

What Would You do? An Ethical AI Quiz

Wei Teo*, Ze Teoh*, Dayang Abang Arabi*, Morad Aboushadi*, Khairrenn Lai*, Zhe Ng*,
Aastha Pant[†], Rashina Hoda[†], Chakkrit Tantithamthavorn[†], Burak Turhan^{†‡}

*Faculty of Engineering, Monash University, Subang Jaya, Malaysia

[†]Faculty of IT, Monash University, Melbourne, Australia

[‡]Faculty of Info. Tech. and Electrical Eng, University of Oulu, Finland

Email: {aastha.pant, rashina.hoda, chakkrit}@monash.edu, burak.turhan@oulu.fi

Abstract—The resurgence of Artificial Intelligence (AI) has been accompanied by a rise in ethical issues. AI practitioners either face challenges in making ethical choices when designing AI-based systems or are not aware of such challenges in the first place. Increasing the level of awareness and understanding of the perceptions of those who develop AI systems is a critical step toward mitigating ethical issues in AI development. Motivated by these challenges, needs, and the lack of engaging approaches to address these, we developed an interactive, scenario-based *ethical AI quiz*. It allows AI practitioners, including software engineers who develop AI systems, to self-assess their awareness and perceptions about AI ethics. The experience of taking the quiz, and the feedback it provides, will help AI practitioners understand the gap areas, and improve their overall ethical practice in everyday development scenarios. To demonstrate these expected outcomes and the relevance of our tool, we also share a preliminary user study. The video demo can be found at <https://zenodo.org/record/7601169#.Y9xgA-xBxhF>.

Index Terms—ethics, AI ethics, AI practitioners, self-assessment tools, ethical AI quiz

I. INTRODUCTION

With the increased use of AI in different domains such as health, education, transport, and banking, there is also a rise in the number of associated ethical issues. For example, Google’s machine learning (ML) algorithms turned out to be gender-biased, associating men with Science, Technology, Engineering, and Mathematics (STEM) careers more frequently than women [20]. While GitHub’s unauthorised and unlicensed use of copyrighted source code as training data for their ML-powered *GitHub Copilot* is an example of a privacy issue [1].

Research shows that people involved in the design and development of AI-based systems, such as AI developers, AI experts, AI scientists, and AI engineers – collectively referred to as AI practitioners – use various strategies to enhance the implementation of ethics and ethical principles during the development of AI-based systems [4], [5], [6]. For example, to create transparency, AI practitioners use strategies such as *conducting audits* [2] and *code documentation* [7]. *Speculating on socio-ethical impacts* [3], and *group discussion* with colleagues [4] are other strategies to improve the awareness of ethical issues. Although several tools are available in the market, such as *Deon* [9], *AI Fairness 360* [10], *Fairlearn* [11], there is little information on their practical use. Regardless, AI practitioners continue to face challenges in making ethical choices on a daily basis [5].

It is essential to understand AI practitioners’ viewpoints [2], awareness [2], [6], and perception [5], [8] of AI ethics, to mitigate ethical issues. Ethical AI education and training aims to address this need [21]. However, traditional classroom-based and didactic approaches typically lack engagement and can be ineffective [19]. These aspects, along with the current lack of support for AI practitioners to self-assess, motivated us to design and develop an interactive *ethical AI quiz*¹, to allow AI practitioners to self-assess their awareness, understanding, and perceptions of AI ethics and ethical principles. In this paper, we share the details of the *quiz* design and development, a limited trial with software engineering students, and implications for practice.

II. ETHICAL AI QUIZ – DESIGN

We aimed to develop an interactive *ethical AI quiz* in the form of a web application to help AI practitioners self-assess their knowledge of AI ethics and ethical principles. In taking the quiz, AI practitioners face a set of 4 scenarios with 3-4 questions each (13 in total) that present ethical dilemmas to varying degrees. Users of the quiz can select from four answer options for each question. The quiz includes a *Result Summary* page at the end that enables AI practitioners to view their overall results. The tool also lets participants view the summary of the result for each question and receive feedback. The key research and design challenges we faced in designing the quiz included: drafting the quiz content (scenarios, questions, and answer options); coming up with a meaningful scoring and feedback system; and designing an appealing UI design and engaging experience, as presented below.

A. Quiz Content

Our primary focus in the initial design phase was brainstorming the potential quiz questions and finalising the contents of the quiz. We decided that our quiz would consist of scenario-based questions that pose hypothetical events involving AI ethics, and users must select answers that best fit their ethical values. Scenario-based questions are the best to ask, as they are more engaging and can help stimulate a person’s thought and decision-making process [12]. Two of the four scenarios were designed based on real-world

¹Ethical AI Quiz: <https://interactive-ai-ethics-quiz.herokuapp.com/>

AI ethical issues and incidents, whereas two scenarios were designed by the team using the Design Science Research Methodology (DSRM). The main aim of a DSRM is to develop knowledge and apply it to designing effective artifacts [18]. Care was taken to use language that was unbiased and non-leading. Developing the scenarios, questions, and answers involved brainstorming among the team members (six students and four supervisors). Since *ethics* is a topic with multiple grey areas and interpretations, the team carefully modeled the scenarios, questions, and answer options through multiple rounds of brainstorming and referencing real-life scenarios. After multiple rounds of discussions, reviews, and revisions, we finalised the quiz scenarios, questions, and answer options.

We referred to *Australia’s AI ethics principles* [13] to create our quiz content because our primary focus is to evaluate the quiz with AI practitioners based in Australia. Initially, we planned to develop the quiz scenarios and questions based on all eight ethical principles. However, we decided to focus only on four ethical principles: accountability, fairness, privacy, and human, societal & environmental well-being because of two reasons: (i) we did not want to clutter the *quiz* with too many scenarios and questions, and (ii) it was challenging to design the scenarios and quiz questions based on other principles, an aspect we aim to explore in future studies. All scenarios and questions are available on the free online *quiz*.

The questions allow AI practitioners to self-assess their awareness, perceptions, and understanding of AI ethics. By responding to scenarios posing ethical dilemmas, the quiz can test the AI practitioner’s knowledge and application of ethical principles. Additionally, the quiz can help reveal if there are certain biases for particular ethical principles present across the industry.

B. Scoring and Feedback

We wanted to maintain an interactive and ‘quiz-like’ feel, so the users can engage in the process. We also wanted to provide AI practitioners with feedback on their answer choices so they can learn about AI ethics through the process of taking the quiz. We designed a scoring system to achieve the first and a feedback system to achieve the second point. The answer options were designed to be sufficiently nuanced to avoid clear right and wrong answers. After much deliberation, we settled on ‘least desirable’, ‘bearable’, ‘less than ideal’, and ‘ideal’, with a score of 1, 2, 3, and 4, respectively. The final score is obtained by adding each question’s score, and the overall feedback is provided based on the final score. The score is not visible to the users, but the feedback is made available at the end of the quiz. The results are stored in the database, where further processing can be performed. Table I describes the overall feedback provided to users based on their scores. Additionally, they receive feedback for each question and can view the answer they selected and the ideal answer for each question. Users can review how they did for each question as the justification for the correct answer is explained to boost experiential learning (see Fig. 2).

TABLE I
SCORES AND FEEDBACK OF THE QUIZ

Score	Feedback
76-100%	Excellent attempt. Congratulations! You are an AI ethics expert.
51-75%	Very good attempt. You are on track to becoming an AI ethics expert.
26-50%	Satisfactory attempt. With little more knowledge, you can become an AI ethics expert.
0-25%	Unsatisfactory attempt. There is much room for improvement.

If users want to participate in a future study, they can choose to share their emails, through a text box at the bottom of the result page, to set up an interview with the research team. In the interview, they can share their thoughts and opinions on ethics and their personal experiences with the researchers.

C. Interface and Engagement

Another important aspect of our design process involved researching modern user interface (UI) and user experience (UX) approaches. For the first step in our UX research, we looked at the different design principles available to us. Specifically, we explored and incorporated these three design principles into our web design: Norman’s Design Principles [14], Schneiderman’s 8 Golden Rules of Design [15], and Fitt’s Law [16]. To integrate accessibility into our design, we also followed the Web Content Accessibility Guidelines (WCAG) 2.0 [17] to ensure usability for diverse users. These design principles and guidelines played a vital role in developing our quiz as they focused on providing a positive experience to users when taking the quiz. For the design of the quiz interface, a series of wireframes of each screen highlighting different design principles were fabricated by the students and reviewed by the supervisory team. Continuous and iterative feedback from the supervisory team led to the refinement and review of each wireframe before reaching a consensus on the final wireframe design. We discussed the design considerations regarding the placement of scenarios, questions, and answer options and what images and background themes should be incorporated to make the *quiz* look aesthetically pleasing and engaging for users. Figure 1 shows the final wireframe design along with an example of one of the scenarios used.

The flow of the *quiz* starts with a friendly and minimalist home page prompting users to read the explanatory statement. Once accepted, the user is redirected to the scenario and quiz questions, where each question takes up the entire screen to enhance readability. The user receives feedback on their overall performance upon completing the *quiz* and can access feedback on individual answers.

III. ETHICAL AI QUIZ – DEVELOPMENT

Three major components comprise our interactive *ethical AI quiz*, which includes:

²<https://storyset.com/>

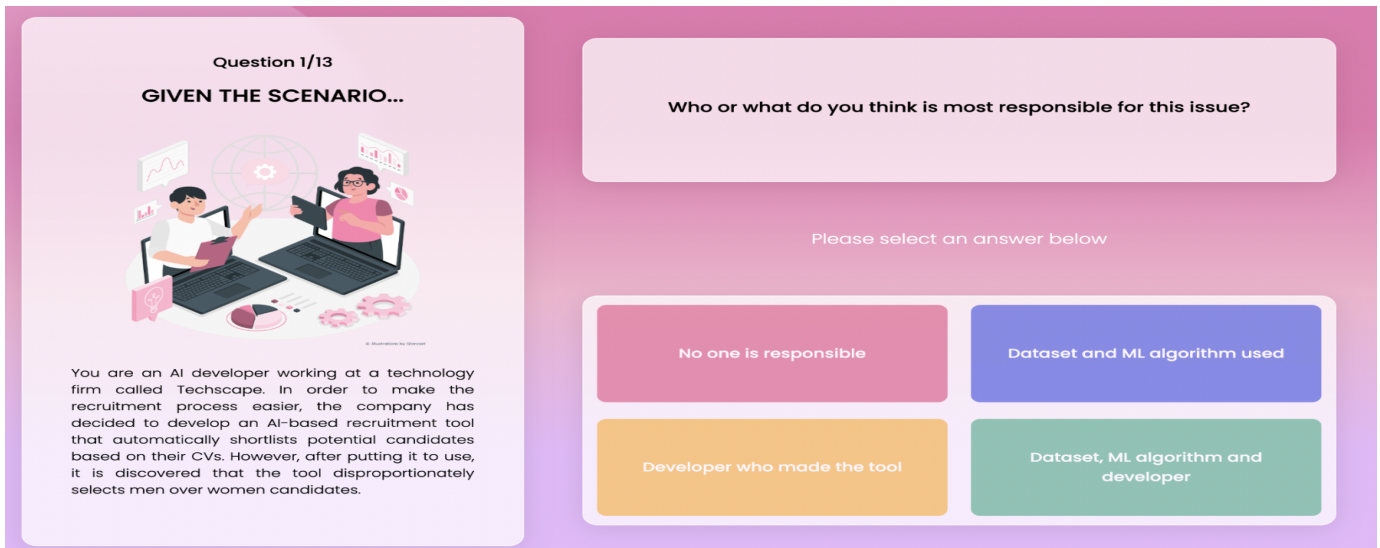


Fig. 1. A scenario and associated question and answer options in the *Ethical AI Quiz*.²

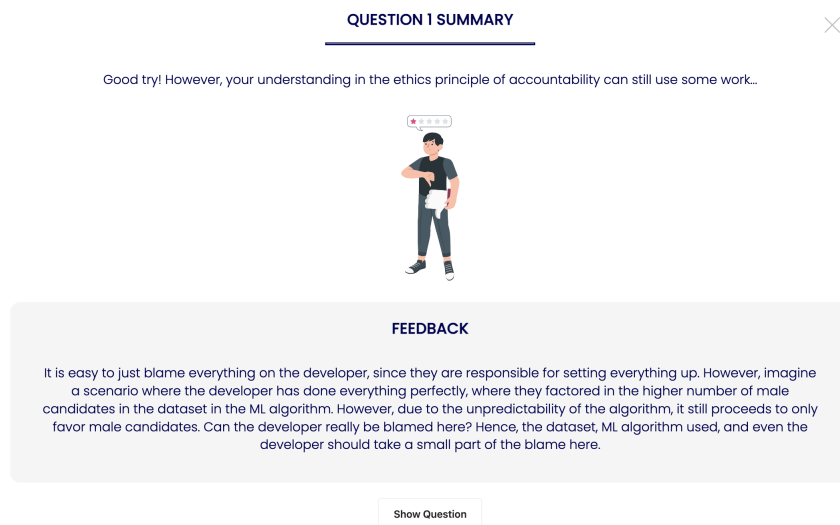


Fig. 2. Feedback on the selected answer, highlighting underlying ethical AI principles, the rationale for why it is not the ideal answer, and sharing the ideal answer.²

Figma³ - It is a design tool that helps to create designs for mobile and web interfaces or any design that one can think of. Using Figma, our team could picture the flow and design a friendly and engaging *quiz* application.

React⁴ - The web app was built using React, a JavaScript library for building user interfaces. We used this front-end framework to develop the UI code for our interactive *quiz*. We hosted our *quiz* at Heroku⁵, a cloud platform that enables developers to build, run and operate applications entirely in the cloud. We could use Heroku for free because our *quiz* was lightweight and did not require much bandwidth.

³<https://www.figma.com/>

⁴<https://reactjs.org/>

⁵<https://id.heroku.com/login>

Firestore⁶ - We required a database to store the user information and their scores and feedback, for which we used Firestore. Firestore was chosen as the database engine because the team members had experience using it, which was favorable for development.

IV. TRIAL AND REFINEMENT

A. Evaluation method

Due to the time constraints of this research project, the *quiz* was evaluated with the students who had taken AI courses in the past and enrolled in a Bachelor of Software Engineering (Honours) at Monash University undertaking Software Engineering Research Project unit. While preparing for user

⁶<https://console.firebase.google.com/u/0/>

evaluation, ethics approval was taken from the Monash University Human Research Ethics Committee (Approval Number: 35521). Upon getting ethics approval, we requested the unit coordinators to post the *quiz* link on the unit forum. The participation from the students was completely voluntary. We gathered a total of 50 responses from students. We exported the data from the database in the .csv file format and analysed the data.

B. Evaluation results

1) *Overall results:* According to the results, only 20% of the participants scored between 76%-100% whereas 76% of the participants scored between 51%-75%. It was worth noting that only 4% of the participants scored between 26%-50%, and none scored below 25%, indicating a basic understanding of ethical AI principles.

2) *Results based on ethical principles:* The results show that questions related to ethical principles like ‘*privacy*’ and ‘*human, societal, and environmental well-being*’ had the highest number of ‘*ideal*’ answers i.e. 35.33% for each ethical principle. It indicates that participants were more confident in answering questions related to these two ethical principles. On the other hand, questions related to ‘*accountability*’ had the least number of ‘*ideal*’ answers, i.e. 27.34%.

Furthermore, the results show that questions related to ‘*fairness*’ had the highest number of ‘*least desirable*’ answers, i.e. 28.5% followed by ‘*privacy*’ with 25.33%. Questions related to the ‘*human, societal and environmental well-being*’ had the lowest percentage of ‘*least desirable*’ responses with 20.67%. Figure 3 shows the percentage of all responses based on ethical principles.




Fig. 3. Percentage of all responses based on ethical principles

3) *Results based on quiz questions:* Results show that a question (*Q4*) had the most ‘*ideal*’ answers (50%), followed by two questions (*Q8 and Q10*) with 42% each. Furthermore, *Q5* had the highest percentage of ‘*less than ideal*’ answers with 44% followed by *Q2* with 38%.

In addition, results show that *Q3* had the highest percentage of the ‘*bearable*’ answers, i.e. 36% whereas, *Q1* had the highest percentage of the ‘*least desirable*’ answers, i.e. 66% followed by *Q9* with 44%. *Q8* had the lowest ‘*least desirable*’ answers, i.e. only 12%, which indicates that most of the participants were confident in answering it. Figure 4 shows the percentage of all responses based on the *quiz* questions. This is the link to the pdf document containing screenshots of the whole *quiz*: <https://zenodo.org/record/7601202#.Y9xTtexBxhE>.




Fig. 4. Percentage of all responses based on *quiz* questions

C. Refinement

After evaluating it with students, we made minor refinements to the back end, such as collecting the date and time the quiz was taken and including an option to select the user type at the beginning of the quiz: AI student, AI practitioner, AI researcher, Working in AI Domain, and other. This was done in preparation for making the quiz available beyond students, AI practitioners, and others as a next step. Ethics approval for the same has been also sought. We plan to make the quiz widely available in the industry and to the general public, and assess how the findings vary by user type.

V. IMPLICATIONS FOR PRACTICE

Our *ethical AI quiz* allows AI practitioners to consider scenario-based questions and select an answer from four answer options. As a result, they can self-assess their knowledge and understanding of AI ethics and ethical principles and better understand ethical requirements.

The experience of taking the quiz, and the feedback it provides, will help AI practitioners understand the gap areas, debate and discuss ethical AI principles and issues together, and improve their overall ethical practice in everyday development scenarios.

VI. CONCLUSION

In this paper, we have presented an interactive *ethical AI quiz* that helps AI practitioners self-assess their knowledge of ethics in AI. Using scenario-based questions, the *quiz* assesses AI practitioners’ knowledge of AI ethics from two angles, namely, their perception and approach to ethics in AI. In addition, it summarises results to help AI practitioners understand how ethical they are when facing challenging ethical AI scenarios. We evaluated the *quiz* with software engineering students and plan to evaluate it with real-world AI practitioners in the future.

REFERENCES

- [1] Matthew Butterick. “Github Copilot Litigation”. (2022). <https://githubcopilotlitigation.com/> [accessed 8- November-2022].
- [2] Ville Vakkuri, Kai-Kristian Kemell, Joni Kultanen, and Pekka Abrahamsson. 2020. “The current state of industrial practice in artificial intelligence ethics.” *IEEE Software*, 37(4), 50–57.
- [3] Ville Vakkuri, Kai-Kristian Kemell, and Pekka Abrahamsson. 2019. “Implementing ethics in AI: Initial results of an industrial multiple case study,” in *International Conference on Product-Focused Software Process Improvement*, 331–338.

- [4] Ville Vakkuri, Kai-Kristian Kemell, and Pekka Abrahamsson. 2019. "Ethically aligned design: An empirical evaluation of the Resolvedd-strategy in software and systems development context," in *45th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, 46–50.
- [5] Will Orr and Jenny L. Davis. 2020. "Attributions of ethical responsibility by artificial intelligence practitioners," *Information, Communication & Society*, 23(5), 719–735.
- [6] Ryan Mark and Gregory Anya. 2020. "Ethics of using smart city AI and big data: The case of four large European cities," *The ORBIT Journal*, 2(2), 1–36.
- [7] Josephine Seah and Mark Findlay. 2021. "Communicating ethics across the AI ecosystem," *SMU Centre for AI & Data Governance Research Paper*.
- [8] Javier Camacho Ibanez and Monica Villas Olmeda. 2022. "Operationalising AI ethics: How are companies bridging the gap between practice and principles? An exploratory study," *AI & Society*, 37, 1663–1687.
- [9] "Deon". (2022). <https://deon.drivendata.org/> [accessed 28-October-2022].
- [10] "AI Fairness 360". (2022). <https://aif360.mybluemix.net/> [accessed 29-October-2022].
- [11] "Fairlearn". (2022). <https://github.com/fairlearn/fairlearn/> [accessed 31-October-2022].
- [12] Adam Lally, Sugato Bagchi, Michael A. Barborak, David W. Buchanan, Jennifer Chu-Carroll, David A. Ferrucci, Michael R. Glass, Aditya Kalyanpur, Eric T. Mueller, J. William Murdock, Siddarth Patwardhan and John M. Prager. 2017. "WatsonPaths: Scenario-based question answering and inference over unstructured information", *AI Magazine*, 38(2), 59–76.
- [13] "Australia's AI Ethics Principles". 2022. <https://www.industry.gov.au/publications/australias-artificial-intelligence-ethics-framework/australias-ai-ethics-principles/> [accessed 29-October-2022].
- [14] Donald A. Norman. 1983. "Design principles for human-computer interfaces", in *CHI '83: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1-10.
- [15] Ben Shneiderman, Catherine Plaisant, Maxine S. Cohen, Steven Jacobs, Niklas Elmqvist, and Nicholas Diakopoulos. 1998. *Designing the user interface: Strategies for effective human-computer interaction*, Addison Wesley Longman, 6th edn., MA.
- [16] Scott I. MacKenzie. 1992. "Fitts' law as a research and design tool in human-computer interaction", *Human-Computer Interaction*, 7(1), 91-139.
- [17] Ben Caldwell, Michael Cooper, Loretta Guarino Reid, and Gregg Vanderheiden. (2008). Web Content Accessibility Guidelines (WCAG) 2.0. Web Accessibility Initiative (WAI), World Wide Web Consortium (W3C). <http://www.w3.org/TR/WCAG20/> [accessed 31-October-2022].
- [18] Bill Kuechler and Vijay Vaishnavi. 2008. "On theory development in design science research: Anatomy of a research project", *European Journal of Information Systems*, 17(5), 489-504.
- [19] Erin S. Lane and Sara E. Harris. 2015. "A new tool for measuring student behavioral engagement in large university classes", *Journal of College Science Teaching*, 44(6), 83-91.
- [20] Marcelo O R Prates, Pedro H Avelar, Luis C Lamb. 2020. "Assessing gender bias in machine translation: A case study with Google translate," *Neural Computing and Applications*, 32, 6363-6381.
- [21] Veronika Bogina, Alan Hartman, Tsvi Kuflik and Avital Shulner-Tal. 2022. "Educating Software and AI Stakeholders About Algorithmic Fairness, Accountability, Transparency and Ethics", *International Journal of Artificial Intelligence in Education*, 32, 808–833.