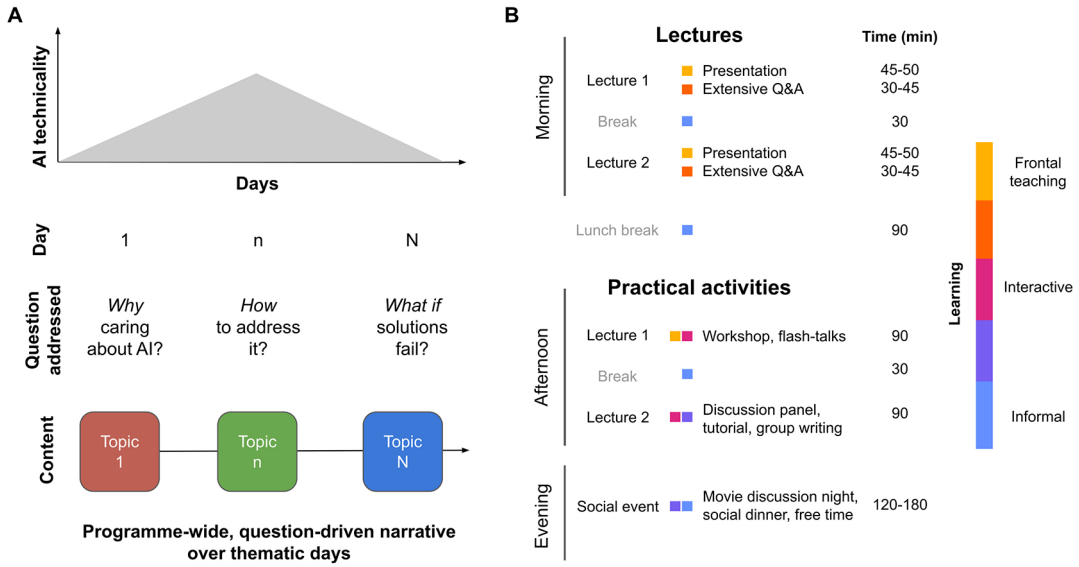

Fig. 1. General programme structure: between-days (A) and within-day (B). The programme was designed around a main narrative organised into thematic days, each addressing a specific question (panel **A**). The general questions driving the narrative were "why", "how", and "what if it fails?" (see Table 1A-B for more details). While the first day aimed to spark motivation, the following days addressed how to address problems and barriers. More specifically, we set the most technical details about AI algorithms mid-programme, following the popular hourglass model for writing. The last day was devoted to discussing what happens if solutions fail, leading to unwanted consequences of AI deployment. Within-day (panel **B**), we planned four main time slots per day. We scheduled more theoretical content and frontal sessions in the first half of the day, to provide foundational knowledge, while more practical and interactive activities in the second half. Each session lasted about 75–90 minutes. In the case of lectures, we shortened time for presentations, allocating more time for Q&A or debate. Sessions were separated by 30-min breaks, and 1.5 h lunch break, to encourage participants to digest information over informal chats, favouring quality over quantity of learning.

2.2.1 First edition programme

In the first edition, we designed a programme over five thematic days (Table 1A).[3] Following our narrative principles, we opened the programme by discussing the importance of engaging with AI as an emerging global infrastructure. Accordingly, the first day focused on the impacts of AI on society, individuals, and the economy. The morning sessions analysed how AI could affect markets, employment and inequalities (Coveri et al. 2021), and prompted a debate on the political and cultural aspects of deskilling, reskilling, and upskilling individuals (Pianta 2020).[4] The afternoon sessions introduced the concept of the "unheimlich" (uncanny) as a sensation humans may experience when interacting with

[3]Programme of the first edition available here: https://www.santannapisa.it/en/alta-formazione/eelisa-spring-school-ethos-tekhne-new-generation-ai-researchers.

[4]For instance, Pianta raised a question on outsourcing cooking to robots. This shift could impact individuals' skills, and in turn affect the cultural heritage and traditions of a whole society, especially in countries with a deeply rooted culinary identity.

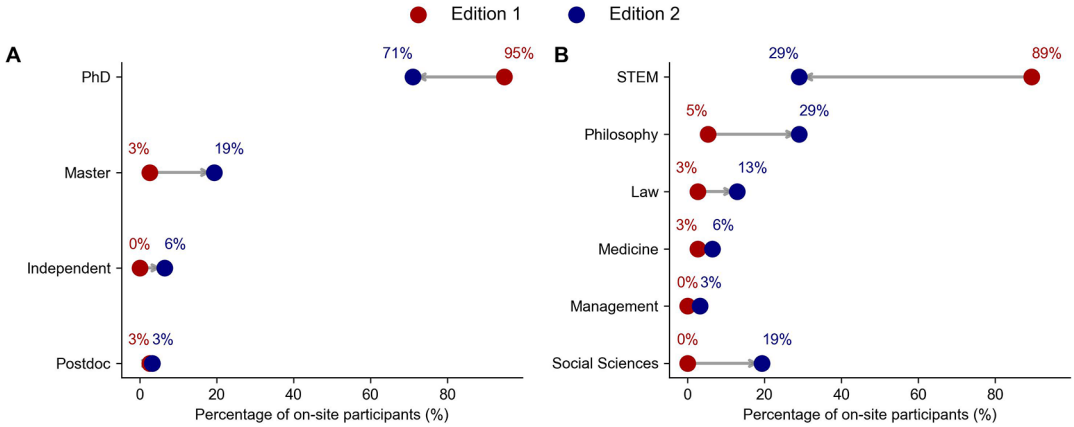

Figure 2. Aggregate participants' profile in first and second edition: current job or educational level (A) and background domain (B). In the first edition (red), the number of total on-site participants was 38. All participants were enrolled as PhD students, with the exception of one master student and one post-doc. Two students came from a background in philosophy, one from law, one from medicine, and the remaining 34 from STEM subjects. In the second edition (blue), 31 participants came together, including 22 PhD students, 6 master students, two independent or aspiring researchers, and one post doc. This edition attracted participants from a broader variety of backgrounds: 9 participants were from STEM fields, 9 from philosophy, 6 from social sciences, four from law, two from medicine and one from management.

robots (Henry 2025), and concluded with a panel moderated by a journalist that welcomed the perspective of big-tech engineers on AI development (Murray 2021).

The second day was designed to discuss the legal boundaries within which AI developers and professionals would operate. The day began with a debate on the differences between the application of ethics and law in AI (Comandé 2022), followed by a technical explanation of person identification from images (Zhou 2019). The afternoon was devoted to the AI Act (The European Parliament and The Council of the European Union 2024). This risk-based EU regulatory framework was first introduced by a policy advisor, and then discussed directly with two MEPs,[5] including the AI Act co-rapporteur, and the founder of an ethical AI consultancy company (Huyskes 2025), in the form of a panel moderated by a journalist.

The third day stepped into the technicalities of the "black box" problem, to better understand the functioning of machine learning and AI models. Lectures presented approaches to explainable AI (XAI) such as post-hoc methods (Samek et al. 2021), and two applicative scenarios, in energy trading (Parginos et al. 2022) and image classification (Balayn et al. 2022). A hands-on tutorial then showed the code to use XAI techniques (Giannotti, Naretto, and Bodria 2023), such as LIME, LORE and SHAP.

The fourth day stressed the importance of algorithmic design, choosing healthcare as the focus domain for its broad appeal. Lectures highlighted open-access solutions against the risk of exacerbating inequalities (Norori et al. 2021), the challenges for the patient-caregiver relationship (de Miguel et al. 2020) and an open-source Python toolkit for algorithmic fairness (Bellamy et al. 2019). Finally, an engineer and a philosopher guided participants in the identification of ethical risks, technical flaws and sources of bias when designing and deploying deep learning algorithms in hospitals (Cacciatore et al. 2022).

On the fifth day, at first, participants debated liability in an AI-related legal case study through a moot court competition, representing the parties involved (Franco 2023). To elaborate on the school's content, participants then wrote an extended abstract in small groups on their favorite programme topic. Finally, participants met citizens for an open discussion, led by a philosopher (Merenda 2020) and a computer science professor (Giannotti, Naretto, and Bodria 2023) at the local cinema. We screened the science-fiction movie Her (Jonze 2013), to begin the debate by exploring how an emotional or sentimental relationship, typically reserved for human-to-human interaction, could instead apply to the relationship between a human and a generative language model.

[5]All political groups at the EU Parliament were contacted by email and invited to participate.

2.2.2 Second edition programme

In the second edition, we developed a programme over three thematic days designed to immerse participants in the multifaceted realities of Generative AI, blending theoretical insights with practical engagement (Table 1B).[6]

The first day began with an interactive quiz that set a collaborative tone, followed by lectures addressing AI's impact on democracy (Cave and Dihal 2020), labour and power structures (Guarascio and Reljic 2025; Coveri et al. 2025), and military uses highlighting urgent regulatory needs on autonomous weapons (Stop Killer Robots 2024). A roundtable critically examined AI's environmental and social costs through diverse perspectives from labour market (Cetrulo and Nuvolari 2019), robotics (Blanchard 2023), and post-growth economics (Vrizzi and O'Neill 2025), moderated by a journalist (De Martino 2025). The day closed with a documentary screening on global AI power dynamics (DW 2024), which sparked informal reflections and peer dialogue.

The second day focused on the technical and ethical foundations of generative AI. Morning lectures explored how these systems function and their ethical challenges (Pappalardo et al. 2024), alongside the importance of standardisation to promote fairness and accountability (Williams 2022). Participants then tackled real-world AI safety issues (Grey and Segerie 2024) and engaged in a hands-on workshop led by PhD students invited from Scuola Normale Superiore's Responsible GenAI group (Muscato et al. 2025), fostering practical skills and collaborative learning. The day ended with a social dinner encouraging informal exchanges and networking.

The final day addressed legal and regulatory aspects, starting with an introduction to the EU AI Act and GDPR by an AI and privacy lawyer (Franco 2023), followed by insights into AI's challenges for copyright law (Sganga 2018). The afternoon featured a two-part workshop on AI compliance, combining theoretical discussions and participant-led lightning talks, allowing attendees to actively apply their knowledge.

2.3 Balance

The third design principle of the Ethos+Tekhne school concerned balance, in terms of space and time to interact (Section 2.3.1), providing not only quantity, but also quality, of interaction (Section 2.3.2). Moreover, we aimed to balance information provided to the participants with information gathered from them. Specifically, we allocated time for participants to process and explain what they had learnt (Section 2.3.3).

2.3.1 Space and time to interact

We limited participant group size for the on-site training to 30–40 individuals, in order to encourage close interaction, while streaming the daily programme for anyone to audit online, to promote open access for anyone interested. With respect to on-site participants, our goal was to provide each individual with enough space and time to actively engage in discussions. For this reason, the programme schedule aimed to balance activity and break time to allow participants to digest the new information and exchange thoughts in informal settings (Sections 2.2.1 and 2.2.2). We typically scheduled four activities per day, each lasting 75–90 minutes each. For lectures, this included 45–50 min presentations complemented by generous Q&A time. We allocated a 30-minute break between sessions, and a long lunch break (1.5 h) between morning and afternoon (Fig. 1B).

2.3.2 Quality of interactions

The initiative's positive impact was reinforced by careful attention to the quality of relationships among participants and instructors, the creation of an open and non-judgmental atmosphere that fostered collaboration and learning, and the deliberate choice to maintain a relatively small number of participants. These settings, combined with social events and sustained interaction over several days, allowed meaningful connections to form, and encouraged the development of a vibrant, supportive learning community.

2.3.3 Group writing to consolidate knowledge

After receiving new information, an important step to transform it into coherent and structured knowledge is writing (Rivard and Straw 2000). At the end of the school, participants were invited to write an extended abstract in groups of 4–5 people on one of the school's topics, to further exchange ideas, digest concepts and cement their knowledge. Groups were

[6]Programme of the second edition available here: https://community.eelisa.eu/activities/ethos-tekhne-a-new-generation-of-ai-researchers-2nd-edition/.

formed based on similar interests, creating the potential to connect them and let them leverage their text for a more formal paper later on. In the first edition, the group writing session took place in the afternoon of the last day, while in the second edition, participants were asked to submit their work during the week following the end of the school, in order to maximise the content offered to them despite the shorter duration of the event.

3. Results

Here, we present three types of results: i) the three main axes of actions that emerged from participant discussions (Section 3.1); ii) participant feedback concretising into a regenerative approach (Section 3.2); iii) the way training programmes in AI can connect Europe (Section 3.3).

3.1 Highlights: three axes of action

Overall, the lectures and discussions highlighted three main axes for ethical AI development: "AI for Good", "Good for AI", and "Good AI". "AI for Good" proposes AI development, applications and deployment that address social challenges (e.g. healthcare, climate change, etc.). "Good for AI" envisages socio-economic (rather than technological) innovation, to create a healthy environment for AI to be developed and applied (O'Neill et al. 2026; Edelman et al. 2025).[7] This axis is based on the idea that the socio-economic system and policies fundamentally shape their own components (here, the AI). Finally, "Good AI" refers to coding AI algorithms that comply with ethical and legal principles. An example is the ALTAI (Assessment List for Trustworthy AI) principles (EU High-Level Expert Group on AI, 2020). These principles are: 1) human agency and oversight, 2) technical robustness and safety, 3) privacy and data governance, 4) transparency, 5) diversity, non-discrimination and fairness, 6) environmental and societal well-being, and 7) accountability. Such algorithms leverage both software accessibility infrastructure and the development of computational tools to analyse AI algorithms' outputs and decision-making processes. Overall, these three axes of action involve experts across diverse skill domains — from more technical researchers (e.g. to develop "Good AI"), to researchers and professionals that can apply AI to specific real-world challenges ("AI for Good"), to researchers and professionals devising policies and system-level change ("Good for AI").

3.2 Participants' feedback

The training took place in March 2023 in Pisa, Italy, and in March 2025 in Paris, France. A post-school survey reported an overall mean satisfaction score of 4.2/5.0 (mean 4.4/5.0, median 4.0/5.0 for the first edition; mean and median 3.0/4.0 for the second edition) (Fig. 3A). The survey also collected suggestions for improvement. In the second edition, we substantially improved at least three of the top four areas for improvement (Fig. 3B), i.e. tutorials, socialising, and availability of teaching material, while identifying further areas for improvement.

3.2.1 Positive feedback

Participant feedback highlighted several strengths of the initiative. Many respondents expressed appreciation for the diverse, knowledgeable speakers and participants. The opportunity to easily engage with this diversity fostered rich discussions and fruitful networking opportunities. Many also valued the daily thematic structure, which offered clarity and a wide range of perspectives on AI. The format of the lectures — with manageable lengths, followed by interactive Q&A sessions — was praised for maintaining focus, encouraging participation, and enhancing understanding. The practical activities — such as group work, abstract writing, workshops, and notably the "moot court competition" session — were especially well received for being insightful and offering deeper engagement with the content. The balance between scientific content and social interaction was seen as a major strength: social moments during breaks were praised for encouraging informal learning and connection-building. Moreover, participants particularly appreciated the welcoming atmosphere, with friendly and approachable organisers. Overall, participants found the event enriching, eye-opening, and well-structured, and they recommended maintaining its interdisciplinary and interactive approach.

3.2.2 Areas of improvement

Participants highlighted three main areas for improvement: i) the intensity of the schedule, despite the long breaks, with suggestions to shorten the number of days or reduce the number of sessions; ii) poor time keeping during the event caused delays and decreased audience attention toward the end of each day — an aspect mainly due to the engaging nature of the talks and the spontaneous discussions that often extended beyond their allotted time; and iii) the late admission confirmations, making travel planning difficult. With respect to the first edition, the second edition improved on: i) some logistical aspects, concerning venue (electric outlets, room temperature, seating) and online connection (affecting

[7]Jonnie Penn provocatively contrasts technological and socio-economic innovation: "*What's the better invention: the Internet or the weekend*?" (Cambridge Festival 2024, minute 26:30).

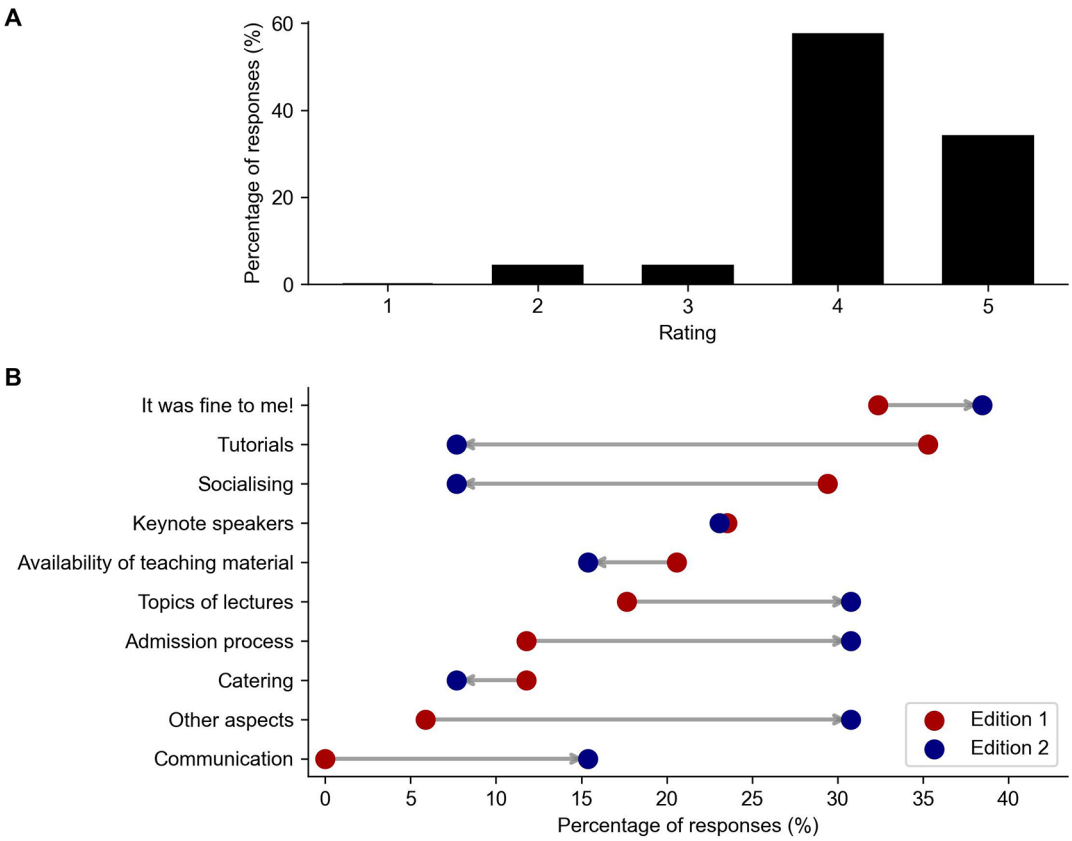

Figure 3. Overall satisfaction scores (A) and points-of-improvement votes (B) by responding participants. We collected feedback from on-site participants 2–4 weeks after the spring school. The survey included a mix of multiple-choice and open answers questions. In the first edition, the survey totalled 11 questions; in the second, 18. In the first edition, 34 participants out of 38 replied; in the second, 13 out of 23 (this total number of participants excludes the 8 members from the organising team). Panel **A** shows the aggregate answers from the two editions to the question *"How would you rate the Spring School?"*. In the first edition, possible ratings ranged from 1 to 5 ("*A waste of time*"=1, "*Excellent*" = 5); in the second, from 1 to 4 ("*Poor*" = 1, "*Adequate*" = 2, "*Good*" = 3, "*Excellent*" = 4). We rescaled values from the second edition to a 1–5 scale. Panel **B** shows the aggregate answers from the two editions to the question "*What should we improve*?". Every time a respondent ticked an option, we counted it as a "*Yes*"; respondents were allowed to select multiple options. The percentage is calculated with respect to the number of respondents for each edition. For this question ("*What should we improve*?"), in the survey for the second edition, we added the option "Communication" in the survey; we set its score to zero for the first edition.

remote speakers); ii) a more accessible level of some technical talks in AI for participants from different backgrounds; iii) longer and more focused time for abstract writing (Section 2.3.3); and iv) the addition of social events in the evening, beyond the movie night. In the future, the school also aims to evaluate its impact, adding an ex-ante evaluation of the participants, to compare it to their ex-post knowledge.

3.2.3 Regenerative nature

All respondents in both editions answered "*Yes*" to "*Do you think there should be a second/third edition?*". While this result may suggest a bias in the collected responses (i.e. only participants who highly enjoyed the school responded), it is a promising achievement underscoring the support and enjoyment of many participants. As evidence, these responses also turned into concrete action: some first-edition participants formed the core of the organising team for the second edition, and in turn, some second-edition and first-edition participants are currently organising the third edition. This regenerative nature is the most tangible result of the engagement and enjoyment of participants with the current design principles of the Ethos+Tekhne training programme.

3.3 Connecting Europe

Additionally, training initiatives are also opportunities to connect Europe. First, this initiative served as a first-time inter-institutional collaboration between three institutions,[8] showing how new initiatives can create new connections across European institutions. Moreover, the new team, who will soon be running the third edition,[9] doubled the number of EELISA partner institutions[10] and joined forces with the United Nations Institute for Training and Research (UNITAR).[11] Second, the European Court of Auditors audited the school online to evaluate the effectiveness of the EU Commission's investments in AI research. This aspect suggests that the impact of training programmes in AI can go beyond learning for participants, benefitting EU institutions as well.

From these promising results, in the long-term, the school aims to create a network of European universities that train young AI researchers in understanding and discussing the European regulatory framework on AI and thus building a responsible AI infrastructure.

4. Conclusions

In line with the latest recommendations in the field of AI (Srikumar et al. 2022),[12] we devised a novel multidisciplinary training for the upcoming generation of AI researchers to align their technical skills with societal impact. We identify three main successful ingredients of our programme: 1) the strong diversification of speakers and learning activities, 2) the creation of a programme-wide, question-driven narrative over thematic days, and 3) the allocation of time and space for participants to digest information, engage and interact with speakers and with each other. Moreover, two further ingredients strengthened our initiative: its bottom-up nature, as it was organised by and for young researchers and professionals, and the creation of cross-country inter-institutional partnerships to organise the school. We believe that all these elements jointly offer solid bases to help prepare the upcoming generation of young researchers and professionals in AI to uphold the responsibility of building a global infrastructure.

Data Availability

The dataset supporting the results is openly available under a CC-BY licence at the following DOI: https://doi.org/10.5281/zenodo.18986243 (Vrizzi et al. 2026).

Acknowledgements

The organising team of the "Ethos+Tekhne: a new generation of AI researchers" EELISA spring school included, in the first edition: Daniele Cavalli, Konstantinos Parginos, Giuseppe Parisi, Raffaele Salvia, Viola Stefanello, and Letizia Lamperti, together with the co-authors Filippo Galli, Laura State, Alberto Arturo Vergani, and Stefano Vrizzi (coordinator); in the second edition: Antoine Bordas, Guido Cassinadri (coordinator), Lorenzo Fontanelli, Jay Luong, Clara Punzi, Stefano Tramacere, together with the co-authors Sara Mazzucato and Amric Trudel (coordinator). Currently, the third edition team includes Sofia Bell, Noemi Luna Carmeno, Mahshid Cheraghi, Marybelle Cherfan (coordinator), Dumitru Denis, Ivan Fino, Ismini Mathioudaki, Gabriele Nanino, Janice Sairus, Maritza Salcido, Dániel Sándor, Barnabás Suciu, Simone Vagnoni. Moreover, the realisation of the school has only been possible thanks to the crucial work of the administration and support team, and the Scientific Committee. The administration and support team included Rossella Raso, Camille Roger, Elisa Guidi, Jean-Baptiste Astruc, Sabrina Bigazzi, Chiara Busnelli, Elena Cambi, Silvia Gargini, Francesco Forino, as well as Viviane Lunardi, who helped manage and collected the feedback of the first edition. We are particularly grateful to Camille Roger and Rossella Raso, who were the key points of reference throughout the two editions. The Scientific Committee was composed of Calogero Oddo, Sara Moccia, Lorenzo Bartalesi, Roberto Pellungrini and Florian Yger in the first edition, Boris Gutkin, Sara Moccia, Calogero Oddo in the second edition, with a special thanks to Calogero Oddo and Sara Moccia for believing in and helping shape this initiative since its early stages. More broadly, we would like to thank the organising institutions Paris Sciences et Lettres, Scuola Normale Superiore, and Scuola Superiore Sant'Anna for hosting the classes and providing facilities. Additionally, we would like to thank Diane Gabeloteau for suggesting to us the inter-institutional funding opportunities from EELISA. Furthermore, we would like to thank Antonio Capellupo at Cinema Arsenale for welcoming us and leading us in the process of organising a movie

[8]Paris Sciences et Lettres, Scuola Normale Superiore, and Scuola Superiore Sant'Anna. Members from other institutions also joined the organising team.
[9]Programme of the third edition available here: https://digitalcampus.eelisa.eu/activities/88638c15-0642-4e84-8bcc-ac4652051b62.
[10]Scuola Superiore Sant'Anna, Paris Sciences et Lettres, Friedrich-Alexander-Universität Erlangen-Nürnberg, Universidad Politécnica de Madrid, Budapesti Műszaki és Gazdaságtudományi Egyetem, and Universitatea Națională de Stiință si Tehnologie Politehnica Bucuresti. Members also joined the team from institutions outside the EELISA network, such as the University of Barcelona and the University of Edinburgh.
[11]The third edition will employ "Everything AI" as theme:https://unitar.org/about/news-stories/featuredarticles/unitar-and-eelisa-ethostekhne-announce-everything-ai-international-event-ai-governance-brussels.
[12]See especially the "Create venues for debate, alignment and collective action" section in (Srikumar et al. 2022).

discussion night in Pisa with an international audience during the first edition. Ethos+Tekhne was organised as part of the Ethics, Social Commitment and Entrepreneurship (ESCE) EELISA community.

We would also like to thank Daniel Walton O'Neill for his feedback on the manuscript, and Mireia Codina Tobias for her feedback on the figures. In particular, Stefano Vrizzi would like to thank Laura State for early discussions on the idea and design of an interdisciplinary summer school on AI.

Finally, we would like to thank all the participants, for their active contributions to the school and for continuing this initiative, and all the invited speakers, for offering their time and sharing their knowledge.

References

Balayn A, Rikalo N, Lofi C, *et al.*: **How can Explainability Methods be Used to Support Bug Identification in Computer Vision Models?** *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems.* 2022; 1–16.
Publisher Full Text

Bellamy RK, Dey K, Hind M, *et al.*: **AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias.** *IBM J. Res. Dev.* 2019; **63**(4/5): 4:1–4:15.
Publisher Full Text

Blanchard A: **Matériels, usages et société: éco-conception multidisciplinaire d'objets.** *Drôles d'objets 2023.* 2023.
Reference Source

Cacciatore A, Migliorelli L, Berardini D, *et al.*: **Some Ethical Remarks on Deep Learning-Based Movements Monitoring for Preterm Infants: Green AI or Red AI?** *Image Analysis and Processing. ICIAP 2022 Workshops: ICIAP International Workshops, Lecce, Italy, May 23–27, 2022, Revised Selected Papers, Part II. N.p.* Springer International Publishing; 2022. Print ISBN: 978-3-031-13323-7; Online ISBN: 978-3-031-13324-4.
Publisher Full Text

Cambridge Festival: **How will AI affect the democratic process?** 2024. dir.
Reference Source

Cave S, Dihal K: **The Whiteness of AI.** *Philos Technol.* 2020; **33**(August): 685–703.
Publisher Full Text

Cetrulo A, Nuvolari A: **Industry 4.0: revolution or hype? Reassessing recent technological trends and their impact on labour.** *J. Ind. Bus. Econ.* 2019; **46**(3): 391–402.
Publisher Full Text

Comandé G, editor: *Elgar encyclopedia of law and data science.* Edward Elgar Publishing; 2022.
Publisher Full Text

Coveri A, Cozza C, Guarascio D: **Monopoly capitalism in the digital era.** *LEM Working Paper Series.* 2021.
Reference Source

Coveri A, Cozza C, Guarascio D: **Blurring boundaries: an analysis of the digital platforms-military nexus.** *Review of Political Economy.* 2025; **37**(4): 1632–1663.
Publisher Full Text

De Martino C: **Monétisation des contenus et IA: l'IAB Tech Lab veut accompagner les éditeurs.** *Mind Media.* 2025.
Reference Source

de Miguel I, Sanz B, Lazcoz G: **Machine learning in the EU health care context: exploring the ethical, legal and social issues.** *Inf. Commun. Soc.* 2020; **23**(8): 1139–1153.
Publisher Full Text

DW, dir: *AI Supremacy: The Artificial Intelligence Battle Between China, USA and Europe.* DW; 2024.

Edelman J, Zhi-Xuan T, Lowe R, *et al.*: **Full-Stack Alignment: Co-Aligning AI and Institutions with Thick Models of Value.** *arXiv preprint arXiv: 2512.03399.* 2025.

EU High-Level Expert Group on AI: **Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment.** *Shaping Europe's digital future.* 2020.
Reference Source

Franco G: **5 cose da sapere sulla nuova legge europea sull'intelligenza artificiale.** *La Repubblica.* 2023.

Giannotti F, Naretto F, Bodria F: **Explainable for Trustworthy AI.** *Human-Centered Artificial Intelligence: Advanced Lectures.* Springer International Publishing; 2023.
Publisher Full Text

Grey M, Segerie C-R: *AI Safety Atlas.* French Center for AI Safety (CeSIA); 2024.
Reference Source

Guarascio D, Reljic J: **AI and Employment in Europe.** *Econ. Lett.* 2025; **247**(0165–1765): 112–183.
Publisher Full Text

Henry B: **Robotic and Cybernetic Plots in the Global Imaginary, on "Iride, Filosofia e discussione pubblica"** 2/2025; pp. 311–334.
Publisher Full Text

High-Level Expert Group on Artificial Intelligence: **Ethics Guidelines for Trustworthy AI.** 2019.
Reference Source

Huyskes D: *Constructing automated societies. Socio-cultural determinants and impacts of automated decision-making in public services.* Magazzini Digitali dedicato alle Tesi di dottorato; 2025. Doctoral dissertation.
Reference Source

Jonze S: *Her.* Annapurna Productions; 2013. dir.

Merenda F: **Reading Arendt to rethink truth, science, and politics in the era of fake news.** *In Democracy and Fake News.* Routledge; 2020; pp. 19–29.
Publisher Full Text

Murray N: **Neural network for generating images trained with a generative adversarial network.** U.S. Patent Application 17/024,030, filed March 25, 2021.

Muscato B, Passaro L, Gezici G, *et al.*: **Perspectives in Play: A Multi-Perspective Approach for More Inclusive NLP Systems.** *arXiv.* 2025.
Publisher Full Text

O'Neill DW, Vrizzi S, Carmeno NL, *et al.*: **The economic alignment problem of artificial intelligence.** *arXiv preprint arXiv:2602.21843.* 2026.

Norori N, Hu Q, Aellen FM, *et al.*: **Addressing bias in big data and AI for health care: A call for open science.** *Patterns.* 2021; **2**.
Publisher Full Text

Pappalardo L, Citraro S, Cornacchia G, *et al.*: **A survey on the impacts of recommender systems on users, items, and human-AI ecosystems.** *arXiv.* 2024.
Publisher Full Text

Parginos K, Bessa R, Camal S, *et al.*: **Interpretable data-driven solar power plant trading strategies.** *2022 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe).* 2022; 1–5.
Publisher Full Text

Pianta M: **Technology and work: key stylized facts for the digital age.** *Handbook of Labor, Human Resources and Population Economics.* Springer, Cham; 2020; **1**(17): 1–17. Print ISBN: 978-3-319-57365-6; Online ISBN: 978-3-319-57365-6.
Publisher Full Text

Rivard LP, Straw SB: **The effect of talk and writing on learning science: An exploratory study.** *Sci. Educ.* 2000; **84**(5): 566–593.
Publisher Full Text

Samek W, Montavon G, Lapuschkin S, *et al.*: **Explaining deep neural networks and beyond: A review of methods and applications.** *Proc. IEEE.* 2021; **109**(3): 247–278.
Publisher Full Text

Seismic Foundation: **On The Razor's Edge: AI vs Everything We Care About.** *Seismic Report 2025.* 2025.
Reference Source

Sganga C: *Propertizing European Copyright: History, Challenges and Opportunities.* Edward Elgar Publishing Limited; 2018.
Reference Source

Sinek S: *Start with Why: How Great Leaders Inspire Everyone to Take Action.* Penguin Publishing Group; 2009.

Srikumar M, Finlay R, Abuhamad G, *et al.*: **Advancing ethics review practices in AI research.** *Nature Machine Intelligence.* 2022; **4**(12): 1061–1064.
Publisher Full Text

Stop Killer Robots: **Submission on Autonomous Weapon Systems to the United Nations Secretary-General.** 2024.
Reference Source | **Reference Source**

The European Parliament and The Council of the European Union. **Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).** 2024.
Reference Source

Vrizzi S, O'Neill DW: **Modelling the Doughnut of social and planetary boundaries with frugal machine learning.** *EurIPS 2025 "Rethinking AI" Workshop.* 2025.
Publisher Full Text

Vrizzi S, Vergani AA, State L, *et al.*: **Ethos+Tekhne Feedback surveys.** 2026.
Publisher Full Text

Williams R: **Using iterative pairwise external validation to contextualize prediction model performance: A use case predicting 1-year heart failure risk in patients with diabetes across five data sources.** *Drug Saf.* 2022; **45**(17): 563–570.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Zhou K, Yang Y, Cavallaro A, *et al.*: **Omni-scale feature learning for person re-identification.** *In Proceedings of the IEEE/CVF international conference on computer vision.* 2019; pp. 3702–3712.

Open Peer Review

Current Peer Review Status: ? ? ? ?

Version 1

Reviewer Report 11 August 2026

https://doi.org/10.21956/openreseurope.21885.r77477

© **2026 Por F.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Fei Ping Por
Tunku Abdul Rahman University of Management and Technology, Kuala Lumpur, Federal Territory of Kuala Lumpur, Malaysia

In the transition between describing the first and second editions (between Sections 2.2.1 and 2.2.2), explicitly explain the rationale for the key design changes — particularly the reduction from five to three days, the shift in disciplinary targeting, and the modified survey instrument. State whether these were responses to participant feedback, resource constraints, or deliberate pedagogical redesign.

Clearly describe the analytical method by which the three axes of action were identified. If they were derived through a structured process (e.g., thematic analysis of discussion notes or written abstracts), describe that process. If they represent the organisers' own conceptual synthesis, present them as part of the programme design framework rather than as empirical findings.

Add a subsection to the Introduction (or expand Section 1.1) that reviews existing AI ethics training initiatives and identifies their limitations, clearly articulating what gap Ethos+Tekhne fills. A comparison table summarising features of two to three comparable programmes would be valuable.

The claim that the second edition "substantially improved" on three of the top four areas for improvement is based on comparing percentages of respondents who flagged each area across editions. But these percentages come from different sample sizes (34 respondents versus 13), different survey instruments (11 questions versus 18, with at least one new option added), and potentially different respondent populations (given the 57% response rate in the second edition). Any shift in percentages could easily reflect sampling variation, survey design changes, or response bias rather than genuine improvement.

Is the background of the case's history and progression described in sufficient detail?
Partly

Is the work clearly and accurately presented and does it cite the current literature?

Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Partly

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Yes

Is the case presented with sufficient detail to be useful for teaching or other practitioners?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Educational Management, Educational Technology

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Reviewer Report 11 August 2026

https://doi.org/10.21956/openreseurope.21885.r77478

© 2026 Sison L. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Louis Robert C. Sison
Bulacan State University, Malolos, Central Luzon, Philippines

The manuscript presents a timely, relevant, and potentially valuable case study of interdisciplinary training at the intersection of artificial intelligence, ethics, law, and society. Its particular strengths are the transparency of the programme description, the operational detail provided for practitioners, the interdisciplinary composition of speakers and participants, and the openly available supporting dataset. The manuscript's discussion of the three design principles, pluralism, narrative, and balance, also offers a useful organizing vocabulary for future programme development.

The principal limitation is that the manuscript currently functions more convincingly as a well-documented programme case report than as an evaluation of educational effectiveness. The evidence primarily concerns participant feedback, programme implementation, and continuation of the initiative. These data are valuable, but they do not yet demonstrate that participation improved ethical reasoning, AI competencies, responsible decision-making, or professional

practices. The distinction between programme acceptance/satisfaction and programme effectiveness should therefore be maintained consistently throughout the manuscript.

A second major concern relates to the theoretical contribution. While the programme is innovative in practice, the justification for pluralism, narrative, and balance is not sufficiently anchored in established educational scholarship. Connecting these principles with contemporary theories of adult learning, experiential learning, interdisciplinary learning, social constructivism, communities of practice, or instructional design would substantially strengthen the manuscript. The authors should also compare Ethos+Tekhne more systematically with existing AI-ethics education initiatives to establish the empirical and pedagogical gap that the case addresses.

Finally, the manuscript would benefit from a clearer evaluation framework. Future iterations should ideally employ pre-/post-measures of knowledge, ethical reasoning, attitudes, or competency, complemented by qualitative evidence of learning and, where possible, longitudinal follow-up. Such evidence would permit the authors to move beyond demonstrating that participants valued the initiative toward establishing whether the programme produces meaningful educational outcomes.

Indeed, this paper has publication value, specifically its transferability.

Is the background of the case's history and progression described in sufficient detail?
Partly

Is the work clearly and accurately presented and does it cite the current literature?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Partly

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Partly

Is the case presented with sufficient detail to be useful for teaching or other practitioners?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Mathematics Education, Pre-service Teaching Education

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Reviewer Report 06 August 2026

https://doi.org/10.21956/openreseurope.21885.r77472

© 2026 Rajiansyah R. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Rajiansyah Rajiansyah
Universitas Mulawarman, Samarinda, East Kalimantan, Indonesia

The manuscript presents a timely and relevant initiative in AI ethics education, and the proposed program is clear and well-organized. However, I identified several areas where the manuscript could be strengthened to enhance its scientific contribution.

First, specify the empirical research gap more clearly. Although the manuscript discusses the growing importance of AI ethics education, it does not critically compare existing AI ethics training programs or clearly explain how the proposed "Ethos+Tekhne" initiative differs from or improves upon previous work. Provide a more comprehensive review of related initiatives and their limitations to better justify the study's novelty.

Second, specify the program design more explicitly in relation to established educational theories. The program is described primarily from a practical perspective, with limited integration of educational theory. Incorporate relevant frameworks, such as Knowles' Adult Learning Theory, Kolb's Experiential Learning Theory, or other instructional design models, to provide a stronger theoretical foundation for the learning activities and pedagogical decisions presented in the manuscript.

Finally, specify why the Diversity, Narrative, and Balance principles were selected as the core design framework. While these principles are presented as central, the manuscript does not adequately explain why they were chosen over other well-established educational approaches. Provide a brief comparison with alternative frameworks, such as project-based learning, problem-based learning, or the flipped classroom model, together with a clear justification for the chosen principles, to strengthen the rationale for the program design.

Address these points by specifying the revisions more clearly to improve the manuscript's clarity, theoretical rigor, and scholarly contribution, and by providing readers with a stronger understanding of the novelty and educational value of the proposed initiative.

Is the background of the case's history and progression described in sufficient detail?
Partly

Is the work clearly and accurately presented and does it cite the current literature?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?

Yes

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Partly

Is the case presented with sufficient detail to be useful for teaching or other practitioners?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Decision Theory, Game Theory, Gamification, AI, Programming

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Reviewer Report 28 July 2026

https://doi.org/10.21956/openreseurope.21885.r77476

© 2026 Daryanes F. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Febblina Daryanes
Universitas Riau, Pekanbaru, Riau, Indonesia

1. The introduction provides a clear rationale for the growing need for multidisciplinary AI education. However, the empirical research gap is not yet sufficiently articulated. While the authors highlight the importance of responsible AI training, they do not adequately synthesize previous studies to demonstrate what has already been achieved, what limitations remain, and how the proposed programme addresses these unresolved issues. Strengthening the empirical gap by comparing existing AI training initiatives and explicitly identifying their limitations would better justify the novelty and scientific contribution of this work.
2. The manuscript is generally well organized and clearly presented, and it cites recent and relevant literature. However, several aspects could be strengthened to enhance its scholarly contribution. First, the programme is described primarily from a practical perspective, with limited integration of established educational theories, such as adult learning theory, experiential learning, or instructional design frameworks. Providing a stronger theoretical foundation would better support the rationale for the programme design. Second, the manuscript does not sufficiently explain why the Diversity–Narrative–Balance principles were selected over other existing educational or training frameworks. Clarifying the theoretical and practical justification for these design principles would strengthen the

conceptual contribution of the study.
3. This article is not quantitative research.

Is the background of the case's history and progression described in sufficient detail?
Partly

Is the work clearly and accurately presented and does it cite the current literature?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Yes

Is the case presented with sufficient detail to be useful for teaching or other practitioners?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Biology Education, Science Education

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 14 Oct 2026

Stefano Vrizzi

1. The introduction provides a clear rationale for the growing need for multidisciplinary AI education. However, the empirical research gap is not yet sufficiently articulated. While the authors highlight the importance of responsible AI training, they do not adequately synthesize previous studies to demonstrate what has already been achieved, what limitations remain, and how the proposed programme addresses these unresolved issues. Strengthening the empirical gap by

comparing existing AI training initiatives and explicitly identifying their limitations would better justify the novelty and scientific contribution of this work.
We thank the reviewer for raising this point. We added one whole section called 'Literature review' (Section 1.2) within our introduction. We discuss relevant schools on AI, ethics, and society, their similarities and differences with Ethos+Tekhne. We highlight the gap that our school and our report aim to fill.

2. The manuscript is generally well organized and clearly presented, and it cites recent and relevant literature. However, several aspects could be strengthened to enhance its scholarly contribution. First, the programme is described primarily from a practical perspective, with limited integration of established educational theories, such as adult learning theory, experiential learning, or instructional design frameworks. Providing a stronger theoretical foundation would better support the rationale for the programme design. Second, the manuscript does not sufficiently explain why the Diversity–Narrative–Balance principles were selected over other existing educational or training frameworks. Clarifying the theoretical and practical justification for these design principles would strengthen the conceptual contribution of the study.
We thank the reviewer for raising this point. We added one whole section called 'Rationale of design principles' (Section 2.4) within our Methods. We present the reasons why we chose those design principles, and we discuss them in relation to established educational theories, such as adult learning theory, experiential learning and more.

Competing Interests: No competing interests were disclosed.