

Youth Perspectives on AI Ethics: An Exploratory Study within the VERFISUM Project

Paulo Duarte Branco¹, Andreia Teles Vieira¹, Maria Potes Barbas²

¹ Instituto Superior de Tecnologias Avançadas de Lisboa,
Centro de Investigação em Artes e Comunicação, Pólo de Literacia Digital e Inclusão Social
paulo.duarte@my.istec.pt, andreia.vieira@my.istec.pt

² Instituto Politécnico de Santarém, Escola Superior de Gestão e Tecnologia de Santarém
Centro de Investigação em Artes e Comunicação, Pólo de Literacia Digital e Inclusão Social
mariapbarbas@gmail.com

***I. Abstract—**As artificial intelligence (AI) becomes increasingly embedded in everyday life, understanding how young people perceive its ethical implications is both urgent and underexplored. This study, conducted within the Erasmus+ project VERFISUM, explores ethical perceptions of AI usage among 91 participants from five European countries (Portugal, Greece, Estonia, Ukraine, and Italy), including both students and educators. The study employed an exploratory online survey with 30 closed-ended questions addressing themes of ethics, responsibility, trust, and regulation.*

Findings reveal a generation that is both digitally immersed and ethically involved. Participants expressed support for ethical principles beyond legal requirements, endorsed the regulation of AI usage, and raised concerns about misinformation, environmental sustainability, and job displacement. On the other hand, results also emphasize worries about responsibility in AI-driven decisions and uncertainty around AI's role in education. The discussion situates these findings within current ethical and educational frameworks, raising awareness for an increased critical digital literacy, participatory policymaking, and ethics-integrated curricula. Finally, the study highlights the importance of treating young people not just as users of AI, but as ethical stakeholders in its development and governance.

II. INTRODUCTION

In the last few years, artificial intelligence (AI) has moved from action movies and science fiction stories into reality and everyday life. From language translation tools to recommendation systems on social media and related websites, from virtual assistants on smartphones to AI-driven decision-making in public services, from customer support lines to self-driving cars, intelligent technologies are increasingly transforming the way we work, learn, and relate to one another. As these systems become more embedded in society, questions regarding their ethical design, distribution, and usage have become impossible to ignore.

At the core of this concern lies a simple but unrelenting dilemma: *just because we can, should we?* This interrogation is especially relevant for young people, who often engage with

digital technologies more intensively than any other age group — and yet may not always have the critical tools to reflect on the systems they use or the values they conceal.

This study is part of the **VERFISUM project**, funded by the Erasmus+ programme, which aims to support the employability of youth and young adults, particularly those with fewer opportunities. But as we are designing the educational tools, we found ourselves asking:

How do young people actually feel about AI?

Do they trust it?

Do they understand it?

And most importantly — do they think it should be ethical?

These questions led us to conduct exploratory research through a survey, not only to measure factual knowledge but also to uncover perceptions, doubts, and hopes about the relationship between ethics and AI. Contemporary debates, including the Asilomar Principles (2017), the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), and the newly adopted European AI Act (2024), provide the conceptual and regulatory framework supporting this research and highlight the ongoing need for critical discussion on AI ethics. Understanding youth perceptions is not only valuable but essential for developing technologies that reflect democratic values and support more inclusive futures.

III. THEORETICAL BACKGROUND

A. Ethics: From Ancient Roots to Technological Challenges

As a philosophical discipline, ethics has long attempted to provide fundamental answers on the proper way for humans to live — both individually and collectively. Virtue, obligation, and reasoned thought were highlighted as means of achieving a fair and fulfilling existence in classical traditions, particularly those of Aristotle and later Immanuel Kant [1], [2]. Despite being

centuries old, these frameworks still provide conceptual support for discussions and debates on technologies and related contemporary issues.

Nowadays, ethical considerations extend beyond human interactions to include **non-human actors** and intelligent computers. This is particularly true in the digital age. Researchers have posed the pressing topic *How do we assign moral responsibility in systems where no individual is fully in control?* As artificial intelligence (AI) systems begin to automate social tasks, influence decision-making, and even mimic human behaviours [3], [4].

Importantly, ethics is now seen as a **dynamic, context-sensitive process** of group reflection rather than a fixed norm or static code [5]. Because of its adaptability, it is especially pertinent to discussions of evolving technologies, which continually challenge the old divisions between tool and agent, intention and algorithm, and harm and malfunction.

B. Artificial Intelligence and the Ethics of Intelligent Systems

One issue with artificial intelligence is that it is not neutral. Its algorithms reflect the biases, intentions, and limitations of those who design, train, and deploy them [6], [7]; as well as the data used for their training. Under the pretext of technological neutrality, AI technologies—from predictive policing to automated hiring systems—are increasingly intertwined with social and political processes, frequently reproducing preexisting inequities [8]. An example of this are the large language models that are being under debate with their biases according to the developer, some of them build in the “occident” like ChatGPT and Gemini (from USA companies like OpenAI and Google) where others were developed in the “orient” like DeepSeek (from the Chinese company with the same name).

AI's ethical field has responded to these challenges by addressing critical areas of interest:

- **Fairness and bias** [9],
- **Transparency and explainability** [10],
- **Accountability and governance** [11],
- **Privacy and data protection** [12].

Various frameworks have been proposed to guide the ethical development and use of AI. The Asilomar AI Principles [13] presented high-end indicators emphasizing safety, transparency, and long-term value alignment. The **UNESCO Recommendation on the Ethics of Artificial Intelligence** [14] marked an essential step in creating a globally endorsed, rights-based framework. Recently, the **European Union's AI Act** introduced the first binding legal framework for classifying AI systems according to risk and imposing appropriate obligations [15].

Despite the growing number of guidelines, it's still common to criticise that such principles often lack mechanisms for *implementation, enforcement, or public accountability* [16]. This “principle-to-practice” gap underscores the need for *ethical*

literacy not only among developers and policymakers but also among citizens—especially younger generations.

C. AI and Ethics in Education: Why Youth Voices Matter

The incorporation of AI into education is not merely a matter of adopting new tools or upgrading infrastructure. It goes much deeper — it's an ethical venture at its core. Whether we're discussing intelligent tutoring systems, predictive analytics, AI-generated feedback, or automated grading, these technologies inevitably invite us to confront challenging questions about *student agency, fairness*, and the very essence of learning [17], [18].

In recent years, there's been a notable shift in the literature from basic digital literacy (being able to operate digital tools) to what's often called *critical digital literacy*. This involves something more demanding: the capacity to interrogate the technologies we use, to ask who built them, for what purpose, and with what unintended consequences [19], [20]. It's about understanding the political economies, cultural assumptions, and systems of power that are embedded in seemingly neutral code.

Today, students are often referred to as “digital natives”. While that term captures something about their fluency, it risks obscuring an equally important point: fluency does not always imply awareness. Familiarity with AI tools doesn't guarantee the ability to think critically about them. Many young people still lack the ethical language — or the conceptual tools — to unpack how these technologies influence privacy, creative expression, identities, and yes, even their democratic rights [21]. That gap concerns me—especially when tools like ChatGPT, Midjourney, or AI-based surveillance are not on the horizon but are already woven into educational spaces and working lives.

Projects such as **VERFISUM** aim to bridge this divide. By integrating ethical inquiry into vocational and digital learning, they offer more than technical upskilling — they create spaces for reflection, dialogue, and value formation. And in doing so, they echo broader movements urging us to place ethics at the heart of education in the 21st century [22], [23]. Because ultimately, our goal shouldn't be to prepare young people to *use* technology — it should be to equip them to *shape* it, critically and collectively, in ways that honour human dignity and shared responsibility.

IV. METHODOLOGY

A. Study Design

This study was envisioned as exploratory. Rather than aiming to test predefined hypotheses or construct statistically generalizable claims, the intention was to listen — to gather honest, unfiltered perspectives from young people and those who work closely with them. The central goal was to map current attitudes and basic understandings regarding ethics and artificial intelligence, and to surface areas of confusion, curiosity, and potential contradictions.

B. Instrument

To gather these insights, a 30-item online questionnaire was developed, carefully structured to create space for reflection rather than constrain respondents within overly rigid categories. Most questions were closed-ended, with three options — “Yes,” “No,” and “Not sure” — allowing for the detection of uncertainty as much as conviction.

The questionnaire was divided into thematic blocks, each intended to probe different ethical dimensions of AI.

The design drew inspiration from several foundational frameworks that have shaped contemporary discourse on AI ethics, including the Asilomar AI Principles (2017), the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), and the recently enacted European AI Act (2024). These were not simply theoretical anchors, but living documents that continue to inform policy, research, and pedagogy across Europe and beyond.

Importantly, this questionnaire was not subjected to formal validation. This was a deliberate decision: the instrument was intended to function more as a conversation starter than as a rigid measurement tool. In that spirit, it served to collect firsthand impressions, provoke internal and group reflection among participants, and generate preliminary patterns to inform future deeper studies.

C. Data Collection and Participants

The survey was conducted online during **June 2024**, with support from partner organizations across five countries involved in the **VERFISUM project: Portugal, Greece, Estonia, Ukraine, and Italy**.

In total, **91 responses were received**. The age range was broad — from 14 to 60 years — reflecting the intergenerational nature of the discussion. Participants were distributed into two key groups:

- **28 identified as educators or youth workers**, individuals often tasked with interpreting technological changes for others,
- **63 identified as students or young people** in secondary or higher education, or vocational training.

Although our sample does not support rigorous cross-country comparative analysis, the geographic spread enabled us to capture a variety of cultural entry points into the conversation about AI and ethics. What emerged was not a unified voice but a **mosaic of perspectives** — one believed to be particularly valuable at this early stage of ethical inquiry in the age of AI.

D. Limitations

As with any exploratory work, there are important limitations to note. The sample, while diverse, is relatively small and not randomly selected, so caution is warranted when generalizing these findings. The instrument itself — intentionally unvalidated — prioritizes openness and breadth over statistical robustness and thus cannot claim psychometric rigor. Moreover, the questionnaire was distributed in English, which

may have introduced interpretive variability depending on participants’ language proficiency and comfort. That said, these limitations are not flaws so much as **contextual conditions**. This was a study designed to illuminate, not to conclude. What was gained was a valuable snapshot — a moment of ethical consciousness among young Europeans and their mentors. These insights now provide a meaningful foundation for designing more targeted research and, perhaps more importantly, for shaping educational tools and spaces within **VERFISUM** that reflect real ethical concerns rather than abstract assumptions.

V. RESULTS AND ANALYSIS

A. Awareness and Familiarity with AI

The results confirm that artificial intelligence is no longer a distant or abstract concept for most participants. A large majority (88%) reported awareness of AI technologies before 2021, indicating that AI has been present in their lives for several years. Direct engagement with AI tools was also widespread: approximately 63% of respondents reported using ChatGPT for professional purposes, with similar usage for personal or entertainment activities.

This level of engagement suggests that participants have moved beyond initial curiosity toward an everyday form of coexistence with AI systems. However, frequent use does not necessarily imply deep understanding. As the results indicate, familiarity with AI tools does not always correspond to high levels of critical or ethical literacy, a tension that emerges across several dimensions of the data.

B. Ethical Intuitions: Stable Values and Evolving Frameworks

Participants demonstrated strong and consistent ethical intuitions. An overwhelming majority (92.3%) agreed that ethical norms should extend beyond legal requirements, while 80.2% affirmed the existence of shared or universal moral principles. These findings indicate a clear normative orientation and a belief that legality alone is insufficient to address ethical challenges posed by AI.

At the same time, 73.3% rejected the idea that ethical principles remain unchanged over time. Rather than indicating inconsistency, these responses suggest that participants perceive ethics as grounded in shared values while still adaptable to changing social and technological contexts. This reflects an ethical perspective that balances continuity with contextual change.

C. Machines and Moral Responsibility

One interesting aspect of this survey was how participants responded to questions about **extending ethical concern to machines**. **60% believed that ethics should encompass machines**, and an even larger majority — **nearly 76%** — believed intelligent agents deserve special ethical consideration.

Such responses echo ongoing philosophical debates: Can machines be treated as moral subjects, or are humans held responsible for how machines are made and used? Although the extent of individual engagement with this question cannot be determined, the responses suggest a **growing willingness to reconsider ethical agency amid technological evolution.**

Participants were also asked whether robots should be designed to behave ethically “by default.” Here, **67% said yes** — a figure that reflects both optimism and responsibility. It implies that ethics is not merely a human attribute but also something that can be embedded — or at least approximated — in code.

Perhaps one of the most revealing sections of the survey dealt with **accountability in AI-driven systems** — particularly autonomous cars. When asked who should be responsible for an AI system’s potentially harmful decisions, responses were scattered across programmers (50.5%), vendors (50.5%), and drivers (61.1%). A significant portion remained unsure.

This diffusion of answers is, in itself, telling. It suggests that participants perceive AI not as a tool wielded by a single actor, but as something **produced by many hands, governed by many layers, and capable of slipping through traditional liability models.** Philosophers like Coeckelbergh [4] have called this the “responsibility gap” — and clearly, that gap is felt at a public level, not just in academic discourse.

D. Risks, Trust, and Perceived Impacts

Participants expressed significant concern about the ethical risks associated with AI. A large majority (81.3%) believed that ethical problems extend beyond physical harm. The most frequently identified risks included privacy violations, intelligent weapons, fake news, job displacement, and environmental degradation. These concerns indicate awareness of concrete, ongoing societal challenges rather than speculative or purely futuristic threats.

At the same time, participants expressed optimism regarding AI’s potential benefits. While 84.3% believed that AI would improve work efficiency, 77.5% also anticipated significant job losses. This coexistence of optimism and anxiety highlights a central tension in contemporary AI ethics, reflecting both hope for technological progress and concern about its social consequences.

E. Regulation and Public Awareness

Support for regulation emerged as one of the clearest findings. A substantial majority of respondents (77.5%) supported implementing a shared EU-level legal framework for AI, while 61.1% also endorsed national-level restrictions. These results suggest that participants view ethics not only as a personal concern but as a matter requiring institutional and legal enforcement.

Despite this support, awareness of existing policy initiatives was limited. Only 20% of respondents were aware of the EU’s

AI ethics documentation published in 2019. This gap highlights a disconnect between regulatory efforts and public awareness, raising important questions about communication, education, and democratic participation in AI governance.

F. Personal Ethics and Everyday AI Use

Responses related to personal AI use revealed emerging ethical boundaries. Most participants reported no discomfort using AI for translation tasks. In contrast, approximately one-third expressed guilt or discomfort when using AI to produce work they were expected to complete independently. This distinction suggests an intuitive ethical boundary between assistance and substitution.

Views on the use of AI in education were mixed. When asked about AI-assisted assessment, respondents were nearly evenly divided between approval, rejection, and uncertainty. This ambivalence may reflect concerns about fairness, transparency, and the relational aspects of education, where evaluation is perceived as more than a purely technical process.

G. Skepticism and Critical Attitudes

Finally, participants demonstrated a notable degree of skepticism toward AI-generated and online information. Few respondents believed that tools such as Google or ChatGPT consistently produce reliable answers. This critical stance reflects awareness of the limitations of algorithmic systems and constitutes a crucial ethical disposition in increasingly AI-mediated information environments.

VI. DISCUSSION

The discussion of the results indicates that participants are not merely passive users of AI technologies, but actively engaged in reflecting on their ethical implications. While responses were sometimes uncertain or even contradictory, indifference was rare. This suggests a form of ethical awareness that is still developing but clearly present, offering a meaningful snapshot of youth moral engagement in an age of automation.

A. Between Enthusiasm and Uncertainty

Participants reported widespread and routine use of AI tools, such as ChatGPT, for both work and leisure, indicating that AI has become part of everyday life rather than a speculative technology. However, this familiarity did not translate into complete confidence. Hesitation regarding responsibility and trust was common, particularly in contexts involving automated decision-making.

Rather than indicating ignorance, this hesitation reflects ethical engagement. Participants appear to be navigating complex socio-technical systems whose functioning and implications are not always transparent. This form of “ethical improvisation” highlights the need for education that supports critical reflection rather than offering simplified or definitive answers.

B. Stable Values and Evolving Ethics

Most respondents affirmed that ethical norms should extend beyond legal requirements and supported the existence of universal ethical principles. At the same time, many acknowledged that such principles evolve. This combination suggests a balanced ethical stance, rejecting both rigid moral absolutism and complete relativism. Participants appear willing to negotiate between shared values and contextual adaptation, a key feature of democratic ethical reasoning.

C. Machines, Responsibility, and the Ethical Landscape

A majority of participants supported including machines in ethical reasoning, reflecting a growing awareness that agency in AI-mediated systems extends beyond human actors. However, responsibility attribution in scenarios such as autonomous driving remained fragmented, with programmers, manufacturers, and users all identified as partially responsible.

This diffusion aligns with the concept of the “responsibility gap” identified in the literature, highlighting the difficulty of assigning accountability in distributed socio-technical systems. Importantly, participants were not seeking simple culprits but appeared to intuitively recognize the collective nature of responsibility in AI development and deployment.

D. Regulation: Supported but Poorly Understood

Support for AI regulation at both EU and national levels was strong. Nevertheless, awareness of existing ethical guidelines and policy documents remained limited. This gap points to a disconnect between institutional efforts and public understanding, suggesting that regulatory legitimacy depends not only on policy creation but also on effective communication and civic engagement.

E. Everyday Ethics: Where Help Ends and Cheating Begins

Participants demonstrated emerging personal ethical boundaries in their use of AI. While most were comfortable using AI for translation or support tasks, greater discomfort emerged when AI was used to replace individual work. This distinction between assistance and substitution reflects informal yet meaningful ethical reasoning regarding ownership, effort, and integrity.

Opinions on AI-assisted assessment were divided, indicating concerns about fairness, transparency, and the relational dimensions of education. These findings suggest that the ethical evaluation of AI in education extends beyond technical accuracy to encompass trust in processes and in human judgment.

F. Trust in the Process, Not Just the Output

In high-stakes decisions, participants showed clear reluctance to delegate responsibility to machines. Just 15.9% were willing to accept an AI-based justice system, compared to 52.9% who would accept an AI-developed therapy for Alzheimer’s.

This asymmetry speaks volumes. Even if the outcomes are

reliable, people prefer that decisions about justice be made by humans. What they seek is not just efficiency, but moral legitimacy — the sense that procedures are fair, understandable, and rooted in human values.

AI systems, no matter how accurate, still struggle to earn that trust.

H. Youth as Ethical Stakeholders

What emerges from all of this is a simple but powerful conclusion: young people are already acting as ethical agents in the digital world. They are thinking about fairness, about responsibility, about what kind of technological future they want to be part of.

They are not waiting to be told what is right and wrong. They are working it out — sometimes quietly, sometimes collectively, often without formal guidance.

Projects such as VERFISUM have the potential to make that moral labour visible, supported, and impactful. By placing ethics at the heart of digital education, they recognise that the future of AI is not just technical — it is profoundly human.

G. Everyday Ethics in Educational Contexts

Overall, the findings suggest that young people are already acting as ethical stakeholders in digital environments. They are actively negotiating questions of fairness, responsibility, and acceptable use, often without formal guidance. Initiatives such as the VERFISUM project demonstrate the potential of embedding ethics at the core of digital education, recognizing that the future of AI is not only technical but fundamentally human.

VII. CONCLUSIONS AND RECOMMENDATIONS

This exploratory study aimed to capture how young people and educators perceive the ethical dimensions of AI. Rather than producing definitive conclusions, the findings reveal a dynamic ethical landscape characterized by evolving intuitions, critical awareness, and unresolved tensions.

Participants demonstrated firm ethical commitments, critical attitudes toward AI use, and a willingness to engage with ethical questions despite limited formal guidance. This suggests that AI ethics education should build upon existing moral reasoning rather than assuming an absence of ethical awareness.

A. Educational Recommendations

Integrate ethics into digital curricula

Ethics must not be a footnote to technical training. Whether in vocational programmes, secondary education, or higher education, courses involving AI and digital tools should include explicit ethical components — with time, space, and support for debate, disagreement, and reflection.

Promote critical digital literacy

Going beyond how tools work, students must be equipped to question *why* they exist, *who* they benefit, and *what* values they

encode. This involves understanding not just the technology itself, but also the social and political structures surrounding it (Pangrazio, 2016; Hinrichsen & Coombs, 2013).

Encourage participatory ethics

Rather than treating students as passive recipients of ethical codes, educators can involve them in co-creating classroom guidelines for AI use, or facilitate role-playing debates about public AI policy. Ethics becomes more powerful when it is lived rather than merely taught.

B. Policy Recommendations

Close the awareness gap

The fact that so few participants were aware of existing EU-level AI ethics documents points to a failure in communication. Institutions must translate their policies not only into different languages but into formats that resonate with younger audiences.

Support context-sensitive implementation

Ethics is always culturally situated. EU and national bodies should fund and endorse educational initiatives that adapt ethical guidelines to local realities — including language, digital infrastructure, and educational traditions.

Include young people in AI governance

Youth should not be treated only as beneficiaries or users of technology. They should be **recognised as stakeholders**. Whether through advisory panels, participatory research, or citizen forums, their voices can and should inform the design, implementation, and revision of ethical standards.

C. Research Recommendations

Expand the scope of data collection

Future studies should use larger, stratified samples to enable comparative analysis across demographic groups, countries, and educational settings. This would allow for a deeper understanding of cultural variation in ethical perceptions.

Develop validated ethical literacy instruments

The use of exploratory tools, while valuable, should be complemented in the future by psychometrically robust instruments that assess ethical reasoning and awareness in a replicable manner.

Research longitudinal change

As AI technologies and policies evolve, so too will public perceptions. Longitudinal studies could help track how young people's ethical attitudes shift in response to new technologies, media narratives, and educational interventions.

C. Final Reflections

This study ultimately suggests that the ethical future of AI cannot be built by engineers and lawmakers alone. It must include educators, students, and everyday citizens — those who live with the technology, and who are often the first to sense when something is not quite right.

The young people who participated in this project are not waiting to be told what to think. They are already asking difficult questions about justice, responsibility, and the human costs of automation. What they need now are the tools—intellectual,

emotional, and institutional—to advance those questions.

Projects like VERFISUM demonstrate what is possible when ethics is not treated as an afterthought but as a core component of learning and participation. These findings highlight that preparing young people for a society increasingly shaped by AI also requires promoting human participation in the development and governance of AI systems.

ACKNOWLEDGMENT

This work was developed within the VERFISUM project, co-funded by the European Union under the Erasmus+ Programme. The authors also acknowledge World Talents Programme for its support in the publication of this work.

REFERENCES

- [1] A. MacIntyre, *After Virtue*, 3rd ed. Notre Dame, IN, USA: University of Notre Dame Press, 2007.
- [2] I. Kant, *The Metaphysics of Morals*, M. Gregor, Trans. Cambridge, U.K.: Cambridge University Press, 1996 (original work published 1797).
- [3] P. Boddington, *Towards a Code of Ethics for Artificial Intelligence*. Cham, Switzerland: Springer, 2017.
- [4] M. Coeckelbergh, *AI Ethics*. Cambridge, MA, USA: MIT Press, 2020.
- [5] L. Floridi, *The Ethics of Information*. Oxford, U.K.: Oxford University Press, 2013.
- [6] C. O'Neil, *Weapons of Math Destruction*. New York, NY, USA: Crown Publishing, 2016.
- [7] S. U. Noble, *Algorithms of Oppression*. New York, NY, USA: NYU Press, 2018.
- [8] V. Eubanks, *Automating Inequality*. New York, NY, USA: St. Martin's Press, 2018.
- [9] J. Angwin, J. Larson, S. Mattu, and L. Kirchner, "Machine bias," *ProPublica*, 2016.
- [10] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv:1702.08608*, 2017.
- [11] I. Rahwan et al., "Machine behaviour," *Nature*, vol. 568, no. 7753, pp. 477–486, 2019.
- [12] S. Zuboff, *The Age of Surveillance Capitalism*. New York, NY, USA: PublicAffairs, 2019.
- [13] Future of Life Institute, "Asilomar AI Principles," 2017.
- [14] UNESCO, *Recommendation on the Ethics of Artificial Intelligence*. Paris, France, 2021.
- [15] European Commission, *Regulation ... Artificial Intelligence (AI Act)*. Brussels, Belgium, 2024.
- [16] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," *Nature Machine Intelligence*, vol. 1, no. 9, pp. 389–399, 2019.
- [17] W. Holmes, M. Bialik, and C. Fadel, *Artificial Intelligence in Education*. Boston, MA, USA, 2019.
- [18] R. Luckin, *Machine Learning and Human Intelligence*. London, U.K.: UCL IOE Press, 2018.
- [19] J. Hinrichsen and A. Coombs, "The five resources of critical digital literacy," *Research in Learning Technology*, vol. 21, 2013.
- [20] L. Pangrazio, "Reconceptualising critical digital literacy," *Discourse*, vol. 37, no. 2, pp. 163–174, 2016.
- [21] M. Stoilova, R. Nandagiri, and S. Livingstone, "Children's understanding of personal data and privacy online," *Information, Communication & Society*, vol. 24, no. 4, pp. 547–567, 2021.
- [22] OECD, *AI and the Future of Skills*, Vol. 1. Paris, France: OECD Publishing, 2021.
- [23] European Commission, *Ethics Guidelines for Trustworthy AI*. Brussels, Belgium, 2019.