The Contour to Classification game

Irene Lee[1*] **Safinah Ali**[2*],
[1]MIT STEP Lab
[2]MIT Media Lab
Cambridge, Massachusetts, 02139
ialee@mit.edu, safinah@media.mit.edu

Abstract

The Contour to Classification game teaches middle school students basic concepts in supervised learning through an online variant of the Neural Network game presented at AAAI 2019 Fall Symposium Teaching AI in K-12 track in 2019. The new online version engages students in the classification of images in a multi-step process of compositing salient smaller features to form larger features and ultimately a partial contour of an object that is used to make a classification. After evaluating the classification, information is sent back through the network in processes mimicking back propagation and gradient descent. Additional rounds of the game can be played to witness how the network evolves and gets "better" at classifying images from contours. Through this activity, we aimed for students to learn the structure and components of a neural network and the processes involved in supervised learning. This version of the Neural Network game supports online student learning by providing the Neural Network game experience in a fully digital format. We will describe the evolution of the concept of the game, play the game, then share preliminary findings from implementing the online Neural Network game in a pilot research study. We will conclude with a discussion of if and how the evolving design addresses classroom needs and scaling considerations.

Introduction

The Contour to Classification game was created to help K-12 students and teachers understand how learning is achieved by neural networks. The goal was to engage learners in activities that would help build mental models of the structures and processes in supervised learning on neural networks. In particular, the structure and function of input, hidden, and output nodes are represented, and the processes of feeding forward, evaluation, back propagation and gradient descent were mimicked. This live simulation strategy is similar to others that have shown promise for supporting learners in developing the mental models they need to retain knowledge, and use knowledge adaptively and flexibly (National Academies of Sciences and Medicine 2018). The designers envisioned that having this concrete experience acting as a node participating in the processes of AI would help learners: a) understand how a neural network learns over time and the limits of its understanding; b) think through where classification and prediction can go wrong, and c) develop ideas about how to fix the errors in classification.

Background

The Contour to Classification game originates from an educational activity called the "Human Neural Network Game" developed by Catherine Schuman, Steven Young, Thomas Proffen, Dasha Herrmannova of Oak Ridge National Laboratory for the TechGirlz program (TechGirlz 2018). In the Human Neural Network Game students play the role of nodes in a 3-layer Neural Network. The structure of the neural network is formed by students sitting in predefined rows representing layers of the neural network. Input nodes are provided with an image and must write four words on individual cards that describe the image. Each of the words is distributed to each of the four hidden layer nodes. Hidden layer nodes select two words from their set of four to pass on to the output node. The output node creates a caption using four of the eight words it has received from the hidden layer nodes. Subsequently the "unveiling" takes place wherein the original image and its caption are exposed for all to see. In this live action game, only the feed forward process was simulated and the game mimics the testing process in supervised learning.

In 2019, Lee and Martin extended the live action game to introduce the processes of back propagation and gradient descent to mimic the training process in supervised learning (Irene Lee 2019). After the unveiling of the original image and its caption, students come up with an evaluation function to assess how well the network performed on captioning, then feedback is provided to nodes by passing circled words (if the word appeared in the original caption) or uncircled words back to their originators. After a discussion of the feedback and possible adjustments to the node's word selection behavior, the students can play additional rounds with new images and captions to see if/how the neural network learns to get better at captioning. During a wrap-up discussion, the facilitators reinforce that students were modeling an artificial neural network and review the analogies made between the actions in the game and processes in supervised learning. In this version of the game, back propagation described as sending information back through the network and gradient descent is described by as adjusting

*Equal contribution.

one's algorithm for choosing words. This live-action version called the "Artificial Neural Network game" was tested with educators at MIT and visiting middle school students from Excel Academy, and at AAAI 2019 Teaching AI in K12 Symposium.

While the game proved engaging and enjoyable to participants of all ages, notable difficulties in implementation were encountered. The setup of the game in a physical space was challenging as chairs and tables needed to be moved to create the physical network. Preparation of the materials needed was time consuming - input images needed to be printed and cards for writing down words had to be assembled into packets before hand for each node. Within the game itself, selecting words and composing captions was not familiar to students and how the behavior related to a real world process was obtuse. Furthermore, it was difficult to play multiple rounds in the time allotted (typically 45 minutes) and thus participants were not able to witness the evolution of the network over time simulating learning. These issues led to the reconsideration of the format and the choice of creating captions for images as a target goal of the activity.

In January 2020 the Artificial Neural Network game was adapted to be an online interactive game for use in an online AI Book club style professional development for teachers. The network diagrams were ported into Google Drawing and layered with data boxes for words. Players in an online platform (typically Zoom) were assigned the role of nodes as before, and input nodes were moved to a breakout room to privately view the original image to be captioned prior to returning to the main room to generate descriptive words. A drawback of this instantiation was that all of the selected words were visible to all players during game play thus reducing the element of surprise and possibly impacting the selections made by those acting as hidden layer nodes. This online version was tested with AI book club members, after school club participants in Waltham Public Schools, and within "Developing AI Literacy" summer workshops for students held in 2020.

In an effort to reconceive the goal of the activity to align with convolutional neural networks used in image classification, Lee designed a new version of the online Artificial Neural Network game in which players create partial outlines or contours of images from photographs as input to the neural network rather than words. The image below left Figure 1 was an inspiration to the design. A prototype was developed using clear mylar sheets and thin masking tape pieces (each 2" in length) to outline contours. The compositing of clear mylar sheets with line segments formed larger contours simulating the levels of abstraction in the deep learning process.

During the Summer 2020, Safinah Ali developed the digital prototype for the Contour to Classification game. In this game, the students work together as a neural network to classify images of animals. The game is played by 10 participants taking the role of the nodes of the input layer (4 participants), the hidden layer 1 (4 participants), the hidden layer 2 (1 participant) and the output layer (1 participant). The input layer nodes are presented with an input image on their screens. They can draw a contour line on a transparent canvas overlay on top of the input image with a limited number of pixels (30 px) to select segments of the input image as information to send to the next layer. Each hidden layer receives information from the previous layer, and forms a strategy to pass on some of this information, and discard the rest. While this strategy is likely somewhat random at first, the hidden layers improve their strategy over time to send the most useful information forward. The output layer node receives a final set of information that passes through the hidden layers, and makes a decision to classify the contour as representing one of 6 animals. If the classification is incorrect, all the nodes of the network receive negative feedback, and get a chance to re-strategize. If the classification is correct, the nodes receive positive feedback, and strengthen their previous strategy. Players play multiple sub-rounds until the network reaches the correct answer. Once the neural network is successful, the players can play a new round wherein they receive a new input image. Participants stay in the same role and use their knowledge of the strategies that proved more successful to classify the new image. This interaction replicates how nodes in a neural network performing image classification behave. In the next section we will walk through the game play in detail.

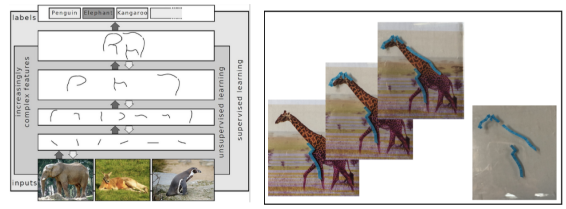

Figure 1: Image classification via composite line segments and prototype.

The Contour to Classification game

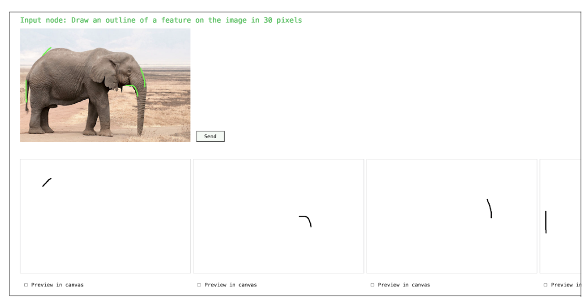

Figure 2: The Contour to Classification game input layer node view

Learning Objectives

In this activity, the ultimate learning goal is for students to understand how neural networks learn. Students should also be able to recall the different components and processes of Neural Networks and explain how they work (using the analogies presented in the game). Students should be able to gain the following takeaways from the lesson:

- Neural networks is a kind of supervised learning.
- The process we played out in the activity is called the "training phase."
- Neural networks consist of an input layer, one or more hidden layers and an output layer which all consists of nodes.
- The nodes are connected by channels (or links) each of which has a weight associated with it.
- Neural networks use feed forward, evaluation and back progression processes to tune the network over time which helps the neural network learn to classify images.
- Neural networks are designed for a specific purpose.
- Training a neural network is a multi-step process of tuning weights and adjusting algorithms.

Target age group

Middle School students (Grades 6 - 8)

Time needed

60 minutes

Materials needed

Access to a computer with a web browser. (Mobile devices also work but are not ideal since they do not afford a large drawing canvas).

Setup

The game is built for the web and hosted on Heroku platform (Heroku Platform 2007) which runs a Node.js web server (Node.js 2009). Students can play the game using any browser with a functional internet with no additional setup required. Students are sent links to the server that designate their role as input, hidden or output layer nodes.

Prerequisite knowledge

Typically, in a workshop setting, this game is played after students have learned about Decision Trees and have experience with Supervised Learning (Classification algorithms) through engaging with Google's Teachable Machine (Google 2019). This sequencing while recommended, is not mandatory.

Connecting to students' prior experiences

Many students are familiar with the "telephone game" in which players line up in a single file then the person at one end of the line whispers a message to the next person in line then the receiver whispers what they heard to the next person in line and so on. Once the message reaches the last person in line, the message received is compared with the original message. Often the message gets horribly garbled along the way.

Introduction to Neural Networks

A short slide presentation introduces the activity and the concept of neural networks as follows: Artificial neural networks,or simply called neural networks (NN), are computing structures and algorithms that are inspired by the biological neural networks of the human brain. An NN is based on a collection of connected units or nodes called artificial neurons, which loosely model the neurons in a biological brain. A typical NN consists of an input layer, several hidden layers, and an output layer, all of which consist of multiple artificial neurons. Each layer is connected to the next layer using channels that have corresponding weights. Weights can represent the quality or "goodness" of information sent along that channel. Further, a NN consists of three main processes: feed forward (or forward propagation), evaluation and back propagation.

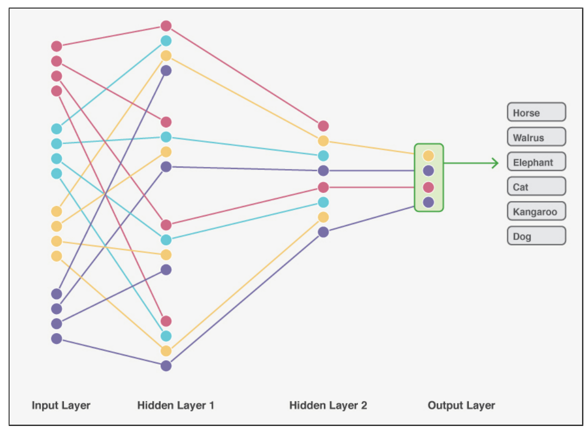

Figure 3: The Contour to Classification game mini-map showing the schematic of the neural network

Players and roles

The game is played by 10 participants. Participants log in, and select a role (pick from the ones remaining). Participants take on the roles of input nodes (4 participants), hidden layer 1 nodes (4 participants), hidden layer 2 nodes (1 participant) and the output node (1 participant).

The interface and game play

Input layer node view: All 4 input nodes see an input image. Each participant can draw 4 contours with a maximum of 30 pixels each on top of the image. They contours appear on blank canvases (image left, lower portion) that show the outline indicators of the input image. A minimap view (image right) shows the player's position in the neural network.

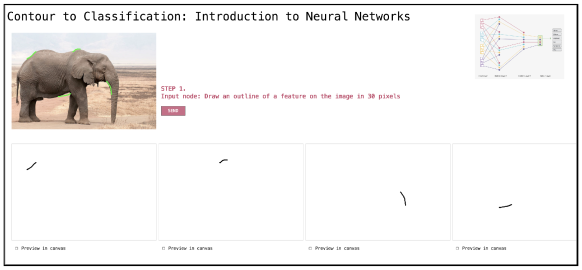

Figure 4: Input layer node view

Next, the Hidden layer 1 nodes take their turn: 4 Players in the hidden layer 1 receive a total of 4 contours, one from each input node. They also have a blank canvas that they can use to composite these contours. They click to select contours and see the composite in the canvas in the upper left of the screen.

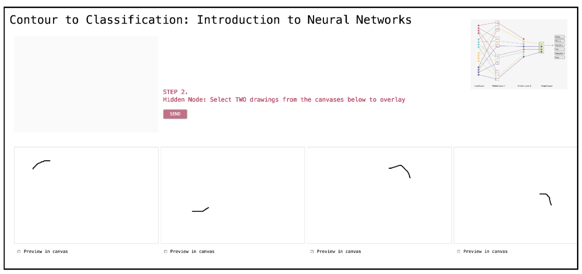

Figure 5: Hidden layer 1 views 4 inputs from the input layers

When they are satisfied with their selection of a combination of 2 out of the 4 contours, they click on the "send" button to send to the next hidden layer.

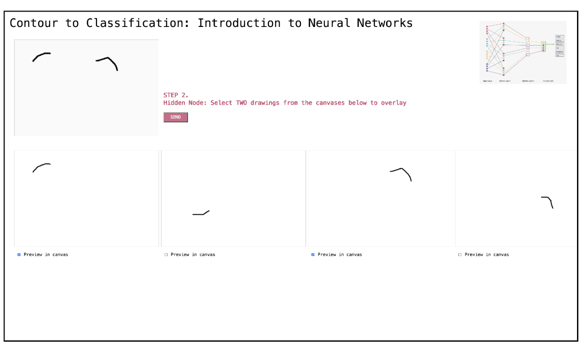

Figure 6: Hidden layer 1 forms a composite of 2 contours to send to hidden layer 2

Hidden layer 2 nodes take their turn: Players in the hidden layer 2 receive 8 contours from the hidden layer 1 (2 from each node in Hidden layer 1) . Players can composite 4 of these contours to form a final shape (seen in the composite view in the upper left). Then the output layer node takes its turn: it receives 4 composited images and has to now classify what kind of animal is represented. The output layer node chooses from the 6 animal options provided.

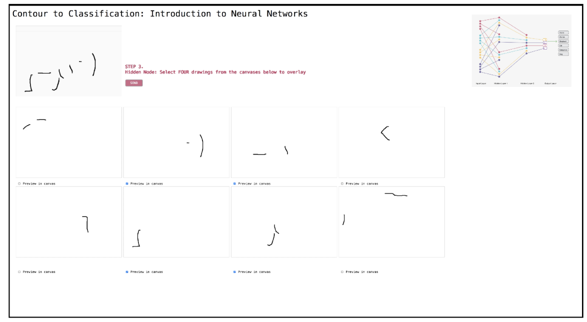

Figure 7: Hidden layer 2 received 8 contours and forms a composite of 4 contours to send to the output layer

Evaluation and back propagation step: After the classification is made an evaluation of correct or incorrect is given then the input nodes and hidden nodes receive a positive/negative feedback based on the evaluation. Information from each layer that reached the last level is marked red (incorrect) or green (correct) based on the classification's correctness. All nodes are asked to come up with a new strategy to make a better prediction prior to a next round of play. For instance, input nodes may strategize that it is beneficial to mark the ears or trunk since they are more characteristic of elephants. The hidden nodes may strategize that to select inputs with contours that are closer together.

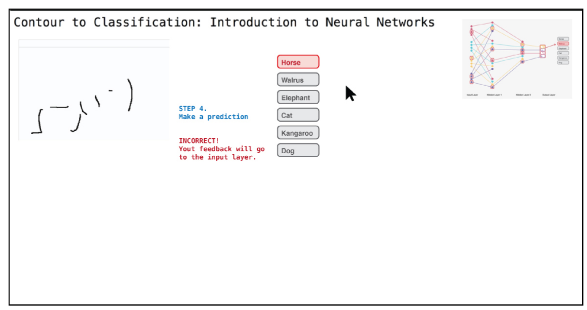

Figure 8: Evaluation and propagation. The output layer makes an incorrect classification.

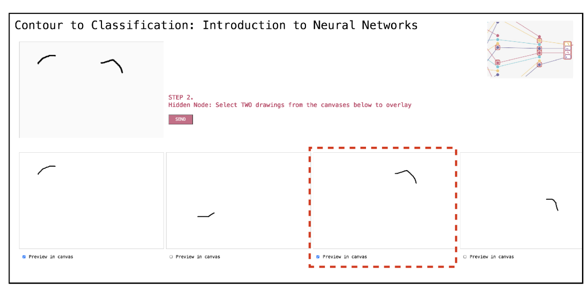

Figure 9: An incorrect feedback is sent to the input and hidden layers

Multiple sub-round play: Based on their re-strategization, players play another sub-round and send new combinations to the output node. Players keep repeating this process until the output node classifies correctly.

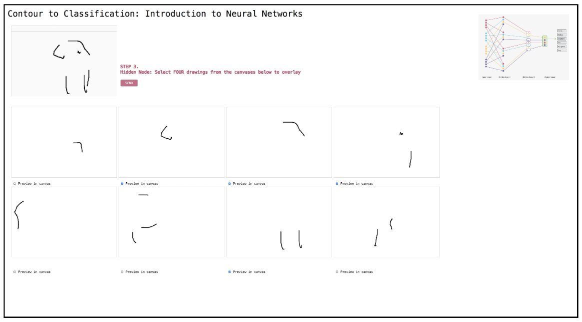

Figure 10: Players re-strategize and play another sub-round

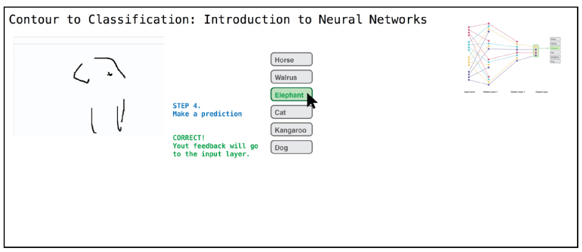

Figure 11: Players keep playing multiple sub-rounds till the output node classifies correctly

Learning to classify new images: Based on this knowledge of what strategies work well for classifying animals, students play another round with a new input image. Input and hidden nodes are now empowered with better strategies of sending information to the hidden nodes to help classify the image correctly. These strategies, or knowledge of what works and what does not work is essentially what constitutes *weights* in a neural network.

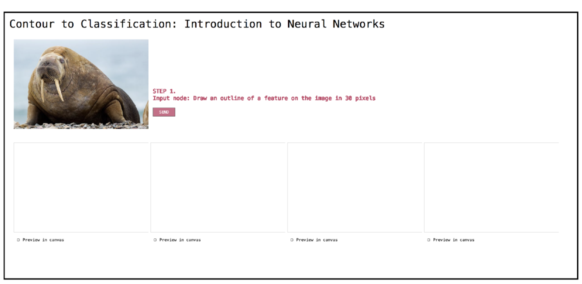

Figure 12: The NN has now learned to classify new images. Input nodes see a new input image in the next round.

Findings from a preliminary study

The online Artificial Neural Network game (a precursor to the Contour to Classification game) was tested in a pilot research study. In partnership with two youth serving community organizations, the study was conducted in three AI summer workshops offered in the summer of 2020 reaching a total of 65 students between the ages of 10 and 16. Due to the COVID-19 pandemic, the workshops were offered virtually and students were instructed online on Zoom.

During a one hour session, the game was introduced by a team of researchers and educators then played within a Google Drawing. Participants were grouped by age into three groups of 10 or 11 individuals into small groups followed by a whole group discussion. The discussion included relating the activity to ethical implications, and its connections to AI careers. Observation notes were collected during the game as well as students' audio and chat data from their discussion after the activity. Students completed a daily reflection and were interviewed about their recollection of the activity at the end of the workshop.

Students reported that they had fun playing the activity. In their responses to "did you find the Neural Network game engaging?" they pointed to the element of surprise, coming up with words to describe an image, and the interactive nature of the game as engaging.

> "...It was pretty fun because some people didn't know what it was and some people knew what it was. You know, it was kind of difficult, but it was actually pretty, really fun."

> "...Describing the picture was cool, and picking the words because putting it through the different layers with the other people was cool because they won't see the photo."

> "Everybody would have their own sets of words and others would try to base their own imagination using those words and try to come up with one single answer. So that was really fun."

> "That one was a lot of fun, because, as I said again, I was interacting. Anything for me that was very interacting and engaging I had a really good fun time with that activity."

Though some students reported having difficulties within the social dynamics in their small groups.

> "Yeah, but the thing is, a few kids in my group, they just weren't participating, really. So it made it boring, I guess. But otherwise I had fun and it was a pretty fun activity."

Others commented on the game being difficult or hard.

> "So the person who has to put like all the words together to make the caption, I feel like it's harder for them because they don't actually know what the caption is and they have to think what's the most reasonable answer. What's the one that makes sense. And they just have to think harder. They probably need to train more. Yeah."

> "...If you're the hidden layer or the output layer, it's like, you don't really know what to pick because you can['t] see the picture. But I guess that's how computers think so if that makes sense."

> "At first the neutral networks one was pretty hard, but then I understood it more. Because I didn't understand

> the instructions that you guys were talking about. So it made it difficult until it came to my turn. Then I kind of got the hang of it."

End of day reflections submitted on the day the game was played mentioned learning about neural networks. Twenty-two of 43 respondents mentioned neural networks when answering the question "what did you learn today?" Six out of 43 reported learning that neural networks were used in / or was a kind of supervised learning. Five out of 43 reported learning "how neural networks work" or "operate". Three out of 43 responded learning about the specific layers in the neural network and what each one does.

In terms of how they learned, students spoke of playing an active part in the live simulation, seeing how the network changes over time, and relating it to an earlier game on decision trees.

> "at first I didn't really know what's going on, but after we went back, I could see how everyone played a part and that's how computers that's one way computers collect their data. So that was pretty good to see that."

> "...that one was one of my favorites because we got to be like real life AIs passing on. We got to do it again, train more and see how it changes, because the layers they had to do it. . . "

> "I really liked it was because it was using hilarious images and it was fun to pass down information. Kind of going from the top of a tree, passing down information to the bottom of a tree. And then they guess what it is. That was just interesting."

When asked how they might improve the activity, students mentioned needing more time, more repetitions, clearer graphics, and a different mechanism to show only the input nodes the image to be captioned.

> "Maybe we should spend more time on it because the pace, it was too fast. Yeah. Right. Maybe if we slow it down and spend more time explaining how the game is supposed to work"

> "It was just the first time we'd actually seen how neural networks saw or worked and all the arrows made it really confusing. But I guess in the end it was better explained."

> "I didn't like the way we did it. I like making the caption and stuff, but I feel like there could have been a link sent to, instead of doing the breakout rooms, because we only saw the image for so long and I actually missed the image, so I didn't get to see it. So maybe someone could have sent a link privately to the people that were going to see the image and then we could have gone into breakout rooms, and then did the description of the image, instead of going into two different breakout groups.

The evolution of the design from the Artificial Neural Network game to the Contours to Classification game addresses issues of playability and usability barriers our learning objectives, and scaling considerations. In terms of playability, we removed some barriers that students encountered during game play in the online Artificial Neural Network game. The drawing of lines to indicate salient features of the input images may be more accessible to young students than coming up with descriptive words. Additionally, we removed the need for students to position their descriptors (either words or contours) in specific correct slots for delivery to subsequent nodes, thereby reducing confusion. Most importantly, in the Contour to Classification game, game play and feedback loops are rapid, thus allowing multiple rounds to be played in a 60 minute session. This advance will enable students toget a sense of how the neural network learns over time by changing their strategy after receiving feedback. These multiple rounds would effectively simulate the feed-forward, evaluation and back-propagation processes in a neural network. The move to a fully online version of the game reduces barriers to implementation, and can even be implemented in remote online classrooms. The ease of setup and implementation was increased since no physical setup of the space or printing of input image cards, as well as no sophisticated software setup besides a browser window is necessary. The new game is also more flexible - new input images are readily available and can be easily added. Further, the online Contour to Classification is accessible to students around the world with low bandwidth overhead thus leading to greater scaling potential.

In terms of the game's potential to support student learning about neural networks and supervised learning, the visual nature of the compositing contour information into layers of abstraction is readily accessible to students. An analogy to how convolutional neural networks are used in image classification can be made thus providing a reference example of the application of neural networks in their lives. Game play features such as feedback on which contours lead to correct classifications and the revision of contours or of the selection of contours in subsequent rounds of game play may increase students' understanding of how channels get strengthened or weakened and how nodes adjust their algorithms for choosing contours to send on. Additionally, students may gain an understanding of the neural network's limit of understanding, where and how the classifications can go wrong, and how the neural network can get better through learning from errors.

A potential disadvantage of the game might be that it requires exactly 10 players, and might not be adaptive to variable class sizes. One way to mitigate this is to allow the network to adjust the number of layers, and neurons per layer, similar to what happens when training a real neural network. Another potential limitation is the requirement of internet connection to play the game synchronously with other participants. Our next steps are to play-test the Contour to Classification game with students in the fall and enable students to choose images or themes that they find relevant and interesting. For example, students may want to use poses from dance moves as input images. Further, we will conduct a study to determine if and how students learned through the experience of playing the Contour to Classification game.

Conclusion

In this paper, we describe the design of the Contours to Classification game that is designed to teach middle school students about Neural Networks. We describe the previous versions of the Neural Network game and its evolution to the Contours to Classification game.

Acknowledgments

*W*e are grateful to the teachers and students of { removed for anonymous review } programs for participating in our activities, and allowing us to pilot the curriculum. We are grateful to { removed for anonymous review } for helping us run the study and collect data. We are especially grateful to our research sponsors and { removed for anonymous review }.

References

Google. 2019. Teachable Machines. URL https://teachablemachine.withgoogle.com/.

Heroku Platform, S. 2007. Heroku. URL https://www.heroku.com/.

Irene Lee, F. M. 2019. What do you meme? : A participatory simulation for teaching about neural networks and machine learning. *2019 "Teaching AI in K-12" Symposium* "Activity".

National Academies of Sciences, E.; and Medicine. 2018. *How people learn II: Learners, contexts, and cultures*. National Academies Press.

Node.js. 2009. Node.js. URL https://nodejs.org/.

TechGirlz. 2018. A Look Inside The Human Neural Network Activity From Our 'Artificial Intelligence: How Computers Learn' TechShop. URL https://www.techgirlz.org/a-look-inside-the-human-neural-network-activity-from-our-artificial-intelligence-how-computers-learn-techshop/.