

## **Preservice Teachers' Mechanistic Reasoning about Machine Learning and Artificial Intelligence**

**Dr. Amy Voss Farris, Pennsylvania State University**

Amy Voss Farris is currently an Assistant Professor of Science Education at the Pennsylvania State University. She investigates the intersections of scientific modeling and computing in elementary and middle school classrooms and seeks to understand how learners' and teachers' experiences in scientific computing can support their development of ideas and practices across STEM disciplines. Her teaching encompasses engineering education, preservice teacher preparation, and computational literacies in the Learning Sciences. She is an active member in the International Society of the Learning Sciences (ISLS) and has published numerous journal articles and conference papers on children's scientific and computational modeling in school settings. Her recent work argues that understanding computational thinking requires accounting for the perspectival, material, and embodied experiences in which children's computing work is grounded.

**Ms. Anna Eunji Kim, Pennsylvania State University**

Anna Eunji Kim is a doctoral student in the Science Education program at Penn State. Her research focuses on the interdisciplinary study of developing the assessment for ELL and NES students in Chemistry Education and developing educational curricula within the context of Artificial Intelligence.

# **Teacher Candidates' Mechanistic Reasoning about Machine Learning and Artificial Intelligence**

*Abstract:* Modern humans are frequently embedded in contexts in which machines learn from their everyday actions. Examples include encountering predictive text when texting a friend and facial recognition in personal digital photographs. However, explanations that account for the underlying causal mechanisms of machine learning systems require learners to consider parts of the system and the relationships among the parts of that system. While mechanistic reasoning [3] is a foundational body of abilities germane to engineering, children and many adults rarely conceptualize their interactions with machines in ways that are consistent with the complex and dynamic nature of machine learning systems.

We investigated how undergraduate teacher candidates (TCs) explained an example of machine learning, *Google Quick, Draw!* [13] after playing the game and participating in a series of machine learning investigations. We found that even among an elite group of non-science major undergraduate students, initial explanations of how the computer recognized images rarely focused on how events that lead to a correct guess are linked to one another. In contrast, given opportunities to read and think together about the mechanisms of machine learning, TCs' descriptions of how *Quick, Draw!* works became more sophisticated in terms of the sense of mechanism and the chaining [3] of events that lead to a correct guess. For example, in students' final explanations at the end of our activities, 12 of 22 (54.5%) described the importance of the beginning "stroke" in their doodles or described patterns of key features of doodles as important to how image recognition is accomplished.

Secondly, we explored how building basic understanding of machine learning can support engineering across disciplinary boundaries in K12 contexts. We asked the same preservice teachers to think about how machine learning could be relevant to the content and practices in their area of disciplinary specialization, and to create an initial lesson design that could be used with middle school students (U.S. Grades 4 – 8). The participating preservice teachers' disciplinary specializations were Social Studies (n = 3), English Language Arts (n = 8), and Mathematics (n = 12). We found that all students portrayed that learning goals about artificial intelligence (in general) and machine learning (in particular) were relevant to their focal disciplinary areas and their understanding of literate participation in society. Additionally, some TCs focused on students' understandings of the social and ethical dimensions of artificial intelligence technologies. This included perceptions of the ethical dimensions of AI and the diverse cultural contexts in which machine learning operates. We report the connections they saw and discuss the relevance of machine learning as an example of reasoning about complex engineered systems for young students and for teachers.

## **Introduction**

Humans are frequently embedded in contexts in which machines learn from their everyday actions. However, constructing explanations about the underlying causal mechanisms of a machine learning system (such as predictive text when texting a friend, or image recognition in collections of photographs) requires individuals to examine an engineered system in ways that consider parts of the system and the relationships among multiple parts of that system [1], [2]. These reasoning abilities are germane to engineering disciplines.

Science and engineering educators have characterized mechanistic explanations that focus on the processes that underline cause and effect relationships within a system. In this line of research, more sophisticated mechanistic explanations explain how components of one part of a system affect other parts of the system, or the behaviors of the system as a whole [1], [3], [6], rather than descriptions of the whole that do not account for how those patterns arise. *Mechanistic reasoning* [2], [3], provides an analytical lens that can be used to categorize a learner's causal explanations of phenomena they have experienced. Studies of mechanistic reasoning in school settings have been concerned with ecological systems (e.g., [7]) kinetic toys [2], and computational models of invisible processes in physical systems (e.g., [8]).

As argued by Gupta and colleagues [9], a more complete engineering education also demands a deeper focus on the “macroethics” of engineering, including the social, ethical, and political impacts of engineers' scientific and technological pursuits. However, children and many adults rarely conceptualize their interactions with machines in ways that are consistent with the complex and dynamic nature of machine learning systems [4], [5]. While mechanistic reasoning is often used to account for thinking about the behavior within systems, we also apply a mechanistic lens to think about ethical dimensions of the sociopolitical contexts in which machine learning systems operate. We build on this research by extending mechanistic reasoning as a lens to understand how non-science major undergraduates in a teacher education program explain machine learning examples—including sociopolitical impacts of those systems. We also examine how the teacher candidates (TCs) construe preliminary lesson plans for how they might integrate machine learning in their teaching with young adolescents.

## **Method**

We conducted a design-based research study [10], [11], [12] within a science education methods course at a large land grant university in the northeastern United States. The course is a required course for TCs who will teach middle level grades (grades 4 – 8). Twenty-three of the 25 TCs enrolled in the course consented to provide their assignments and reflections for this study. None of the TCs were science or engineering majors, and all were specializing in one of three content areas: Social Studies (n = 3), English Language Arts (n = 8), and Mathematics (n = 12). All the TCs were in their final semester of coursework and teaching practicum before their full-time student teaching semester. The TCs included 15 women and 8 men.

A key goal of the course was that the preservice teachers learn to recognize and value the science and engineering learning opportunities that are embedded within their areas of content specialization. During Weeks 12 and 13 of the fifteen-week course, learning aims focused on interdisciplinary connections of engineering and technology across the middle level content domains. We integrated two workshops focused on machine learning and artificial intelligence into the course. The design conjectures [12] guiding our work were that machine learning was novel to participants in the class and that investigating machine learning together would generate discussions about what young adolescents think about machine learning. Secondly, we also conjectured that these discussions could support TCs' thinking about the relevance to the interdisciplinary nature of learning about artificial intelligence as well as learning about how technologies impact society.

## Sources of Data

We considered how the learning activities unfolded and all artifacts TCs created in Weeks 12 – 14 in this analysis. This included their classwork throughout these sessions and instructor and researcher fieldnotes. The first author of this manuscript was the primary instructor, the second author was a teaching and research assistant. Written artifacts from the TCs include the explanations of *Quick, Draw!* [13] and a conceptual draft of one lesson plan using one of four AI-related resources that we had introduced during the workshops. The TCs were also asked to respond to short reflective writing prompts regarding the reason why they chose a specific activity in the lesson plan, how the activity they described in the lesson plan allows students to access the learning goals, and how TCs perceive applications of science and technology as important or relevant to students' lives or to TCs' work as teachers.

## Analytic Approach

The following research questions guided our analysis.

1. What are preservice middle school teachers' everyday explanations of machine learning (ML)? Do they see ML as relevant to the disciplinary content that they teach?
2. In what ways can an introduction to machine learning provide a learning experience in scientific and technological literacy for non-science major students in their certification areas?

We conducted a thematic analysis [14] and identified patterns across the data that were relevant to our research questions. We then identified illustrative cases that further explained these key themes.

## Progression of Learning Activities

The learning activities are summarized in Table 1. The first day consisted of three learning activities. In the first part of the first day, members of the class played *Quick, Draw!* [13] for 10 minutes, taking turns. They then created initial explanations of how the computer recognizes drawings and shared their emergent explanations with their peers in groups. During small group discussions, the instructor provided to supplementary resources in the form of two blogs, one concerning the patterns of cultural difference among players' doodles [15] and one overview of a recurrent neural network used to train a machine to draw based on hand-drawn images [16]—which we primarily used for discussion of the images, without discussing of the recurrent neural network. We also had a brief discussion in which TCs shared their thoughts about how artificial intelligence (AI) and machine learning (ML) interact in the daily life of teachers and students and how AI/ML might be relevant to teaching practices.

Table 1. Summary of Learning Activities

Workshop 1 TCs played *Quick, Draw!* and constructed initial explanations of how they thought it worked.

Instructors introduce to *Machine Learning for Kids* materials and assign TCs to train and build one of two ML image recognition models *Snap!* or *Car or Cup?*

TCs build and troubleshoot their ML systems.

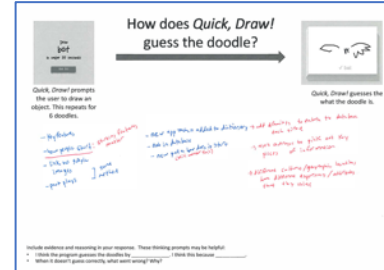

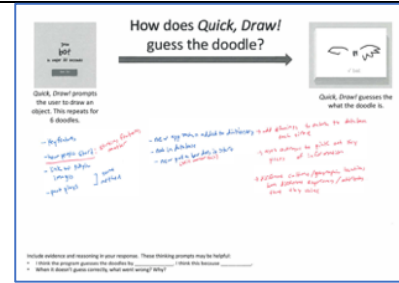
Workshop 2 Review of AI concepts and examples from Workshop 1. Students share their code from Day 1. TCs discussed the confidence rate of the models and their relationship with the data training process.

Video: *What is Machine Learning?* (Code.org, 2020) (3 min)

Revisions to *Quick, Draw!* explanations

TCs played *Guess Who?* game as a whole class first and shared their thoughts about the idea of an algorithmic score for each character.

TCs drafted their lesson plan using one of the resources they experienced during the workshop.



In the second half of the first workshop, TCs carried out one of the two learning activities (i.e., *Car or Cup?*, *Snap!*) from *Machine Learning for Kids* [17], a learning environment that

provides detailed instructions for students to generate machine learning projects in Scratch and other programming languages. The researchers selected activities since the activities had image classification and supervised learning components. The instructor assigned approximately half of classroom to each activity and invited students to move to a different table if they wished to do the activity that they were not assigned, leaving agency for students to select their activity. No students changed tables. *Car or Cup?* is a game in which students select images (sourced from the internet) of cars and cups and use those images to train a machine learning model that sorts cars and cups on a Cartesian plane. Successful training requires choosing images that represent a wide variety of possibilities for cars or cups. The pre-built Scratch code includes special variables and code blocks to call on their trained model, but students must modify the code for it to work properly. Students modified the pre-built Scratch code to test their systems.

The alternate activity (conducted by approximately one-half of the class) was *Snap!*. In this game, players hand-draw the symbol for each suite in a standard deck of playing cards (heart, diamond, spade, club) on an index card, and use their drawing to train a machine learning model. As with *Car or Cup?*, the pre-built Scratch code includes special variables and code blocks to call on their trained model, and students must make modest changes to their code so that it correctly calls on their training dataset. The game then times players with a visible timer as they search for and pull up their hand-drawn card and hold it in front of the computer camera for image recognition. Successful models require students to hold up the appropriate suite within fractions of a second.

On the second day of the workshop TCs first reviewed key ideas from the first day of the workshop. A student demonstrated her *Car or Cup?* model with the associated Scratch codes. The instructor problematized that the model can recognize some images but is often incorrect, especially when presented with a “real” car or cup in front of the camera. A second student presented his *Snap!* model with the associated Scratch codes. The instructor also presented her buggy model as an example, in which she had erroneously reversed the variables for “car” and “cup”, so that the system was almost always incorrect. No students were able to recognize the bug without the instructor showing the error explicitly. Through demonstration, TCs discussed the confidence rate reported by their models and their relationship with the data training process, making the connection that the quality and diversity of the dataset and the specific features of the images. For example, that noise can greatly impact how an image is interpreted. The relationship between training and the performance of the model was unclear up to this point for the TCs. The discussion goal was that students would begin to attribute the accuracy of their image recognition models to specific features of the training process and to the code. We see this discussion as a key turning point in the ways that class members were reasoning about the general mechanisms through which image recognition can work in machine learning systems.

Next, the students watched a short video about machine learning concepts (Code.org, 2020) and reflected on their initial ideas about the relevance of ML to teaching. TCs then revised their initial model of how *Quick, Draw!* recognizes the drawing of users.

After a short break, students transitioned to a new game: Students played *Guess Who?*, a popular logic board game available in the US since the 1970s. To play the game in the whole class format, the instructor used a web-based version of the game. In *Guess Who?*, players

secretly select one of approximately 24 characters, and their opponent asks simple questions about the character. The objective of the game is to identify which character your opponent has selected before the other player identifies the character whom you selected. For example, if the opponent asks, “Does your character have glasses?” and your character does wear glasses, you would answer, “Yes, my character is wearing glasses.” The other player can then eliminate all characters who are not wearing glasses. We played two rounds and the instructor asked the TCs to share their strategies for winning the game. Students ultimately agreed that identifying a character who is most similar to the typical character was most strategic for winning the game. For example, choosing a women was a minimally strategic choice, as most of the characters appear to be renderings of men. If asked, “Is your character a man?” responding “yes” makes it more difficult for your opponent to win the game.

The instructor then briefly introduced a lesson overview from the publicly available lesson resource on Machine Bias and *Guess Who?* [18]. The lesson plan demonstrates how a system used in the criminal justice system unfairly informs sentencing decisions by assigning a numerical score to offenders based on their perceived “likelihood of... committing a future crime” [19]. People who are Black consistently are scored with a higher numerical “risk” score, creating a systemic injustice that lands Black offenders with tougher sentences. The instructor summarized the article associated with the text and posted the article for the TCs to review outside of class. Class members connected this article to ML systems used to recognize faces in photographs.

The culminating assignment was for students to consider any of the four major learning activities (i.e., *Car or Cup?*, *Snap!*, *Guess Who?*, and *Quick, Draw!*) they had participated in during the workshop and think about prospective learning goals for their middle school students. Students were asked to demonstrate that they could meaningfully connect ML or AI to standards-based learning objectives at the middle level, and to explain their rationale for how they did so. They created draft lesson plans from template provided by the instructor and explained their design decisions in three open-ended reflection questions.

## Analysis of Teacher Candidate Assignments

*Mechanisms of Image Recognition in the Quick, Draw! game.* In the model-based explanations assessment, students drew or wrote about their initial explanations of how does *Quick, Draw!* “guess” the doodle. We prompted students to think about the following in crafting their responses:

- I think the program guesses the doodles by \_\_\_\_\_. I think this is because \_\_\_\_\_.
- When it doesn’t guess correctly, what went wrong? Why?

After the Workshop 2 activities, we asked students to revise their initial explanations, again using the same prompts. A sample is shown in Figure 1, where the blue text is the original explanation, and the red text is what the student added. We analyzed initial and final models to understand both their initial explanations and what salient changes were made to their initial explanations. In the initial models, most students struggled to identify how the game works, and



many suggested that it is making comparisons to images available on the web, for example, by sourcing Google Images. In the revisions, more normative answers focused on key features of the drawing, sourced from an expansive database of past plays. Two students (i.e., Annabel, Helen) mentioned that *Quick, Draw!* might generalize/distinguish key shapes that distinguish the very beginning of the doodle from others. Additionally, a couple of students (e.g., Helen, Michael) recognized that finding similarities and differences between other drawings is critical to the success of the game for guessing the doodles. Notably, no student was able to connect their playing of the game to key language used to describe machine learning, such as “neural networks” or “training.”

Our analysis of the explanations indicates that all students explained how the software works with some level of mechanistic reasoning, partially driven by the prompts. While all students who were present to generate initial explanations (n = 21) wrote something about particular entities or activities in the system (e.g., the “database”, patterns in the database), 18 responded with some indication of chaining (n = 15, including causal relationships between any entities and activities) or discussed specific properties of entities (n = 12; e.g., the first stroke of a drawing) that might be associated with that chaining. In their revised models, all students who were present (n = 22) indicated at least chaining or individual properties of entities. Although the initial model and revised model did not show a significant difference in the number of more sophisticated codes using Russ’s [3] hierarchy of codes for mechanistic reasoning, revised explanations were overall more specified and linked to the important mechanisms of machine learning, evidencing a closer coupling between parts of the explanation.

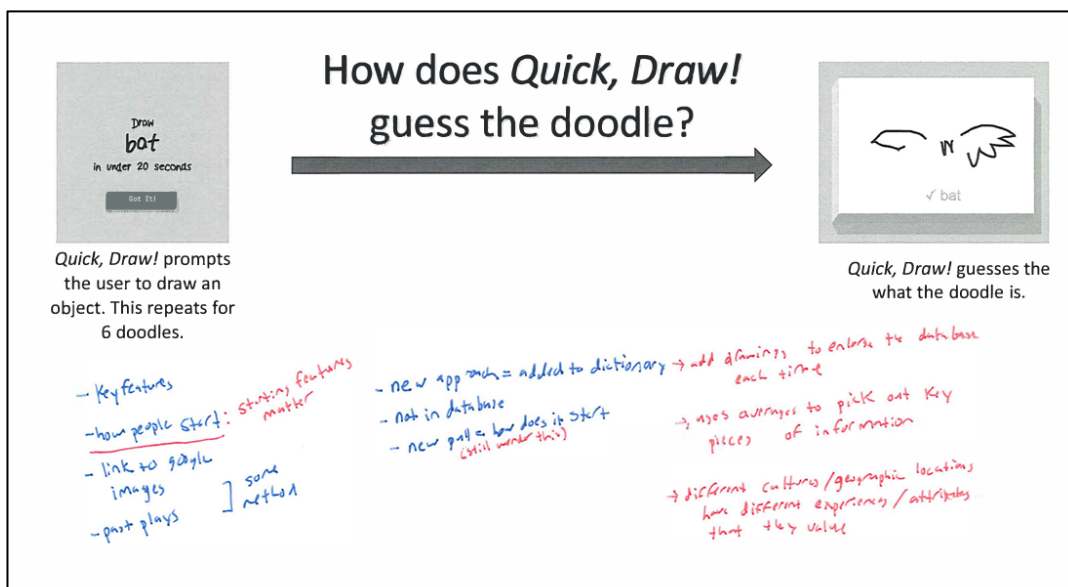


Figure 1. Students made initial and revised explanations of how they thought *Quick, Draw!* works.

In summary, TCs revised models near the end of the second workshop demonstrated a greater recognition of how particular entities and activities are coupled together in mechanisms that lead to a correct “guess.” For example, three students wrote about how the first stroke of the drawing matters in terms of the game’s ability to quickly guess what the player was drawing

(Kathy, Catherine, Alex). Additional students generalized that the software’s ability to recognize similarities and differences between other drawings in the reference dataset is critical for guessing the doodles.

In the lesson plans, all students employ one of the major activities we did in the class in a lesson plan for middle grades. Students from different certification areas built lesson plans related to ML/AI with different tools, including Quick, Draw!, Machine Learning for Kids, and the Guess Who teacher resource [18]). The usage of tools was distributed across all resources (*Quick, Draw!*,  $n = 12$ ; *Machine Learning for Kids*,  $n = 7$ ; *Guess Who?*,  $n = 2$ ). While students focused on different aspects (ethics of image recognition, cultural difference in sketching, using variables in block-based code) each supported an elementary explanation of mechanisms of ML. Their lesson plans are aimed for different lesson goals related to their areas. All students meaningfully connect their lesson plans with relevant ML/AI concepts when provided with a “bank” of possible standards for science, engineering, technology and computing, and the Common Core State standards. In what follows, we describe two primary themes that students recognized that connect their emergent understanding of ML systems to disciplinary ideas that they teach about in the classroom.

### ***Case 1: “Averages” of Doodles as a Feature of Mechanism***

In the days between the first and second workshop, Kathy (all names are pseudonyms) used *Quick, Draw!* in her field placement. She was eager to report to the authors that she had already used it in her placement and the students thought the game was intriguing and told us before the beginning of class. Her practicum placement was a 5<sup>th</sup> grade classroom in a rural community. Due to the pandemic, her class was hybrid. She, the mentor teacher, and the majority of her students attended in person and some students attended remotely. Kathy chose *Quick, Draw!* for an interdisciplinary period within the existing structure of the day and defined the primary learning aim as “Computers can learn from patterns in photos, images, or drawings to recognize new images.” She also reported that she selected the game because it helped facilitate community across her students in the classroom and learning from home.

Kathy commented that the most interesting thing that happened in her enactment of the lesson was that students were able to reason about how patterns in the large dataset of doodles lead to the software making correct guess about the doodle. When Kathy shared an image from the one of the blogs we had reviewed in the Workshop (Figure 2), a 5<sup>th</sup> grade student described the blurriness of the image as a kind of “average” that was calculated from a large set of images, indicating that this young adolescent was able to reason about the blurry lines as representative of a distribution of doodles, with the most salient areas indicating some kind of central tendency of the spatial relationships in the doodles, and furthermore, that this central tendency is important for how the software makes a correct “guess” about the doodle. This instance demonstrates the brilliance of this teacher candidate and the young people she was teaching, and the potential for machine learning systems to generate learning contexts in which students begin to exercise complex mathematical and statistical reasoning even before they have normative language to describe these concepts.




Figure 2. An overlay view of thousands of doodles of chairs from 48 countries. A learner in Kathy’s classroom identified that blurry representations of overlaid doodles is an “average” of the doodles. Image from Kyle McDonald, <https://twitter.com/kcimc/status/902229769919406085>. Used with permission.

### ***Case 2: Linking ML Explanations to Middle Level Classrooms***

The second major theme regards how students were able to link explanations of ML to content that they found valuable to teach in middle level classrooms, particularly in classes that are not focused on science, engineering, or technology. We focus in this case on the two students who selected to plan their lesson plans around *Guess Who?* and bias in machine learning.

Jamie selected the *Guess Who?* game as a means to teach young adolescents about critical consciousness for living in technologically driven societies:

*This lesson is valuable for social growth. Students use stereotypes and bias[es] every day, likely without knowing it. Becoming aware of our personal bias and stereotypes will make us more considerate and empathetic individuals.*

In their description, Jamie wrote that in the integrative lesson she was planning, students will first play *Guess Who?* multiple times, and she will ask them to record the characteristics of the characters that they chose in the instances in which they won the game. This is to highlight that winning the game means drawing on the biases of the game. She further explained that they would prompt students to think and write about why the characteristics of the character may have helped them win that round of the game. She wrote that, “Ideally, students will conclude that “winning characters are the ones with indistinguishable or common features. Finally, I’ll connect student responses to the bias article [that we had skimmed in class]. In the focal machine

learning system used in criminal justice, “Individuals are stereotyped as a result of their distinguishable features.”

Jamie also saw the social and psychological relevance of the activity as a way to connect to the lived world of students:

*Psychology was always my favorite subject because of how closely it relates with our world outside of the classroom. My work as a teacher will always mirror my philosophical value of connecting content to the real-world.*

Similarly, Joshua wrote a lesson in which young adolescents would learn about how systems that use artificial intelligence can impact life in ways that are helpful, as well as ways that represent people unfairly. Joshua also identified that they wanted students in their class to think about data bias and how machine learning systems may produce content treats some groups unfairly and reinforces inequities. In their reflection, Joshua wrote:

*The students will be able to see how the AI determines the threat level which is done by physical appearance which we as a class or on our own can determine what other variables does the AI forget to account for and how can we make a better AI for this task or can we improve the AI that already exists...*

## **Summary of Findings**

Regarding RQ1, concerning preservice middle school teachers’ everyday explanations of machine learning, we found overwhelmingly that the teacher candidates in this course had intuitive but naïve ideas about examples of AI and machine learning in their everyday lives, and often confused examples of machine learning with non-learning systems. After playing with *Quick, Draw!*, we prompted them to create initial models of how the system works. Many students thought that the game drew on collections of images (such as they would see if they completed an image search). In the revised model, all TCs indicated some causal chaining among entities and activities or wrote specifics about properties of entities. Some students began to think about the most prominent features of sketches as ways that the computer makes a guess. Overall, students’ explanations increasingly took on a mechanistic approach to explaining how the software works.

Concerning RQ2, we conclude that the introduction to ML provided space for broad perspectives of AI. In initial discussions about how ML plays a role in the lives of the young adolescents that they teach, the TCs were able to give some examples but were largely unable to think of connections to the primary disciplines that they teach. However, final lesson plans all include meaningful connections to learning about AI, contextualized in learning goals that are prominent in middle grades classrooms.

## **Discussion and Implications**

Humans work alongside machines in problem solving, and artificial intelligence and machine learning applications have changed the landscape of possibilities for how humans “think with”

machines. We have described how a non-science major teacher candidates began to reason about machine learning and think about implications for their teaching. In our study, teacher candidates considered how they—and the young adolescents they teach—think about how machines learn in games like *Quick, Draw!*, troubleshooting their own machine learning code and data, and thinking about the social and ethical implications in machine bias used in the criminal justice system.

We contend that while the TCs explanations of ML at the end of the second workshop were quite modest from a technical perspective, the lesson plans they developed demonstrated curiosity and deep engagement with everyday explanations of ML. Through their experiences with a handful of ML tools across just a few hours, they began to see many substantive curricular connections and relevance to students' out-of-school lives. Secondly, the emergence of ethics and data bias as learning goals arose from only a rudimentary understanding of *how* machines learn. We do not believe that critical perspectives about machine bias would have been possible without a basic mechanistic explanation of the processes involved in machine learning.

### **Acknowledgements**

The authors are grateful to the teacher candidates who creatively engaged in this work and to Mehrdad Mahdavi and Swaroop Ghosh for inviting us to think together about machine learning systems in drug discovery contexts. This work is partially supported by the National Science Foundation NSF OIA-2040667.

## References

- [1] Bolger, M. S., Kobiela, M., Weinberg, P. J., & Lehrer, R. (2012). Children's mechanistic reasoning. *Cognition and Instruction, 30*(2), 170-206.
- [2] Weinberg, P. J. (2020, June). A Pathway Towards STEM Integration: Embodiment, Mathematization, and Mechanistic Reasoning. In *2020 ASEE Virtual Annual Conference Content Access*.
- [3] Russ, R. S., Scherr, R. E., Hammer, D., & Mikeska, J. (2008). Recognizing mechanistic reasoning in student scientific inquiry: A framework for discourse analysis developed from philosophy of science. *Science Education, 92*(3), 499-525.
- [4] Royal Society. 2017. Machine learning: The power and promise of computers that learn by example. Technical Report.
- [5] Fiebrink, R. (2019). Machine learning education for artists, musicians, and other creative practitioners. *ACM Transactions on Computing Education (TOCE), 19*(4), 1-32.
- [6] Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science, 1* – 25.
- [7] Dickes, A. C., Sengupta, P., Farris, A. V., & Basu, S. (2016). Development of mechanistic reasoning and multilevel explanations of ecology in third grade using agent-based models. *Science Education, 100*(4), 734-776.
- [8] Jacobson, M. J., & Wilensky, U. (2006). Complex systems in education: Scientific and educational importance and implications for the learning sciences. *The Journal of the Learning Sciences, 15*(1), 11-34.
- [9] Gupta, A., Turpen, C., Philip, T., & Elby, A. (2019). Narrative co-construction of stances towards engineers' work in socio-technical contexts. In *Critical, transdisciplinary and embodied approaches in STEM education* (pp. 251-272). Springer, Cham.
- [10] Cobb, P., Confrey, J., diSessa, A., Lehrer, R., & Schauble, L. (2003). Design experiments in educational research. *Educational Researcher, 32*(1), 9-13.
- [11] Design-based Research Collective. (2003). Design-based research: An emerging paradigm for educational inquiry. *Educational Researcher, 32*(1), 5-8.
- [12] Sandoval, W. (2014). Conjecture mapping: An approach to systematic educational design research. *Journal of the Learning Sciences, 23*(1), 18-36.

- [13] Google Quick, Draw! <https://quickdraw.withgoogle.com/>
- [14] Braun, V., & Clarke, V. (2019). Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4), 589-597.
- [15] Jongejan, J. (2017, December 8). A look at one billion drawings from around the world. *The Keyword*. <https://www.blog.google/technology/ai/quick-draw-one-billion-drawings-around-world/>
- [16] Ha, D. (2017, April 13). Teaching machines to draw. *Google AI Blog*. <https://ai.googleblog.com/2017/04/teaching-machines-to-draw.html>
- [17] Lane, D. (2021, December 8). Machine learning for kids. <https://machinelearningforkids.co.uk/>
- [18] Guess Who. Resource from AI for Teachers. <https://aiforteachers.org/resource/guess-who-2>
- [19] Angwin, J., Larson, J., Mattu, S. & Kirchner, L. (2016). Machine bias. *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [20] Code.org. (2020). *What is machine learning?* [Video]. YouTube. <https://youtu.be/KHbwOetbms>