Exploration Project 1: Teachable Machine Classifier

Habitat-based Wildlife Classifier

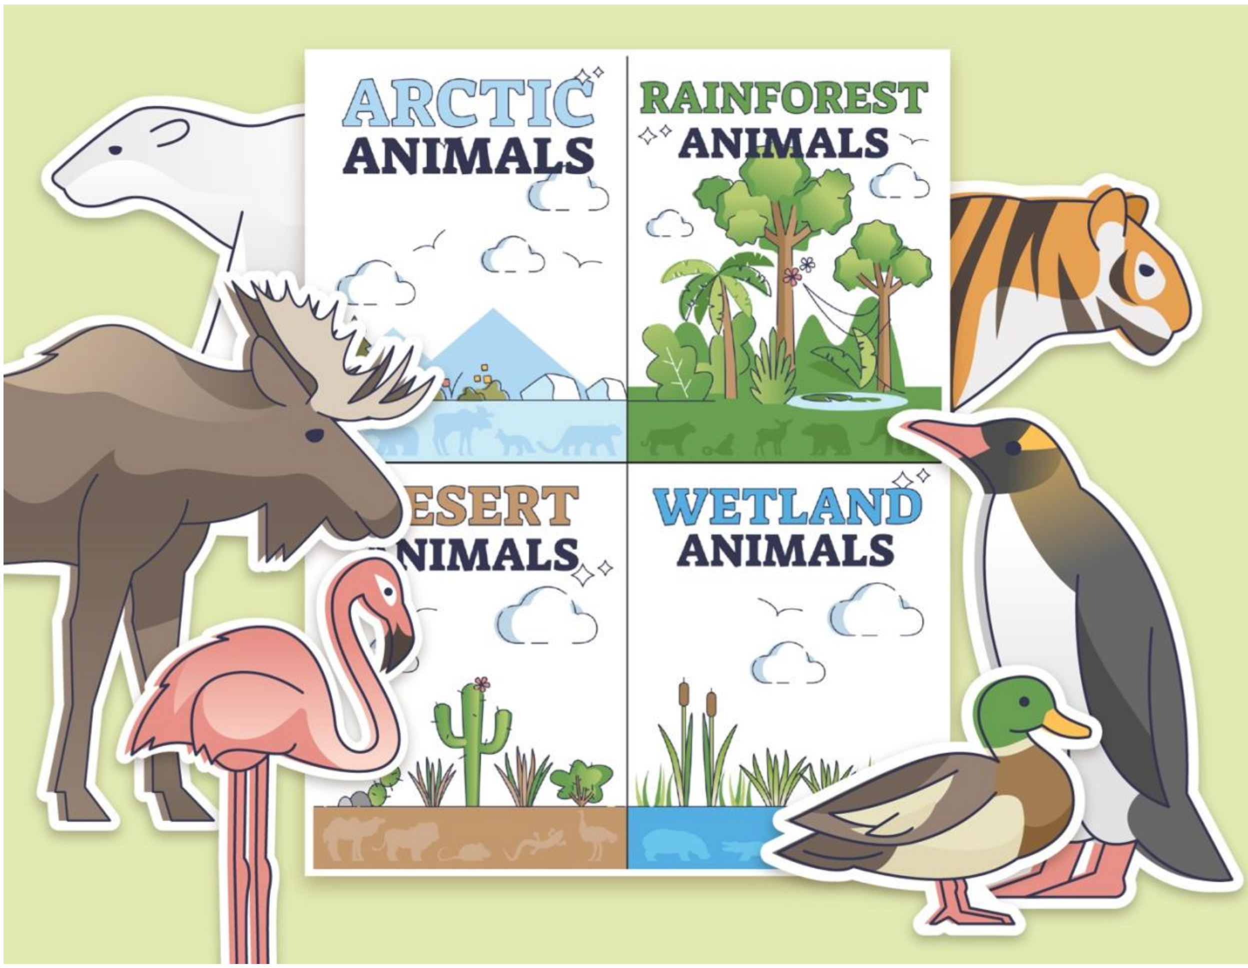

Submitted by Jyoti Poonia (M.S.Ed. LST)

EDUC 5911 Artificial Intelligence for Children and Youth: Learning, Creating, and Understanding

Spring 2025

1. Project description

For this project, I designed a Teachable Machine classifier that categorizes animals based on their natural habitats. Using Illustrator, I repurposed a set of stylized animal illustrations, sourced from Adobe Stock, sorted into four habitat categories—Desert, Arctic, Forest, and Wetlands. The goal of this experiment was to train an AI model (Teachable machines) to recognize an animal's habitat based on visual characteristics such as color, shape, and environmental adaptations. This project is designed for children aged 5-10 years and can be integrated into wildlife education, conservation awareness programs, and interactive museum exhibits. This project explores Big Idea 1(Perception) and Big Idea 2 (Representation and Reasoning) of the AI4K12 initiative (Touretzky et al., 2023).

2. Process

I selected five animals or birds for each habitat, ensuring that they were visually distinct enough from each other, for clear differentiation (see *Figure 1*). I manipulated the illustrations carefully to generate the required graphics for the project. The graphics needed to be both functionally useful for training the model and visually engaging to simulate real-world art-based learning experience.

Figure 1: Habitat categories and respective animals with key features of the images selected

Habitat	Animals	Image Set Features
Arctic	Snow Leopard, Polar Bear, Moose, Penguin, Arctic Fox	White/grey backgrounds, minimal vegetation and elements
Desert	Camel, Desert Rat, Lizard, Lion, Ostrich	Sandy/brown tones, minimal vegetation and elements
Forest	Deer, Tiger, Elephant, Monkey, Brown bear	Green background, trees, busy background
Wetland	Flamingo, Duck, Hippopotamus, Frog, Crocodile	Blue/brown tones, minimal to busy background

To build the dataset for each habitat, I collected real-life images of the illustrated animals from Adobe Stock, compiling approximately eight images per animal. This resulted in about 40 images per habitat and a total of 160 images for the project (see *Figure 2*).

Figure 2: Image sets for teaching the model

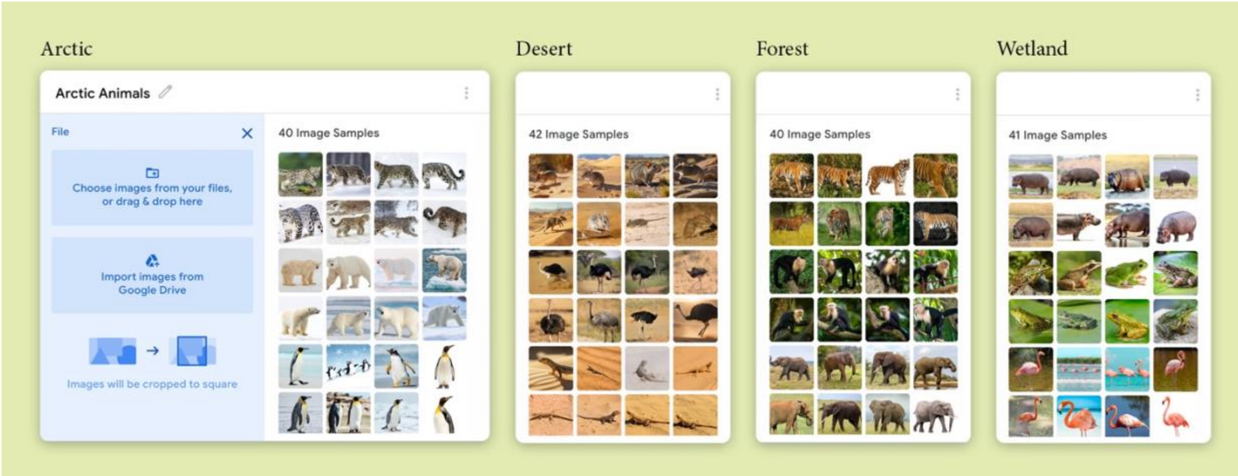

For the first phase of model training, I trained the classifier using only real images. This provided an initial baseline, for the classification, which worked well for some animals (we shall return to this in the reflection section later) and failed for others. For the second phase, I sought to improve accuracy by incorporating the illustrated images into the dataset. This hybrid approach of blending real-world images with artistic representations helped increase accuracy for most of the animals (see *Figure 3*).

Figure 3: Phase 1 and Phase 2 of teaching the model

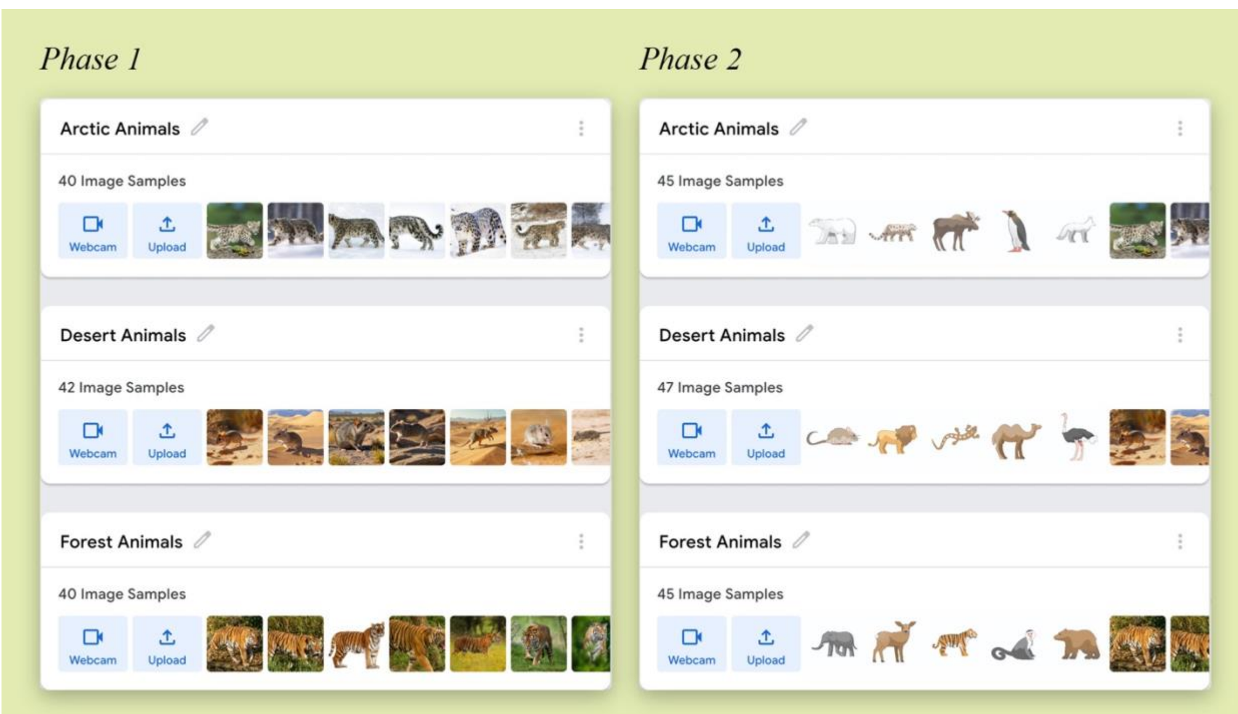

I developed a low-fidelity prototype of how this project can be implemented in an interactive setting. This prototype serves as a proof of concept, demonstrating how the classifier can be integrated into educational tools for its various educational applications (see *Figure 4*).

Figure 4: Low-fidelity prototype

3. Reflection on the learning

This project provided valuable insights into both the technical and conceptual aspects of AI-based classification, particularly in the context of an art-based learning application. I reflect on two broad areas: considerations during data selection and observations from results.

3.1 Considerations during data selection

The data collection process involved several challenges that required thoughtful decision-making.

i. ***Real life images vs. illustrations:*** The intention was to understand limitations and scalability of art-based applications and get insights into the kind of data sets required for such models. Initially, I used real-life images of animals, to check how successfully they would classify the illustrations, ensuring that the images were as close representations as possible of their stylized versions. However, I realized that inconsistencies in background, lighting, and other visual elements introduced bias and affected the classifier's accuracy. I added illustrations to the data set for phase 2 of training, to ensure a more controlled and consistent dataset while also aligning with the project's focus on an art-based learning approach.

ii. ***Scope of the data set:*** Since the goal was to create an interactive experience for children, the model was trained only for the selected species. I selected visually distinct animals or birds for each habitat to avoid overlap in classification. I included some pairs such as a polar bear and a brown bear, or a moose and a deer, which might pose challenges due to their similar features, to see how the model performs. I ensured that the secondary elements (background—white vs green) for such pairs were distinct enough. Additionally, I considered whether to include plant life, such as cacti or trees, in the dataset. However, I ultimately decided to focus solely on animals, as introducing additional elements could have complicated both data collection and classification.

iii. ***Image formatting:*** All the images were either cropped to squares or the animals were centered consistently to help in accurate representation. This tells me how such models have their pre-defined guidelines for data that we should adhere to. (see *Figure 5*)

Figure 5: Image specifications

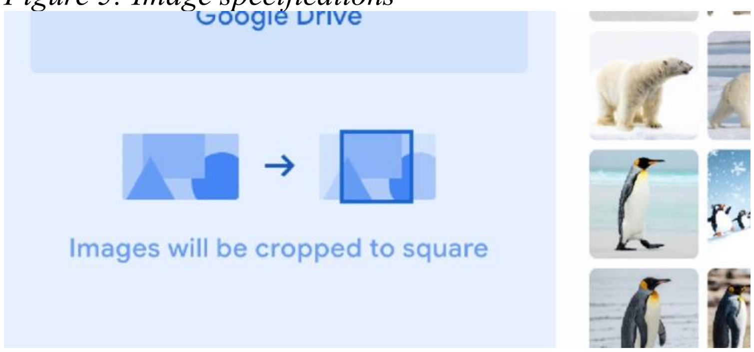

3.2 Observations from results

The results of the classification process revealed interesting patterns. Some animals were identified with near-perfect accuracy from the first round of training, regardless of the background. This suggests that their visual features were distinct enough for the model to recognize them reliably (see *Figure 6*).

Figure 6: Animals with high accuracy after phase 1 of training

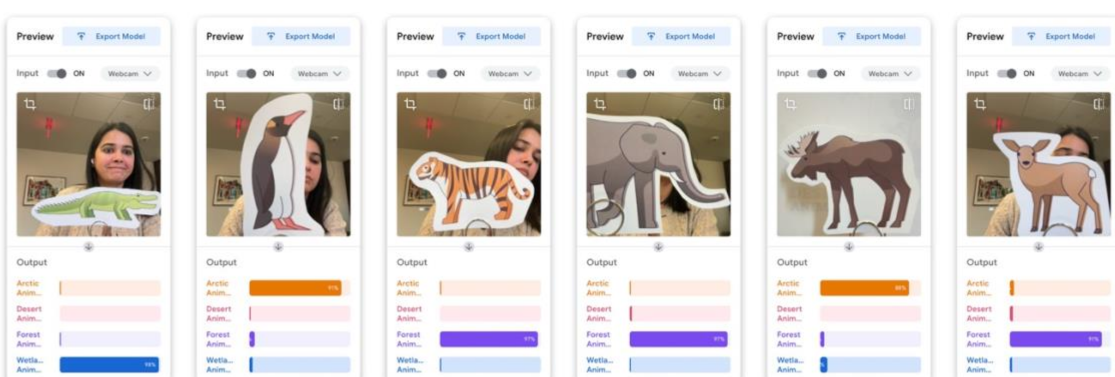

Others had fluctuating accuracy, around 70%, depending on factors such as distance from the screen, lighting conditions, and whether a white background was placed behind the image. This variability indicated that the model was sensitive to environmental factors, which could be mitigated by further refining the dataset (see *Figures 7 and 8*).

Figure 7: Animals with passable accuracy after phase 1 of training

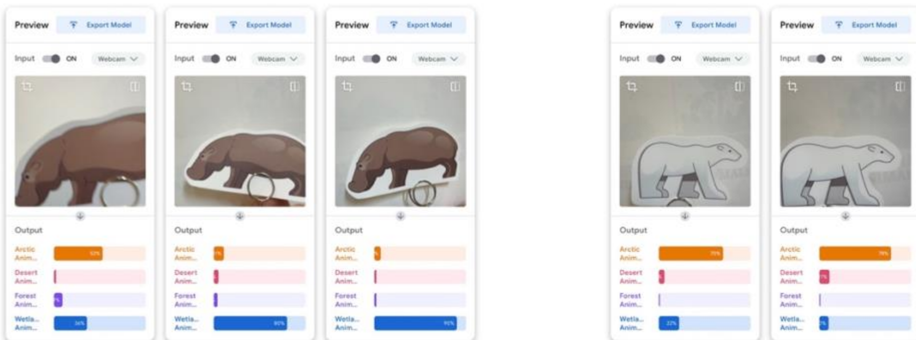

Figure 8: Animals with fluctuating accuracy depending on the background

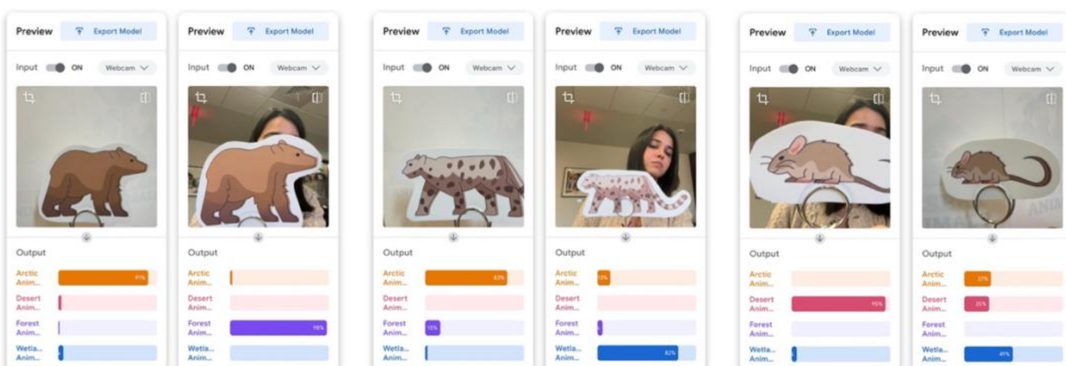

For some animals, accuracy significantly improved after a second round of training, suggesting that the model benefited from additional exposure to the same set of images (see *Figure 9*).

Figure 9: Animals with improved accuracy after phase 2 of training

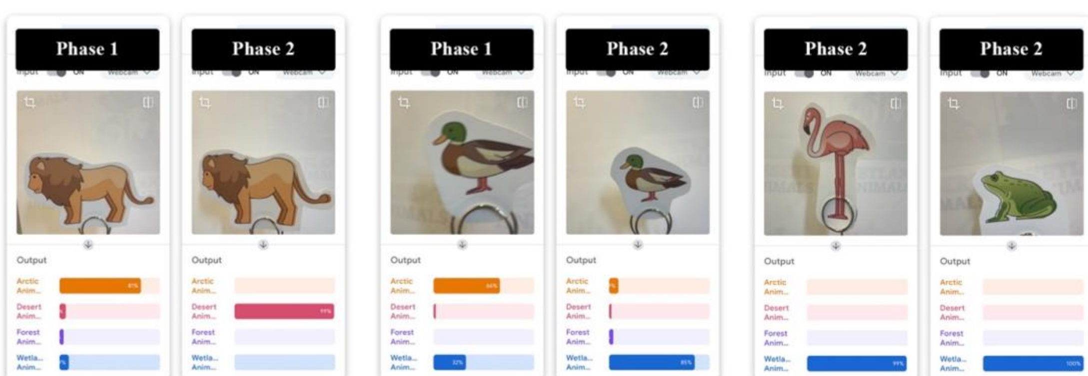

However, certain animals failed to match their intended habitat even after retraining, challenging my assumption that including he same images would result in 100% accuracy (see *Figure 10 and 11*). Further exploration is needed for more insights into how to optimize the dataset.

Figure 10: Animals that failed to be classified after phase 2 of training

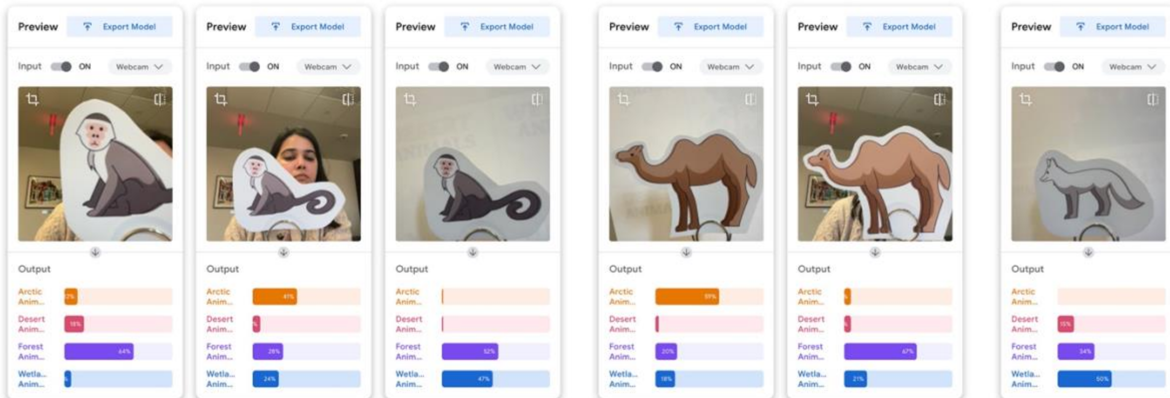

Figure 11: Testing the ostrich after phase 2 of training

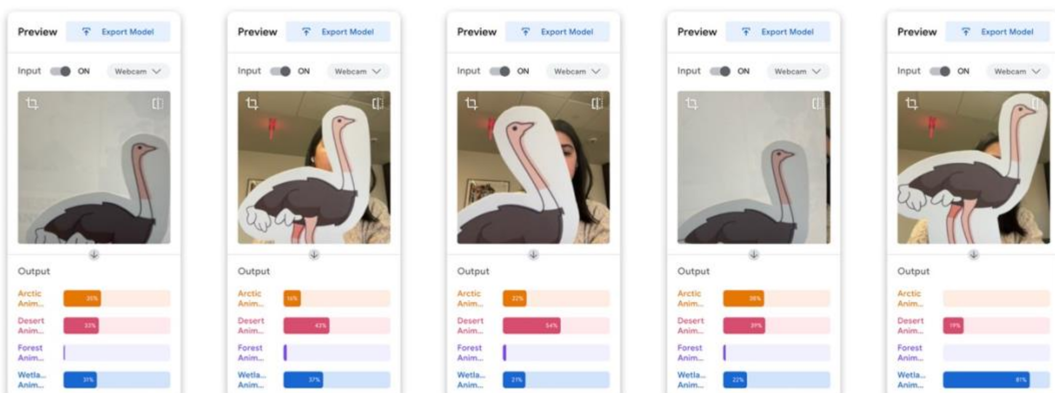

One of my key takeaways is that a larger and more diverse dataset would have likely improved classification accuracy across all categories. Overall, this project reinforced my understanding of how AI models interpret visual data and the impact of data selection on model performance. It also highlighted the importance of balancing technical accuracy with practical usability in an educational setting. Moving forward, I aim to refine the dataset further, incorporate additional training rounds, and explore ways to make the classification model more robust and adaptable.

4. Prototype for a Real-world application

Figure 12 shows the low-fidelity small scale prototype of one possible application. Here, children would match animals with their respective habitats using teachable machines. The habitat sheet also has silhouettes of the animals to scaffold them through the process and in case of inaccuracies in the model. While this project is a 'black box' approach (Long & Magerko,

2020) and falls short of truly being constructivist learning (Kahn & Winters, 2021), it seeks to bridge the gap by making AI exploration into a tangible, hands-on experience.

Figure 12: Prototype

References

Kahn, K., & Winters, N. (2021). Constructionism and AI: A history and possible futures. British Journal of Educational Technology, 52(3), 1130-1142.

Long, D., & Magerko, B. (2020, April). What is AI literacy? Competencies and design considerations. In Proceedings of the 2020 CHI conference on human factors in computing systems (pp. 1-16).

Touretzky, D., Gardner-McCune, C., & Seehorn, D. (2023). Machine learning and the five big ideas in AI. International Journal of Artificial Intelligence in Education, 33(2), 233-266.