

AI Chef Trainer: Introducing Students to the Importance of Data in Machine Learning

Saniya Vahedian Movahed and Fred Martin

University of Texas at San Antonio
saniya.vahedianmovahed@utsa.edu, fred.martin@utsa.edu

Abstract

We developed AI Chef Trainer, an educational web app that introduces children to the role of data in machine learning (ML) through the engaging task of recipe recommendation. We tested our software with middle school students. The results indicated that the students recognized the importance of both data quantity and specificity in the training process. Initially, students tested the AI Chef’s capabilities by selecting from a list of ingredients to see what the system recommended as possible recipes. After observing the recommendations, they contributed by adding their own recipes—each being a set of ingredients and a corresponding recipe-name—which were used to retrain the model and finally re-tested recipe suggestions. This cyclical process of testing, contributing, retraining, and post-training testing provided students with hands-on experience in how AI systems learn and adapt over time based on new data. Forty-five of 52 students entered recipes, and 26 of the 52 tested their own recipes using the specific ingredients they entered. Students were introduced to the concept of confidence percentages via the AI recipe suggestions. Even as the primary focus was on the data’s role in ML, the AI Chef Trainer software also served as a window into students’ cultural expression and personal preferences.

Live Demo — tinyurl.com/AI-Chef

Source Code — github.com/saniavn/AI-Chef-Trainer

Introduction

As artificial intelligence (AI) and machine learning (ML) technologies become more prevalent, it is imperative to equip novices, particularly children, with the skills needed to interact with, navigate, and understand these systems. Cultivating AI literacy—understanding, using, and ethically engaging with AI technologies—has shifted from an advantage to a fundamental necessity for success and informed participation in an AI-augmented world (Druga et al. 2019; Long and Magerko 2020). Researchers call for national AI education centered around Five Big Ideas to prepare children as informed digital citizens (Touretzky et al. 2019).

Machine learning models are heavily reliant on diverse and representative data to perform robustly and generalize effectively (Sambasivan et al. 2021). Recognizing the role of

data in AI/ML technologies, it is crucial to educate students about the impact of data on system effectiveness and biases. This necessity highlights the importance of integrating data literacy into the K-12 curriculum to emphasize data’s significant impact on these technologies (Grover 2024).

Researchers and developers have created tools to introduce students to AI and machine learning. A leading example is Google Teachable Machine, which allows students to create their own training datasets, demonstrating how crucial training data is in shaping AI behavior (Carney et al. 2020). The Machine Learning for Kids platform allows students to train a model and create Scratch programs based on the model’s behavior (Lane 2024). AI for Oceans introduces children to machine learning by having them train classifiers to identify marine life categories. It shows how data impact the accuracy and fairness of AI systems (Code.org 2020). ChemAIstry teaches youth machine learning by letting them train a decision tree classifier to identify safe and unsafe items in a chemistry lab, illustrating the direct impact of data on AI outcomes (Martin et al. 2024).

In further work, youth refine machine learning models to recognize patterns in physical activities, enhancing their understanding of ML’s iterative nature and broader applications (Zimmermann-Niefield et al. 2019). Zhorai, a conversational agent, allows children actively teach the agent about ecosystem topics, and observe how their inputs affect the agent’s learning process (Lin et al. 2020). The Pasta Land activity uses decision trees with the Palmer Penguin dataset for k-nearest neighbors to teach middle school students about machine learning algorithms (Ma et al. 2023).

The present study introduces a novel way to engage children in understanding the impact of training data in ML. Recipe recommendation was deliberately selected to provide a relatable, personalized, and data-centric activity. Our research questions (RQs) are:

1. How can we develop student data literacy skills by allowing students to explore, experiment, and learn from their interactions with an AI model?
2. How can we demonstrate the importance of data in AI by integrating student-generated recipes into an existing dataset?
3. To what extent will student-generated recipes encourage personal expression?

System Design and Implementation

The AI Chef Trainer uses a Naive Bayes classifier to suggest recipes based on student-selected ingredients. Initially, ingredients are treated as free text. They are then analyzed using Term Frequency-Inverse Document Frequency (TF-IDF), which captures their significance across recipes and transforms them into a numeric format that reflects both the frequency and uniqueness of ingredients and serves as the input for the classifier. This transformation allows the ingredients to be treated as discrete variables. The Naive Bayes classifier then calculates each recipe's likelihood by considering the presence of each ingredient as an independent feature. This independence assumption simplifies computations, making Naive Bayes efficient for handling discrete variables in text classification (Wickramasinghe and Kalutara 2021). After obtaining the raw probabilities, these probabilities are normalized so that they sum to 100%.

User Interface Details

The front-end of the AI Chef Trainer is structured across four distinct screens (Figure 1), each designed to facilitate different interactions with the system:

- **Student Information:** On the initial screen, students are prompted to enter their personal information, including their name, grade, and age range. This is for research purposes as discussed in the following section.
- **Pre-Training Test:** This screen has two panels. The left-side panel lets students select ingredients from a checklist or enter new ones, with a submission button below. The right-side panel displays AI-generated recipe recommendations, showing to students the immediate AI response to their inputs.
- **Training (recipe addition and model retraining):** On this screen, students contribute directly to the AI's learning by adding their recipes to the base dataset. They enter ingredients, specify a recipe name, and click the "Add a Recipe" button. After inputting their data, they click the "Train" button to retrain the AI model with the new data.
- **Post-Training Test:** The fourth screen mirrors the second screen in layout and function but is used to retest the AI's performance (recipe recommendations) after it has been re-trained with new data. This allows students to directly observe and evaluate changes in the AI's suggestions, providing a vivid demonstration of the interplay between data input and machine learning.

For each test, the AI suggests three recipe predictions with the highest corresponding probabilities, showing students the likelihood of each based on their ingredients.

System Architecture

The AI Chef Trainer has a Flask-based back-end and a dynamic front-end designed with HTML, CSS, and JavaScript. This combination provides a responsive and interactive platform. The Flask application supports multiple interactive endpoints. These endpoints manage student inputs, model retraining, and data storage, ensuring a secure and responsive experience for each student.



Figure 1: AI Chef Trainer's four screens

grade	number	percent
6th	17	33%
7th	23	44%
8th	12	23%
total	52	100%

Table 1: Student participation by grade

To support concurrent student interactions, each student’s data is managed independently. Specifically, a copy of the base recipe dataset is created for each student, and model retraining is performed independently for each student’s dataset. This approach ensures that individual student data does not accumulate in the base dataset so that each student has the same experience.

The AI Chef Trainer is equipped with a deliberately limited dataset of 79 recipes to demonstrate the impact of data in AI performance. After initial tests, students can improve the system by contributing new recipes and retraining the model via the “Train” button. Then, they retest the system to observe the influence of additional data. The system resets to its initial state after each student’s use.

The research experiment with the software was run locally on laptops; a web-hosted version is now available. To preserve data privacy, after a student exits the system, their interaction log data (JSON format) is downloaded to their local machine and their data file is deleted from the server.

Study Design

In spring 2024, AI Chef Trainer and several other projects were developed as part of a research course. These were tested by middle school students at a STEM public charter school in a major Texas city. The school had 665 students in grades 6–8 (aged 10–14 years). Of these, parents of 222 students consented to their children’s participation and 124 students participated in one of the two afterschool program sessions.

Upon arrival, after providing assent, students filled out a pre-interaction survey. Then, they selected stations to engage with, rotating every 8–10 minutes to a separate project, ensuring exposure to multiple projects.

Time constraints prevented all students from engaging with every project. 63 students ultimately engaged with AI Chef Trainer at one of five or six concurrent workstations. Screen recordings, audio recordings of student conversations, and interaction log data with the software were collected as part of the study. IRB approval was obtained for the study.

The pre-interaction survey was designed to capture students’ attitudes towards AI and their familiarity with specific concepts. The data revealed that students predominantly expressed strong agreement with AI’s potential to positively influence daily life and enhance career opportunities across various fields. The enthusiastic response illustrates their optimism about AI’s potential. Also, most of the students were already familiar with the concepts of probability, however, they were not aware that the AI/ML responses might be associated with probability.

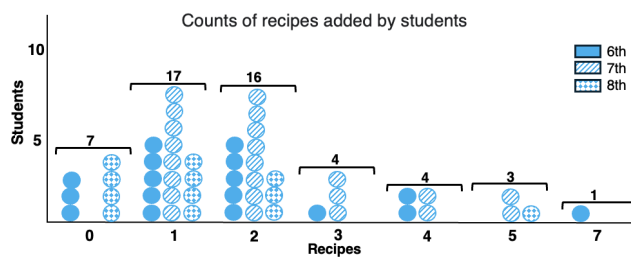


Figure 2: Histogram of students submitting recipes by count

Results

In this section, we present findings derived from analysis of student interaction log data, audio and screen recording, and post-survey data from the middle school study. While 63 students interacted with the software, we considered only 52 of those students for the analysis. We excluded 11 students because they did not complete the post-survey interaction or there were no corresponding log data files. The absence of data suggests that these students likely teamed up with others and interacted with the tool collaboratively, rather than individually. Also, it should be noted that the students did not receive a formal introduction to any AI concepts as part of the study. Table 1 shows student participation by grade.

Student Engagement

We tracked student interaction with AI Chef Trainer using log data to measure their engagement. Students participation included adding recipes and ingredients. This involvement reflects their engagement in the learning process.

Figure 2 shows a histogram of the number of students contributing recipe counts from 0 to 7 recipes submitted. Nearly all students (45 of 52, 87%) submitted at least one recipe. 28 students submitted two or more recipes. One student submitted seven recipes. Overall, this shows good engagement with AI Chef Trainer and understanding of its intended use.

Figure 3 shows the number of tests conducted by students before and after training the model. The majority of students conducted one or two preliminary tests. The post-training testing shows a visible increase in students conducting multiple tests; more students conducted four tests, and one highly engaged student conducted 13 tests.

Understanding AI’s Mechanism

This section explores how active engagement with the system—through experimentation with data and personal recipe contributions—deepened students’ insights into the influence of their inputs on AI performance and accuracy.

Student Experiments and AI Response Patterns

Through an analysis of interaction log data, we identified distinct patterns in how students engaged with the AI’s response mechanisms. Of the 52 students participating in the test, 26 (50%) chose to experiment by testing the AI with subsets of ingredients used during their initial training sessions, or by employing the complete set of ingredients from specific recipes. This investigative approach was evident across different grades.

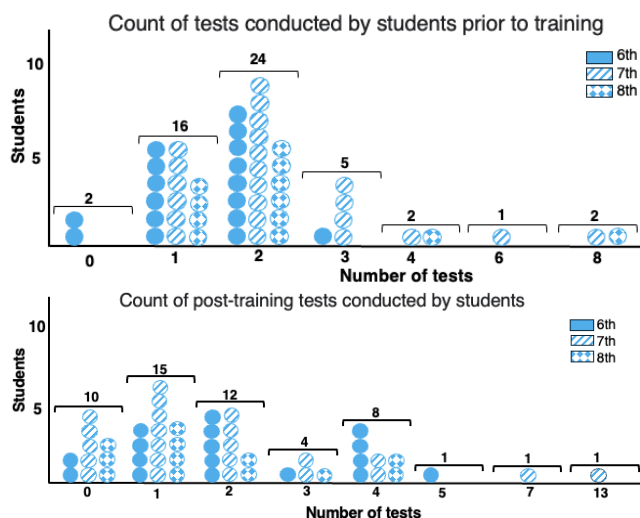


Figure 3: Histogram of number of tests conducted by students **before** and **after** adding their own recipes.

Figure 4 on the next page dives deeper into students' post-training test data behavior, showing a per-student view of how many ingredients were used in each student's post-training tests. For each student, it shows a box plot illustrating the distribution of the count of ingredients they used in their post-training test sessions. Arrows highlight the students (26 of 52) who tested the AI Chef Trainer with their own added recipes. The diagram reveals that 42 of the students did the post-training test at least once. A star marks a student who performed two tests, each with three ingredients, indicating no variation in their average ingredient count. Students represented by a single dot, instead of a box, either conducted only one post-training test or did not conduct a post-training test.

Some students focused on using particular ingredients to probe the AI's ability to process and respond to unique data inputs. This deliberate selection of ingredients reflects a sophisticated investigative approach, revealing students' intent to understand how variations in input data could affect the AI's predictive capabilities.

Here are two examples of students specifically testing their own recipes. From the individual student-testing diagram (Figure 4), Student 3 defined this recipe:

nachos: cheese bell_peppers onion beans corn chips

and then tested with the same ingredients, yielding the recommendations:

nachos (46.71%), Nachos (26.95%), stuff Bell Peppers (26.33%)

The student's recipe—the lower-case *nachos*—was recommended with the highest probability (our system was case-sensitive).

Student 28 defined the recipe:

Ice-cream: cheese tomato eggplant oil bell_peppers onion tortilla rice beef carrot milk potato meat lettuce beans cheese bell_peppers onion beans corn chips

There are repeated ingredients in this recipe (cheese, beans). The software allowed adding ingredients by checkbox and open text; it did not check for repeated ingredients.

Then the student tested with the ingredients:

cheese tomato potato oil flour bell_peppers onion carrot salmon chicken mushroom beef meat lemon olive oil eggs dough sauce salsa milk

which yielded the recommendations:

Ice-cream (37.63), Burger (32.32), Tacos (30.05)

The student was amused that AI Chef accepted their awful recipe for ice cream as being the definition for it.

Learning and Adaption Using data from post-survey interactions, we reflect on how students perceived the AI's ability to incorporate and learn from their input. Figure 5 shows the student distribution responses to the prompt "Do you think the software learned from the ingredients you picked?" As shown, most student responses leaned towards "agree" to "strongly agree" regarding the AI Chef's ability to learn from the ingredients.

From students' responses to the open-ended prompt, "Do you believe that providing more ingredients leads to more accurate recipe suggestions or not? Why?" we identified 18 relevant responses. Many emphasized the importance of providing AI with data: "it learns from the new data provided"; "I learned you have to train an AI by feeding it data for it to learn"; and "I learned that AI doesn't know everything yet and needs to learn from trial and error."

Others focused on data quality: "If you give more information, the ai will predict better"; "I think it would make it more accurate since it learns more things"; and "Yes, it gives the AI a better chance to learn and to receive what i'm thinking."

Data Quantity Regarding students' views on the relationship between the volume/quantity of ingredients provided and the accuracy of the AI's recipe suggestions, we identified responses that reflect on this theme from open-ended questions (totaling 23 responses). Examples of student responses to the prompt "Do you think the software learned from the ingredients you picked?" included "because more ingredients means more data"; "yes, because it has more data" (2 times); "the ai has more to work with"; "More ingredients allowed for a more accurate recipe because there was more training data"; "it teaches it more information"; and "Adding more recipe suggestions leads to more accurate recipe suggestions." These quotes suggest that students perceived a direct correlation between the quantity of ingredients provided and the accuracy of the AI's recipe suggestions.

Data Specificity Through analysis of the open-ended question, we observed that responses from 11 students show an understanding that specific ingredients enhance the AI's

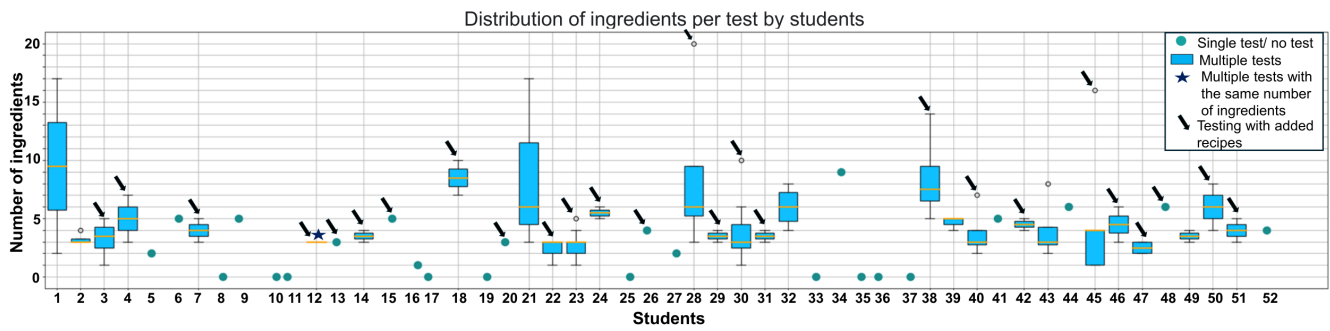


Figure 4: Per-student view of how many ingredients were used in each student’s post-training tests. Box plot shows the average number ingredients used by each student.

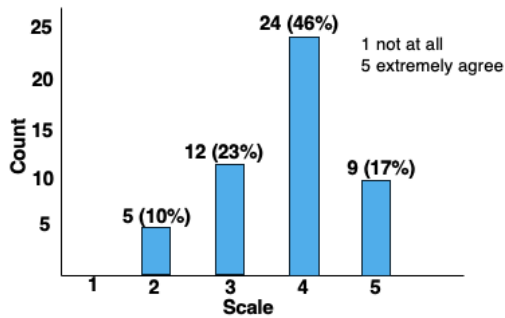


Figure 5: Student responses to “Do you think the software learned from the ingredients you picked?” Two students did not respond.

predictive accuracy. Examples included: “because more detail gives you more knowledge on things”; “Providing more ingredients allows the prediction to be based off of a more specific description”; “I think adding more ingredients creates more of a specific recipe for a recipe. If the recipe is more specific that gives it more of a better description of the recipe and a more accurate recipe because it limits the possibility of other recipes having the same ingredients”; “Ai looks up recipes that have those exact ingredients and that makes it more specific”; “I think adding more ingredients does lead to a more accurate recipe because the more specific the ingredients are the more it can be one recipe than another. recipes with only a few ingredients will match with more things”; and “because some ingredients can narrow down guesses because they are specific to some recipes.”

Additionally to quantify the qualitative insights, we analyzed student responses to a multiple option question, “What do you think it’s important for AI to learn from?” as shown in Figure 6. Most students agreed on the importance of AI learning from a variety of dishes and ingredients.

Understanding AI’s Decision-Making Process To further assess students’ understanding of how data influences on AI behavior, we asked “If the AI often suggests certain types of recipes over others, what do you think could be the reason?” Responses are shown in Figure 7. Students

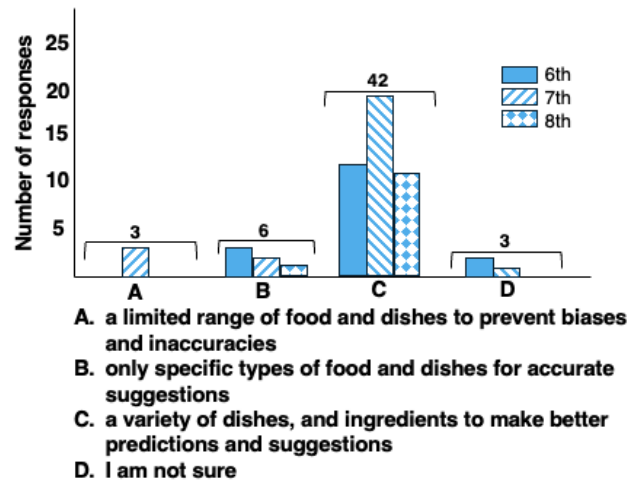


Figure 6: Student responses to the prompt “What do you think it’s important for AI to learn from?” One student selected 3 options (A, B, and C).

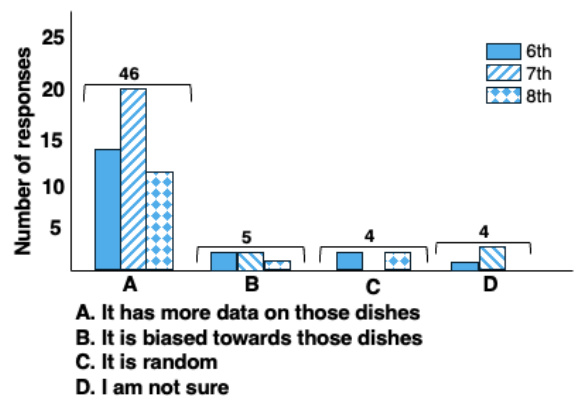


Figure 7: Student responses to prompt “If the AI often suggests certain types of dishes over others, what do you think could be the reason?” Two options were selected by two students and one student did not respond.

Ingredient	Recipe
salt (9), pepper, spice mix, jalapenos, cucumber, ured dahl, basil (3), kale, banana, seaweed, chocolate chips (7), avocado (3), gee, lentils, salsa (2), chives, ranch, tator, guacamole, apple, creamer, waffle, strawberry, cream (2), dough (5), sauce, meatballs, noodles (3), spices, corn chips (2), ice (3), cabbage, tomato sauce, cannoli shells, powdered sugar, vanilla extract, garlic, peas, vingar, bread (3), baking soda, marshmallows, water (2), baisl, mango, pepaers, curry spice, tumeric (2), graham cracker, pepper, panchphoran, sprinkles, sugar (3), wipcream, vanila, ketchup, jeera, chili powder, olives, panchphoron (2), kalo jeera, salsa, cilantro (2), salsa verde, salmon	fish taco, burger, salad (3), taco (5), brownie, orange, dried-lentils, chichadee, ramen, milkshake, S'mores, rice, Ceaser Salad, chicken rice, guacamole, pongal, sandwich, chicken sandwich, nachos (2) noodles, sushi, Bertie, aloor dam, begun bhaji, biryani, enchilada, Carne Asada, stew, curry chicken

Table 2: Ingredients and recipes contributed by students

Note: Student misspellings are preserved

expected that the AI would be more likely to recommend recipes for which it had more data. Students recognized that AI effectiveness hinges on data availability, showing a solid understanding of data dependency in AI outcomes.

A smaller group highlighted potential biases, indicating awareness of data-driven prejudices: “AI has bias”; “AI is able to match things without 100% accuracy and is able to acknowledge that inaccuracy”; and “AI is very predictive.”

In responding to the prompt “Do you believe that providing more ingredients leads to more accurate dish suggestions or not? Why?” sample student responses included “No, because it could change the taste to bad or good”; “no because there are more possibilities”; and “no, because there are so many recipes with the same ingredient.”

Personalization and Cultural Relevance

This section explores how students contributed uniquely. Even a single, well-detailed recipe with unique ingredients can provide insights into student backgrounds. While a pre-defined list of ingredients is provided in the tool, Table 2 presents the ingredients and recipes that students chose to add beyond the given options, alongside the frequency of each entry, as written by the students. This list is divided into two columns: the first column lists the ingredients with the number of times they were input by students (noted in parentheses), and the second column details the recipes with corresponding frequencies. This frequency data serves as a quantitative measure of each item’s popularity and recurrence within the submitted data.

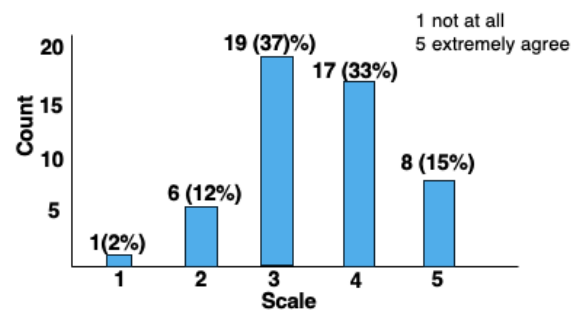


Figure 8: Student responses to the prompt “Did the AI Chef Trainer’s prediction match your expectation?” One student did not respond.

Upon analyzing these entries, the prominence of Mexican and Indian cuisines became evident. This illustrates the deep influence of cultural identity on data inputs. This highlights the dynamic interplay between cultural identity and technological engagement. Examples of such recipes are “Cheese, oil, onion, tortilla, rice, lettuce, Salmon, Avocado, Cabbage” to make “fish taco” and “oil, onion, potato, turmeric, panchphoran, jeera, chili powder” to make “aloor dam.”

In another example, a 7th grader said “I teach AI recipe” and “ai knows my secret family recipes.” His recipe was “eggplant, tortilla, rice, potato, meat” to make “taco.” This both shows the AI’s capability to incorporate and learn from diverse cultural inputs and also reflects the student’s active role in teaching the AI about personal and family culinary traditions. This student also expressed an intention to “mess up with AI” by deliberately inputting unusual ingredient combinations (e.g., meat, beef) for the ice cream recipe.

Student Overall Experience

We also asked, “Did the AI Chef Trainer’s prediction match your expectation?” Responses indicated a generally positive perception of the AI’s performance. Most students rated their satisfaction in the upper range, between average and strongly agree (Figure 8).

We also asked them, “What was your favorite part of using this tool?” 44.4% of students’ favorite part was seeing the AI suggestions; 34.7% liked learning about the AI; and finally, 20.8% of students’ favorite part was adding their recipes.

A Full Student Example Student 31 initially tested the AI Chef six times using ingredient counts of 7, 2, 3, 3, 4, and 3, gaining insights into its predictive capabilities. Subsequently, the student added three recipes and conducted two additional tests with ingredient counts of 3 and 4. In response to the prompts, “Do you think the software learned from the ingredients you picked?” and “Did the AI Chef Trainer’s predictions meet your expectations?” the student rated their experience as 4 on the 5-point scale. The student’s responses to post survey question, “What do you think it’s important for AI to learn from?” was option C, indicating the student’s belief that diverse dishes and ingredients enhanced AI learning. Also, the student believed that adding more ingredients “teaches it more information.”

Supplementary Work

We also tested the software in June 2024 with a small group of 12 high school students. We provided a more formal introduction to the AI Chef Trainer, explaining its functionality and how it works. Here, we familiarized students with the tool's purpose and capabilities before their interaction.

We illustrated the input-process-output model of AI/ML through a simple diagram, depicting the input as a set of ingredients, an "opaque box" representing a machine learning model, and the output as recommended recipes. Then, we unpacked the AI Chef Trainer's recommendation process to demystify the "opaque box" and enhance transparency.

Further slides introduced the binary table format, where ingredients were marked as "1" for present and "0" for absent. This visualization helped explain how Naive Bayes used these data to calculate the probabilities of each recipe being a likely match based on the input ingredients. A key assumption of the Naive Bayes classifier is the independence of features, as represented in our presentation. The model assumes that the presence or absence of each ingredient affects the likelihood of a recipe's occurrence independently of other ingredients.

This approach clarified how the Chef AI works beyond a "opaque box" interaction and involved participants in understanding its computational mechanics. After that, a supplementary video showed the software's interface and navigation, helping students to know how to navigate AI Chef.

Discussion

Our first research question (RQ1) is "How can we develop student data literacy skills by allowing students to explore, experiment, and learn from their interactions with an AI model?" Through the AI Chef Trainer, students engaged in a cyclical process of initial testing with the retrained model, contributing recipes (data), and retraining, followed by post-training tests. This cycle enabled students to observe changes in AI Chef behavior with familiar recipe recommendations. Student quotes show that they recognized the adaptive learning capabilities of the system.

While the students were generally familiar with the concept of probability, they were introduced to the concept of confidence in AI suggestions by seeing that recipe recommendations came with probabilities.

Our RQ2 is "How can we demonstrate the importance of data in AI by integrating student-generated recipes into an existing dataset?" Students (87%) actively taught the AI Chef recipes, and half of the students who contributed to the training process evaluated the AI Chef Trainer's performance using the recipes they had added. Students recognized that adding more ingredients (data) led to more specific and accurate AI recipe recommendations. This reflected their grasp of essential machine learning principles, such as the role of training data in enhancing outputs. Some students distinguished between the quantity and specificity of data, indicating that not just the amount, but the type of data, crucially refined AI Chef recipe recommendations.

Two students expressed concerns that too much data might overwhelm the AI or diminish its accuracy. Student

21 believed that adding more ingredients "allows for more dishes which decreases accuracy." This student tested the AI with 17 ingredients in one test. This showed a critical thinking approach where students contemplate the potential limitations or drawbacks of AI systems. A few student responses varied from seeing AI's learning process as similar to human learning to understanding its reliance on trial and error.

Our RQ3 asked "To what extent will student-generated recipes encourage personal expression?" Here, we discovered ample evidence that students were comfortable sharing their cultural heritages around food. Their recipes allowed us to identify patterns of cultural expression, providing insights into how technology can serve as a bridge in celebrating and maintaining cultural traditions in a digital age.

We did not see notable differences among students of the three grade levels.

Conclusions, Limitations, and Future Work

We developed AI Chef Trainer to introduce students to the role of data in ML through an iterative process of testing, contributing their own recipes, and retesting the system. Students were actively engaged with AI Chef Trainer. 87% of students contributed to the AI's learning process in ways that reflected their personal and cultural identities. The evidence of these cultural contributions became vividly clear through analysis of their interaction logs.

Half of the students tested the AI using either subsets of ingredients from their training sessions or the complete ingredient lists from specific recipes they added during training. This was evidence that they understood the role of training data in machine learning.

The student interactions highlight their growing understanding of the importance of data in AI performance. Many students recognized the relationship between data quantity and specificity in AI Chef Trainer performance.

While this study provides important insights, it is important to acknowledge its limitations. Due to time constraints in the after-school program design, students interacted with the system for only eight to ten minutes. The results of this study need to be supported with additional studies with students of various ages and backgrounds and providing more interaction and reflection time.

The students involved in this study had a high level of enthusiasm for the task. Most of them contributed their own recipes. We discovered that the AI Chef afforded for students to share their cultural heritages around food. This hands-on experimentation is crucial for cultivating a comprehension of AI mechanisms, bridging theoretical knowledge with real-world applications. We encourage adoption of our tool and this culturally expressive approach by AI curriculum designers.

Acknowledgments

We thank the teachers and administrators who facilitated our work and the children who participated in this study. This material is based upon work supported in part by the National Science Foundation under Grant IIS-2112633.

References

- Carney, M.; Webster, B.; Alvarado, I.; Phillips, K.; Howell, N.; Griffith, J.; Jongejan, J.; Pitaru, A.; and Chen, A. 2020. Teachable machine: Approachable Web-based tool for exploring machine learning classification. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–8.
- Code.org. 2020. AI for Oceans. Online. Accessed: 2024-08-12.
- Druga, S.; Vu, S. T.; Likhith, E.; and Qiu, T. 2019. Inclusive AI literacy for kids around the world. In *Proceedings of FabLearn 2019*, 104–111. ACM.
- Grover, S. 2024. Teaching AI to K-12 Learners: Lessons, Issues, and Guidance. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, 422–428.
- Lane, D. 2024. Machine Learning for Kids. Accessed: 2024-05-10.
- Lin, P.; Van Brummelen, J.; Lukin, G.; Williams, R.; and Breazeal, C. 2020. Zhorai: Designing a conversational agent for children to explore machine learning concepts. In *Proceedings of the AAAI Conference on Artificial Intelligence V. 34*, 09, 13381–13388.
- Long, D.; and Magerko, B. 2020. What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16.
- Ma, R.; Sanusi, I. T.; Mahipal, V.; Gonzales, J. E.; and Martin, F. G. 2023. Developing machine learning algorithm literacy with novel plugged and unplugged approaches. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*, 298–304.
- Martin, F.; Mahipal, V.; Jain, G.; Ghosh, S.; and Sanusi, I. T. 2024. ChemAIstry: A Novel Software Tool for Teaching Model Training in K-8 Education. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, 792–798.
- Sambasivan, N.; Kapania, S.; Highfill, H.; Akrong, D.; Paritosh, P.; and Aroyo, L. M. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–15.
- Touretzky, D.; Gardner-McCune, C.; Martin, F.; and Seehorn, D. 2019. Envisioning AI for K-12: What should every child know about AI? In *Proceedings of the AAAI Conference on Artificial Intelligence V. 33*, 1, 9795–9799.
- Wickramasinghe, I.; and Kalutarage, H. 2021. Naive Bayes: applications, variations and vulnerabilities: a review of literature with code snippets for implementation. *Soft Computing*, 25(3): 2277–2293.
- Zimmermann-Niefield, A.; Turner, M.; Murphy, B.; Kane, S. K.; and Shapiro, R. B. 2019. Youth learning machine learning through building models of athletic moves. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*, 121–132.