Journal of Educational Studies in Mathematics and Computer Science

Using Image Classification to Teach Machine Learning to Gifted Elementary School Students

Milena Živković[1, a,*], Dubravka Živković[2, b], Svetlana Radojević[3, c]
[1]University of Kragujevac, Faculty of Science, Kragujevac, Serbia
[2]University of Niš, Faculty of Philosophy, Niš, Serbia
[3]The First Grammar School of Kragujevac, Kragujevac, Serbia
[a] milena.zivkovic@pmf.kg.ac.rs, [b]dubravkazivkovic3@gmail.com, [c]svetlana.radojevic@prvagimnazija.edu.rs

Received: 15 May 2025; **Accepted:** 22 July 2025

Abstract

Artificial intelligence and machine learning are becoming increasingly integral to modern life, yet these concepts remain predominantly absent from elementary education. This study presents a project-based learning activity in which talented students enrolled in the Centre for Talented Youth Kragujevac trained their image classification models using Google's Teachable Machine, a no-code web tool. During five sessions, students (aged 13-14) engaged in data collection, model training, testing, and reflection without writing a single line of code. The activity aimed to demystify artificial intelligence by making it hands-on, collaborative, and accessible. Students demonstrated strong engagement and produced models with accuracy ranging from 84% to 96% while reflecting critically on challenges such as dataset bias, overfitting, and the ethical use of artificial intelligence. The findings suggest that integrating creative, no-code platforms can foster foundational digital literacy, promote ethical awareness, and spark student interest in emerging technologies.

Keywords: Project-Based Learning, No-Code Tools, Education, Elementary School, Teachable Machine, Image Classification.

MSC2020: 97P80

1. Introduction

Artificial Intelligence (AI) and Machine Learning (ML) technologies are rapidly transforming the ways people interact with the world, influencing sectors ranging from healthcare and agriculture to education and creative industries (Alowais & Alowais, 2023; Topol, 2019; Zawacki-Richter et al., 2019; Chen et al., 2020). As these technologies become increasingly embedded in everyday life, educational systems must equip students with critical thinking skills to understand and engage with them. Despite the growing relevance of AI, systematic instruction in these areas is still largely absent from elementary education. Barriers such as the perceived technical complexity of AI, insufficient teacher training, and the lack of age-appropriate tools have hindered its early integration into the classroom environment (Luckin et al., 2016; Ormrod, 2016). As a result, students often encounter AI at the university level for the first time, missing valuable opportunities to develop foundational digital literacy, conceptual understanding, and ethical reasoning at an earlier stage (Ormrod, 2016; Kurz et al., 2024). Recent research has highlighted the potential of no-code educational tools that allow young learners to explore AI through intuitive, hands-on experiences.

* Corresponding author
E-mail address: milena.zivkovic@pmf.kg.ac.rs

Original Research Article

One of these tools is Google's Teachable Machine (TM), which has gained popularity due to its ease of use and accessibility. It enables users to train models for image, sound, and pose classification directly in a web navigator without requiring any programming experience or advanced technical knowledge (Google, n.d.). Studies have demonstrated that TM is not only technically effective, achieving over 90% accuracy in domains such as plant disease identification (Odabaş et al., 2024), herbal medicine classification (Malahina et al., 2024), and health diagnostics (Alowais & Alowais, 2023), but also pedagogically valuable. In educational contexts, research demonstrates that TM promotes student engagement, fosters interdisciplinary learning, and supports the development of computational thinking skills (Kurz et al., 2024; Relmasira & Vartiainen, 2025).

In recent years, several initiatives across Europe and globally have aimed to bring artificial intelligence into K-12 education using age-appropriate methods. For example, the European Commission has supported the development of the AI4K12 framework in collaboration with international partners to create guidelines and resources for teaching AI in schools (Touretzky et al., 2019). Similarly, the "Elements of AI" course, developed in Finland by the University of Helsinki and Reaktor, has been widely adopted to increase AI literacy among both students and teachers (Vuorikari et al., 2016). In the UK, the Raspberry Pi Foundation and Google's "Machine Learning for Kids" project have made machine learning concepts accessible through block-based coding environments and classroom activities (Lane, 2021). These initiatives highlight the growing interest in demystifying AI at an early age and align with the present study's goal of using no-code tools like Teachable Machine to foster early understanding of AI concepts.

In Serbia, the national Strategy for the Development of Artificial Intelligence 2020 – 2025 calls for the gradual introduction of AI-related content into the school system. As part of this effort, seventh-grade compulsory subjects "Engineering and Technology" and "Informatics and Computing" now include basic principles of intelligent algorithms and data processing; however, no standalone modules on these topics are offered in lower elementary grades (Government of the Republic of Serbia, 2020). Moreover, although two optional AI-fundamentals courses have been made available at the primary level and three electives at the secondary (grammar school) level, explicit instruction in these technologies remains peripheral to the core curriculum and is often delivered only through extracurricular or pilot initiatives (Badža, 2024).

This paper presents an exploratory project involving gifted elementary school students who created image classification models using TM. Over five sessions, students engaged in the complete machine learning cycle: selecting a topic, collecting training data, building a model, testing its performance, and reflecting on results – all without writing code. Alongside the technical process, students discussed key concepts such as bias, generalization, and overfitting and explored the ethical implications of using facial images in AI systems. Moreover, the project invited comparisons between machine learning processes and classical theories of human learning, such as Thorndike's trial-and-error learning and Köhler's insight learning, offering a unique interdisciplinary framework for understanding how both humans and machines acquire knowledge (Slavin, 2018; Thorndike, 1911; Köhler, 1925). The primary aim of this study was to demystify AI by making it tangible, reflective, and age-appropriate, thereby fostering technical awareness, ethical sensitivity, and cognitive curiosity. The findings suggest that even young students, with appropriate guided pedagogical strategies and tools, can meaningfully engage with artificial intelligence in educational and empowering ways.

2. Psychological and Ethical reflections in AI-Based learning

In addition to technical and pedagogical considerations, the integration of artificial intelligence into student learning environments raises ethical and psychological questions. Some projects involved the use of facial images to classify emotions or expressions. While this helped students understand how AI can interpret visual cues, it also prompted crucial discussions about data privacy, especially when collecting images of children. Questions such as *Where do these images go?*, *Can*

AI keep them forever?, or *Is it safe to share my face with a machine?,* became natural entry points to explore broader concerns related to data storage, biometric surveillance, and the long-term implications of training AI on personal data. These concerns align with ongoing global debates about the ethical use of facial recognition technology, particularly among vulnerable populations such as minors. Recent reviews underscore the complexity of these issues, highlighting the need for clear guidelines and teacher training when deploying AI in K-12 settings (Gouseti et al., 2024). On the psychological side, we can meaningfully compare AI's learning mechanisms with theories of human learning. For example, AI models improve performance through trial and error, gradually adjusting internal parameters to minimize error, a process that mirrors Thorndike's theory of learning by trial and error. In his work with animals, Thorndike demonstrated that repeated attempts, reinforced by success, led to more efficient behavior, a dynamic similar to how AI optimizes performance through repeated iterations and feedback loops (Thorndike, 1911). Alternatively, some aspects of machine learning align with insight learning, as described by Wolfgang Köhler, where problem-solving emerges from a sudden reorganization of information (Köhler, 1925). While AI typically does not reason in the human sense, more advanced models capable of pattern recognition and abstraction can mimic this form of sudden generalization after sufficient exposure to structured input data. These parallels offer a compelling framework for students to reflect on how AI works, how humans learn, and how both systems improve through structured experiences, feedback, and adaptation. Encouraging students to explore these similarities can deepen their metacognitive awareness and foster a richer understanding of learning as a concept that transcends both biology and technology.

3. Material and Methods

In this study, the participants were students (four students aged 13-14) enrolled in a Center for Talented Youth Kragujevac in Kragujevac, Serbia. Parental consent was obtained for all participants, and no identifiable personal data was collected or stored. These students are part of a specialized program designed for gifted youth with a demonstrated interest and aptitude in science, technology, engineering, and mathematics (STEM). The activity was conducted over five consecutive 90-minute sessions. Most of them had no prior experience with artificial intelligence or programming. The goal was to introduce them to core machine learning concepts through practical, hands-on experimentation using a no-code platform.

The platform selected for the project was TM (Google, n.d.), a free, browser-based tool developed by Google. It enables users to train simple image, audio, or pose classification models using a graphical interface and real-time webcam input. Upon starting a new project, users can choose between different types of models, such as image, audio, or pose recognition. The platform offers two options for image-based classification: a standard image model optimized for general use (224×224px color images) and an embedded model suitable for microcontrollers using smaller grayscale images (96×96px). Users capture images directly via webcam or upload files, label them into categories, and initiate training with a single click (Fig. 1). The tool then provides real-time feedback on model performance and exports the trained model to various formats such as TensorFlow, TensorFlow Lite, or TF.js for integration into applications or microcontroller environments.

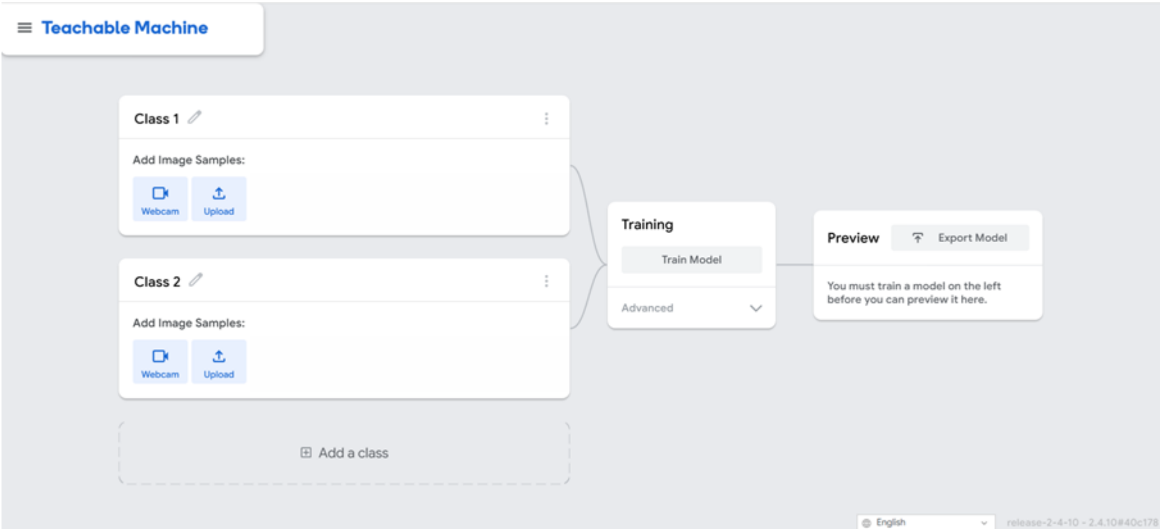

Fig. 1. TM starting interface.

This study exclusively used the image-classification module. Table 1 summarizes details of each group's seven project topics (A-G), number of classes, and dataset size.

Table 1. Overview of student projects: topics, number of classes, and dataset size

Model	Topic	Number of Classes	Images per Class
A	Facial expressions (happy/sad)	2	~150
B	Recyclables (plastic/metal/paper)	3	~150
C	Hand gestures (rock/paper/scissors)	3	~70
D	Fruits (apple/banana/orange)	3	~150
E	Rock types (igneous/sedimentary/metamorphic)	3	~100
F	Emotions (happy/sad/angry)	3	~100
G	School objects (pen/eraser/ruler)	3	~100

The four participating students worked as a single collaborative group and explored a total of seven different project themes during the activity. Rather than each student selecting a single unique project, they worked individually or in pairs on multiple topics, depending on their interests and available time. Some projects were initiated by one student, while others were jointly developed and tested by the group. Examples of chosen topics included classifying emotions based on facial expressions, differentiating hand gestures used in sign language, identifying types of recyclable materials such as plastic, metal, and paper, and recognizing fruits or everyday classroom objects. Each group followed a structured workflow (Table 2).

Table 2. Project workflow phases in the machine learning activity.

Phase	Description of Activities
1. Planning Phase	Students defined their classification goal, brainstormed class labels (categories), and anticipated the types of data needed for effective model training.
2. Data Collection	Using TM's built-in webcam tool, students captured between 50 and 150 images per class, depending on the group and topic complexity. They varied the lighting, background, and angles to prevent overfitting.
3. Model Training	Students trained their models using default settings, observing accuracy and loss metrics. Most groups iterated the process to improve performance.
4. Evaluation and Testing	Students tested their models with new, unseen inputs (sometimes from peers) to assess generalization and robustness.
5. Documentation and Reflection	Each student maintained a learning journal documenting their process, model behavior, challenges (e.g., class imbalance), and improvement ideas.

Throughout the activity, brief plenary discussions were held between sessions to introduce key AI concepts (e.g., training data, overfitting, bias, validation) and to allow peer sharing of results and ideas. Facilitators guided students with probing questions rather than direct solutions, encouraging critical thinking and iterative experimentation.

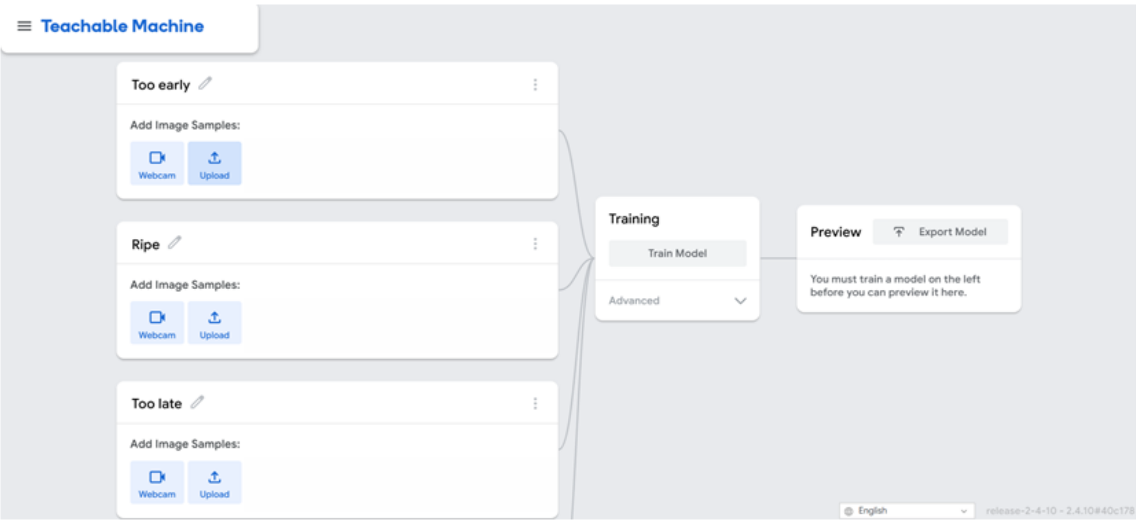

Fig. 2. Example of a student project using TM to classify fruit ripeness into three categories: "Too early," "Ripe," and "Too late."

One notable example involved a student group that trained a model to recognize fruit ripeness based on visual cues. They created three categories – *Too early, Ripe,* and *Too late* – and collected dozens of images for each class using a webcam and printed fruit images (Fig. 2). This project demonstrated the model's ability to differentiate subtle color and shape changes associated with ripening.

Even though TM is designed as a beginner-friendly tool, it provides limited access to these training parameters, allowing for structured experimentation within a controlled classroom setting.

Table 3 illustrates a potential configuration of parameter values that could be used for such exploratory studies.

Table 3. Testing parameters with epoch value, learning rate and batch size.

Parameter	Suggested values					
Epoch	10	50	100	250	750	1000
Learning rate	0.00001	0.0001	0.001	0.01	0.1	1
Batch size	16	32	64	128	256	512

Incorporating controlled manipulation of these parameters into the learning process could enable students to go beyond basic model creation and begin to understand the underlying dynamics of training processes.

4. Results

Across the activity, the four students created a total of seven image classification models. While they collaborated as a group during discussions and feedback sessions, the hands-on work on each model was conducted either individually or in pairs, depending on student interest and project complexity. Some students initiated more than one model, and in some cases peers assisted with data collection and testing. Table 4 presents an overview of model performance based on self-reported accuracy using TM's built-in validation tool.

Table 4. Overview of model configuration and performance for each student project.

Model	Topic	Number of classes	Total images	Accuracy (%)
A	Facial expressions	2	200–400	93%
B	Recyclables	3	200–400	87%
C	Hand gestures (RPS)	3	100–200	91%
D	Fruits	3	300–400	84%
E	Rock types	3	200–400	96%
F	Emotions (happy/sad/angry)	3	100–200	89%
G	School objects	3	100–200	85%

Table 4 presents the technical setup and observed performance of each student project. All models were trained using TM's default parameters, which typically included 50 training epochs and a 16-batch size. No modifications to the learning rate or other hyperparameters were made, as the primary goal was to focus on the intuitive and hands-on nature of training AI models without coding. The projects covered a range of classification tasks, such as recognizing facial expressions, sorting recyclables, interpreting hand gestures, and identifying school-related objects or natural items like fruits and rocks. Each project involved between 200 and 400 images, split across 2 to 3 classes depending on the chosen theme. Despite the simplicity of the models and limited training data, student projects achieved high reported accuracy, with most models scoring between 84% and 96%.

Several students recognized the impact of dataset size and balance on model accuracy. Many also experimented with re-training their models after adjusting conditions such as lighting or adding more samples to underrepresented classes.

a)

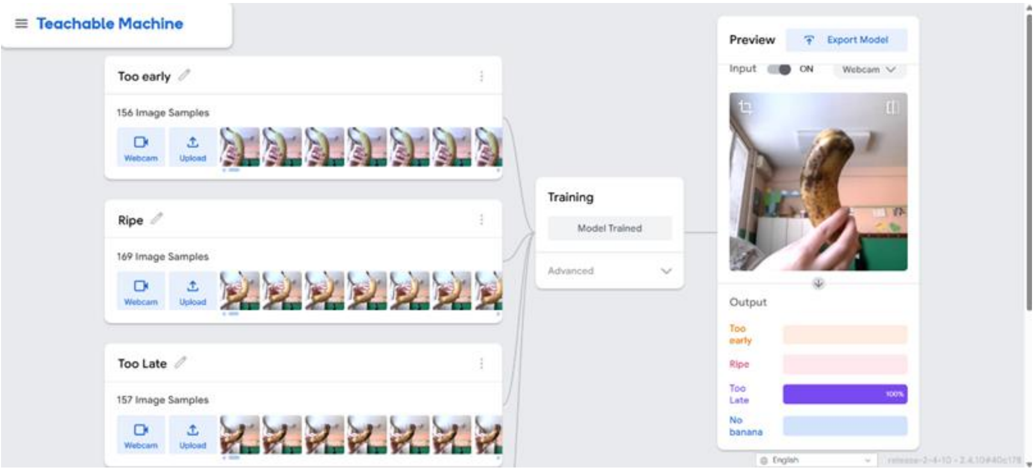

b)

Fig. 3. Example of a trained model used to classify banana ripeness
(a) Dataset samples showing different ripeness stages of bananas.
(b) Model output correctly classifying an overripe banana with 100% confidence.

Fig. 3a and 3b present a practical example from one of the student projects focused on classifying banana ripeness into three stages: *Too early, Ripe*, and *Too late.* The dataset included over 150 image samples per class, collected via webcam. The model achieved high confidence in classification, as shown in the example where a banana at a late ripeness stage was correctly identified with 100% certainty.

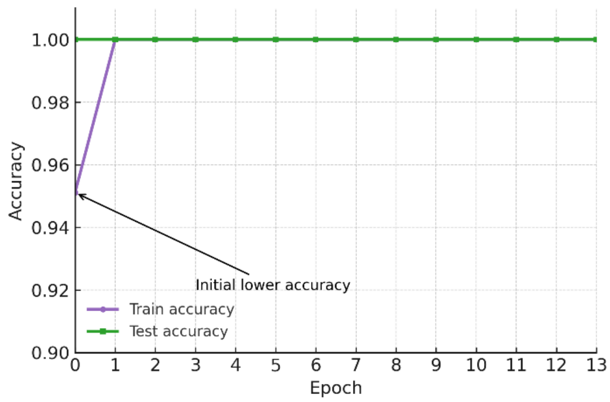

a)

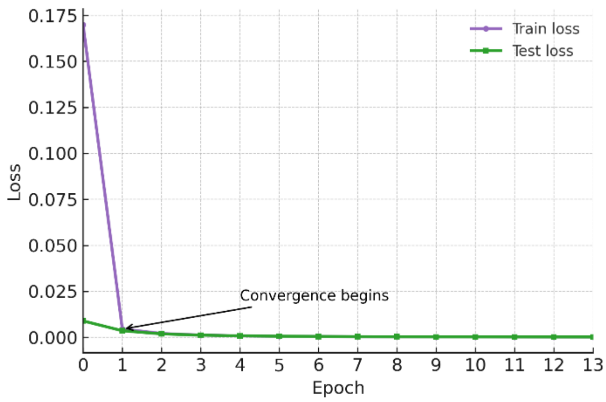

b)

Fig. 4. Training results for the banana ripeness classification model
(a) Accuracy per epoch showing perfect classification achieved by epoch 16.
(b) Loss per epoch indicating rapid convergence with minimal overfitting.

Fig. 4a–b summarize the banana-ripeness model's behavior under TM's default settings: convergence on perfect accuracy and loss reduction within the first few epochs.

5. Discussion

While elementary school is rarely thought of as a setting for AI education, our project demonstrates that hands-on, no-code tools can bridge that gap. Students not only trained seven distinct classifiers with impressively high accuracy but also developed a nuanced understanding of why some models generalize better than others. First, generalization performance, how well the model performed on new, unseen data, varied depending on the nature of the visual content. Projects involving more stable, static images (e.g., rocks, school objects) tended to generalize better compared to models working with more dynamic or nuanced inputs like facial expressions, which were more sensitive to changes in lighting and subtle differences in the data. Students who used consistent lighting and clean backgrounds achieved enhanced accuracy. For example, Group E (rock types) gathered data outdoors under natural lighting with static object positions, which helped ensure reliable model performance. On the contrary, students using variable camera angles or facial data (Groups A, F) encountered minor misclassifications due to differences in expressions, lighting, and occlusion. Student learning journals confirmed these findings:

- "*Initially, our model misclassified smiles and neutral faces, but after incorporating a greater variety of training images, its performance improved significantly.*"
- "*It was surprising that changing the lighting affected accuracy so much—it made me realize how sensitive AI can be.*"

5.1 Convergence and overfitting

Beyond generalization, we also examined the models' learning dynamics over training epochs. In Figure 4a, training and test accuracy reach 100% by epoch 16 and remain stable during the rest of the training process. This process indicates that the model learned to perfectly distinguish between the three classes (*Too early, Ripe, Too late*) based on the input data. Figure 4b shows a corresponding sharp drop in training and test loss, approaching near-zero values very early in the training process. The close alignment between the training and test curves in both graphs suggests that the model achieved excellent performance with no signs of overfitting. Although students were not directly involved in tuning training parameters or analysing performance metrics in depth, these visualizations introduced key machine learning concepts such as convergence, overfitting, and generalization in an age-appropriate and intuitive manner.

5.2 Pedagogical implications

The immediate visual feedback on accuracy and loss empowered students to connect abstract machine-learning concepts with tangible outcomes. By offloading coding syntax to the platform, learners could focus on experimental design, hypothesis testing, and iterative refinement, key elements of computational thinking. Small-group collaboration fostered peer learning and social constructivism: students compared strategies for data collection, debated trade-offs between dataset diversity and noise, and negotiated when to stop training. These interactions not only deepened technical fluency but also promoted metacognitive skills as students articulated their reasoning and evaluated model behavior.

5.3 Ethical and theoretical reflections

Integrating facial-image projects naturally prompted discussions about data privacy, consent, and the permanence of digital records, issues aligned with global frameworks such as AI4K12. Situating these debates alongside classical learning theories created a rich interdisciplinary context: trial-and-error optimization in models mirrored Thorndike's puzzle-box experiments, while rapid convergence evoked Köhler's insight learning by illustrating how structured feedback can trigger conceptual leaps. Encouraging students to draw these parallels fostered a holistic understanding of learning as both a human and algorithmic process, strengthening their ethical awareness and critical stance toward emerging technologies. Together, these findings suggest that no-code AI experiences can do more than teach technical skills; they can cultivate collaborative inquiry, ethical sensitivity, and a theory-informed appreciation of learning mechanisms.

5.4 Limitations of the study

While the results of this study are encouraging, there are several limitations to consider. Firstly, the sample size was small, consisting only of four students enrolled in a regional center for gifted education. Consequently, the findings may not be relevant to broader student populations with diverse academic backgrounds and abilities. Secondly, the study relied exclusively on the TM platform, which, although user-friendly, offers limited control over model architecture, hyperparameters, and performance metrics beyond fundamental accuracy. More advanced insights, such as confusion matrices or loss curves, were unavailable through the interface. Additionally, the models were tested using small, self-collected datasets, which may have introduced bias or overfitting despite the students' efforts to vary lighting and background conditions. Finally, the short duration of the project limited the depth of exploration, and long-term retention or understanding of AI concepts was not formally assessed.

6. Conclusion and future work

This study demonstrates that elementary level students, even without prior coding experience, can effectively engage with core machine learning concepts through hands on platforms that require no code. Over five sessions, participants designed and trained seven different image classification models, covering tasks such as facial expressions, recyclable materials and fruit ripeness, with reported accuracies between 84 % and 96 %. Beyond achieving strong performance metrics, students developed critical awareness of issues such as dataset bias, overfitting and ethical considerations around personal data use. By removing technical barriers and emphasizing collaborative experimentation, the activity fostered foundational digital literacies including iterative problem solving, reflective analysis and informed dialogue. Students questioned how model inputs, training conditions and real world contexts influence outcomes and thereby laid the groundwork for responsible use of artificial intelligence. Future work should broaden participation to include more diverse learner populations and implement formal assessments to quantify gains in both conceptual understanding and ethical reasoning. Comparing different hands on environments, such as Scratch + Machine Learning for Kids, MIT App Inventor with AI extensions, and beginner-friendly

Python notebooks, could provide differentiated learning pathways tailored to student interests and experience levels. Incorporating richer analytics, including confusion matrices, parameter sensitivity studies and structured ethical case critiques, would deepen technical insight while maintaining accessibility. Longitudinal studies are also needed to evaluate how early AI exposure shapes students' sustained interest in STEM fields, their digital competence and their development as informed digital citizens. Embedding accessible, project based AI experiences at the elementary level can equip the next generation with both the skills and the critical mindset needed to navigate and contribute to an increasingly intelligent world.

Acknowledgments

The authors express profound gratitude to the Center for Talented Youth Kragujevac for organizing and supporting this project. We are thankful to the gifted students who participated with enthusiasm, creativity, and curiosity, making this learning experience meaningful and inspiring. We extend our appreciation to the partnering schools and mentors whose guidance and collaboration contributed significantly to the success of the activities. We would especially like to acknowledge the outstanding contributions of the participating students: Staša Milićević, Đurđa Đurđević, Elena Novović, and Luka Vasović, whose dedication, originality, and passion for learning were truly remarkable. These students, from Sveti Sava Elementary School and Mirko Jovanović Elementary School, exemplify the spirit of inquiry and excellence that this project aimed to foster. It was a privilege to work with such bright young minds, and we are confident that their talents will continue to shine in the future.

References

Alowais, S. A., & Alowais, A. A. (2023). Revolutionizing healthcare: The role of artificial intelligence in clinical practice. *BMC Medical Education, 23*(1), 469. https://doi.org/10.1186/s12909-023-04698-z

Badža, S. R. (2024). Development of artificial intelligence in Serbia: Serbia as the regional leader. *Napredak*, *5*. https://doi.org/10.5937/napredak5-52537

Chen, L., Chen, P. & Lin, Z. (2020). Artificial intelligence in education: A review. *IEEE Access, 8*, 75264–75278. https://doi.org/10.1109/ACCESS.2020.2988510

Google. nd. Teachable Machine. Retrieved March 8, 2025, from https://teachablemachine.withgoogle.com

Gouseti, A., James, F., Fallin, L., & Burden, K. (2024). The ethics of using AI in K–12 education: A systematic literature review, *Technology, Pedagogy and Education, 33*(1), 161–182. https://doi.org/10.1080/1475939X.2024.2428601

Government of the Republic of Serbia, Ministry of Trade, Tourism and Telecommunications. *Strategy for the development of artificial intelligence 2020–2025.* Retrieved July 16, 2025, from https://ai.gov.rs

Kohler, W. (1925). The mentality of apes. Harcourt, Brace & World.

Kurz, T. L., Jayasuriya, S., Swisher, K., Mativo, J., Pidaparti, R., & Robinson, D. T. (2024). The impact of Teachable Machine on middle school teachers' perceptions of science lessons after professional development. *Education Sciences*, *14*(4), 417. https://doi.org/10.3390/educsci14040417

Lane, D. (2020). Machine Learning for Kids: A platform to teach AI using Scratch and Python. Retrieved March 8, 2025, from https://machinelearningforkids.co.uk

Luckin, R., Holmes, W., Griffiths, M. & Forcier, L. B (2016). Intelligence Unleashed: An Argument for AI in Education. Pearson.

Malahina, A. U., Saitakela, M., Bulan, S. J., Lamabelawa, M. I. J., & Belutowe, Y. S. (2024) Teachable Machine: Optimization of herbal plant image classification based on epoch value, batch size and learning rate. *Journal of Applied Data Sciences*, *5*(2), 532–545. https://doi.org/10.47738/jads.v5i2.206

Odabas, M. S., Şenyer, N., & Kaya, S. O. (2024). Classification of apple diseases and pests using the Google.com powered Teachable Machine. *Journal of Agricultural Faculty of Gaziosmanpaşa University*, *41*(2), 66–71. https://doi.org/10.55507/gopzfd.1287389

Omrod, J. E. (2024). Human learning (7th global ed.). Pearson.

Relmasira, T., & Vartiainen, H. (2025). Artificial intelligence literacy education in primary schools: A review. *International Journal of Technology and Design Education*. https://doi.org/10.1007/s10798-025-09979-w

Ruuska, T., & Salminen, J. E. (2020). Massive-scale online course (MOOC) to teach AI to the general public: Case Elements of AI. In 2020 *IEEE Frontiers in Education Conference (FIE)*. https://doi.org/10.1109/FIE44824.2020.9274208

Slavin, R. E. (2018). Educational psychology: Theory and practice (12th ed.). Pearson.

Thorndike, E. L. (1911). Animal intelligence: Experimental studies. Macmillan.

Topol, E. J. (2019). Deep medicine: How artificial intelligence can make healthcare human again. (2019). Envisioning AI for K–12: What should every child know about AI? *In Proceedings of the AAAI Conference on Artificial Intelligence*, *33*(01), 9795–9799. https://doi.org/10.1609/aaai.v33i01.33019795

Touretzky, D. S., Gardner-McCune, C., Martin, F., & Seehorn, D. Tran, T. A., Nguyen, T. H., Nguyen, T. H. N., & Tran, H. A. (2024). Detecting dental caries through captured images using the machine learning technology Teachable Machine. *Asian Journal of Health Research*, *3*(1), 20–23. https://doi.org/10.55561/ajhr.v3i1.148

Vuorikari, R., Punie, Y., Carretero Gomez, S. & Van Den Brande, G. (2016). DigComp 2.0: The Digital Competence Framework for Citizens. Update Phase 1: the Conceptual Reference Model. EUR 27948 EN. Luxembourg (Luxembourg): *Publications Office of the European Union*; JRC101254. https://dx.doi.org/10.2791/607218

Zawacki-Richter, O., Marin, V. I., Bond, M., & Gouverneur, F. (2019) Systematic review of research on artificial intelligence applications in higher education – where are the educators? *International Journal of Educational Technology in Higher Education*, *16*(1), 39. https://doi.org/10.1186/s41239-019-0171-0