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 is a browser-based game that teaches middle school students basic concepts in supervised learning. The game is an online variant of the Neural Network game that was presented at AAAI Fall Symposium Teaching AI in K-12 track in 2019. We share preliminary findings from implementing the online version of the original Neural Network game in a pilot research study and describe the game's evolution to the Contour to Classification game. The new game uses a simulation of a neural network to engage students, through digital drawing and selection interactions, in the classification of images. The players act as nodes 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 prediction. After evaluating the prediction, 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 game, we aimed for students to learn the structure, components, and functioning of a neural network, and the processes involved in supervised learning. The Contour to Classification game supports online student learning by providing the image classification experience using purely visual inputs to each layer. 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 artificial neural networks (Hecht-Nielsen 1992). 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.

[*]Equal contribution.

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 (accessed December 17, 2020). 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 one's algorithm for choosing words. This live-action version called the "Artificial Neural Network game" was tested with educators at MIT, middle school students visiting MIT, and at AAAI 2019 Teaching AI in K12 Symposium.

In January 2020 the Artificial Neural Network game was adapted to be an online interactive game for use in online professional development workshop 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 teachers, after school club participants in a public school in the Boston area, and within "Developing AI Literacy" (DAILy) summer workshops for students held in 2020.

In the next section we will describe the research findings associated with the online Artificial Neural Network game before moving forward to describe the design of the new Contour to Classification game.

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 two AI summer workshops offered during the summer of 2020. Due to the COVID-19 pandemic, the workshops were offered virtually and students were instructed online on Zoom. The first workshop was offered over three weeks in July 2020, at two hours a day, by a team of researchers and graduate students involved in the DAILy project. The second workshop was offered over four weeks spanning mid July through August 2020, at one or two hours a day, by experienced computer science educators who served as the partner organization's summer instructional staff. Prior to teaching the DAILy curriculum, these instructors observed a workshop in June 2020 and had debriefed each day with the workshop staff.

Thirty-eight students who attended the workshop consented to participate in our research study. The consenting students were between the ages of 10 and 16. An equal number of females and males participated, the most recent grades in school completed were 6th, 7th and 9th graders (21%, 21% and 26% respectively. Noteably, the population of students reached was 89% from underrepresented groups in computing and STEM including female, Black, Hispanic/Latinx, and students representing two or more races/ethnicities including an underrepresented group. As a proxy for prior experience with computing, students were asked for their familiarity with "Scratch", the popular tool for introducing K-12 students to computing. Thirty-two percent reported having some familiarity with Scratch; 68% reported having no prior experience with Scratch.

As part of the study, an AI concept inventory was administered to consenting students before the start of the workshop and at the workshop's conclusion (Lee et al. 2021). Within the AI concept inventory, the neural network scale consisted of five of the forty multiple choice items. The analysis of student learning about neural networks centered around four questions: 1) Did students learn the names of the layers in a neural network? 2) Did students learn the ordering of processes in neural networks (two items addressed this question)?, 3) did students understand which processes contributed to a neural network's learning and 4) Did students understand in which phase (training or testing) learning takes place in neural networks?

In each of the workshops in this study, online Neural Network game was offered as a one-hour activity within the 30-hour Developing AI Literacy or "DAILy" curriculum (Lee et al. 2021). Prior to this activity, students engaged in exploring what is and is not AI, and considering stakeholders and their goals in the design of AI systems, as well as the ethics of AI. Students' first exposure to AI algorithms was through a unit on logic systems that featured decision trees as an example of a logic system. In a participatory simulation called "PastaLand" (Lee, Ali, and DiPaola 2020 (accessed December 17, 2020) students built, tested, and investigated their building process and the decision trees themselves for where biases were embedded. This decision tree activity was the students' first introduction to AI as a knowledge structure and process. Through the experience, students also gained familiarity with using Google Drawings as a platform for game play and collaborative interaction with peers online.

During a one hour session, the Neural Network game was introduced by a team of researchers and educators and collaboratively played within a Google Drawing followed by a whole group discussion. Participants were grouped by age into three groups of 10 or 11 individuals into small groups. 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' utterances and chat data from their discussion after the activity. Students also completed a daily reflection and were interviewed about their recollection of the activity at the end of the workshop.

Analysis of the neural network scale of the AI Concept Inventory showed an overall average gain of 0.329 (from 3.316 to 3.645) out of a maximum of 7 points but this difference was not statistically significant ($t(37)=1.742$, $p=.090$). Within the scale, only one item, the labeling of the layers in a neural network (input, hidden, and output layers), showed a statistically significant gain between the pre- and post-test ($t(37)=2.458$, $p=0.019$). The four other items (the stages in training a neural network; the actions that comprise learning; and the identification of the training stage as when learning takes place) showed small but insignificant gains. These results are not surprising given that the intervention was 1 hour in duration and suggest a longer exposure and multiple rounds of gameplay may be necessary for significant gains in learning to occur.

Despite limited neural network learning gains (as measured by the AI concept inventory), students' engagement and reactions to the neural network game are reasons to continue to develop the neural network activity. Students reported that they had fun playing the activity. In their responses to the post-test interview question "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."

> "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."

In end of day reflections submitted on the day the game was played students mentioned learning about neural networks. Twenty-two of 43 respondents mentioned neural networks when answering the open-ended 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 that they learned 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 .. 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."

Re-conceptualizing the Neural Network Game

While the game proved engaging and enjoyable to participants of all ages, notable difficulties were encountered. The goal of the game, to make a caption for an image, was confusing for some students. Selecting words and composing captions was not familiar to students and how the behavior related to a real-world process was obtuse. Further, students could not connect their activity of captioning the image to how information passes between layers in an actual neural network and were oblivious to the role of the nodes in other groups (or layers). There were usability concerns that we found in gameplay, such as, the need to move students acting as input nodes to a zoom breakout room to view the image, which was time consuming and broke the flow of the activity, and reduced the memorability of the input image. 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 activity and of creating captions for images as a target goal of the activity. A suggestion made by one of the expert panelists who reviewed our curriculum hinted at a possible redesign: ”To make the connection between Day 4 (neural networks) and Day 5 (image classification) a bit stronger, I wonder if the game should be to label an image (instead of caption)... Also consider a slight variation: instead of the hidden layer receiving words, it receives smaller crops of the image (aka puzzle pieces) which is a bit closer to what the actual convolution layers see. Then the hidden nodes label those pieces independently and pass the labels to the final layer. The final layer needs to output 1 label which is a summary of the individual labels."

In an effort to reconceive the goal of the activity to align with Convolutional Neural Networks (CNNs) used in image classification, Lee experimented with a paper prototype of the online Artificial Neural Network game in which players create partial outlines or contours of images from photographs as the information passed through the neural network rather than words. The paper prototype used thin masking tape pieces (each 2" in length) to outline contours on clear mylar sheets. The compositing of the mylar sheets with taped segments formed larger contours simulating the levels of abstraction in the deep learning process, as shown in Figure 1. The compositing of contours to create higher level representations of the image is shown in Figure 2.

Figure 1: Paper prototype of the Neural Network game

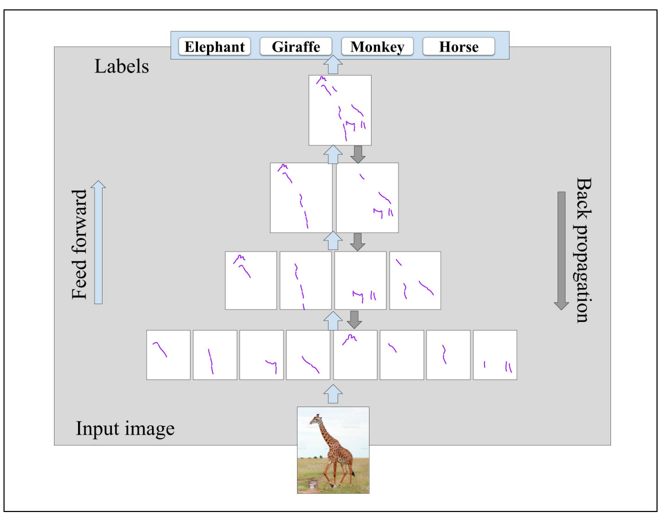

Figure 2: Image classification via composite line segments

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. Ali conceived of the neural network as having four layers: the input layer, the hidden layer 1, the hidden layer 2 and the output layer. Each input layer node sees the original image and draws 4 separate contours. They draw contour lines on a transparent canvas overlay on top of the input image with a limited number of pixels (30 pixels per contour by default). The 16 contours (4 from each input nodes) are distributed among the hidden layer 1 nodes (but the hidden layer nodes do not see the original image). Each hidden layer 1 node can composite the contours they receive and select which 2 of the 4 contours to pass along to the hidden layer 2 node. The hidden layer 2 node (a single node) receives 8 contours and can composite them before selecting which 4 of the eight contours to send to the output layer node. The output layer node views the composite of 4 contours and classifies it as representing one of 6 animals as shown in Figure 3. In short, the input nodes develop a strategy for drawing contours; the hidden layer nodes form a strategy for selecting which contours are sent on and which are discarded; and the output node comes up with a strategy for classifying the composite of the contours correctly.

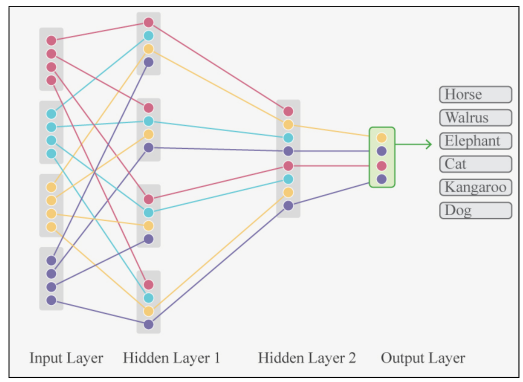

Figure 3: The Contour to Classification game schematic of the neural network

If the predicted classification is incorrect, all the nodes of the network receive an indication of which contours were selected and which were discarded, an indication of which links were strengthened and weakened on the minimap, and get a chance to re-strategize. If the classification is correct, the nodes receive feedback that their selections led to a successful classification and an indication of which links were strengthened, thus reinforcing their strategy. Players can play multiple sub-rounds until the network reaches the correct answer. While the strategies employed are likely to be somewhat random at first, the nodes in each layers can improve their strategies over time to send the most useful information forward. Once the neural network is successful, the players can play a new round wherein the input layer nodes receive a new input image. Participants stay in the same role and apply 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.

The Contour to Classification Game

Learning Objectives

In this game, 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 gain the following understandings from the lesson:

- Neural networks can be used for supervised learning.
- Neural networks consist of an input layer, one or more hidden layers, and an output layer.
- Each layer is made up of nodes.
- The nodes from one layer are connected to the nodes in the next layer through channels (or links) creating a network of nodes.
- Each link has a weight or ”goodness” associated with it.
- Neural networks use feed forward, evaluation and back propagation processes to tune the network over time which helps the neural network get better at classifying images (i.e. learning).
- The process we played out in the game, wherein the network was tuned, is called the “training phase.”
- 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 and connected to the internet. Mobile devices also work but are not ideal since they do not afford a large drawing canvas. Teachers can set up the game and students can play the game using any browser with a functional internet connection with no additional setup required.

Technical Details

The game itself is built for the web using HTML/CSS and Javascript and and uses a Node.js web server.

Prerequisite Knowledge

Typically, in a workshop setting, this game is played after students have learned about Decision Trees and have experience with Supervised Learning through engaging with Google’s Teachable Machine (Google 2019 (accessed December 17, 2020). If students learned about decision trees through playing the ”PastaLand” game (Lee, Ali, and DiPaola 2020 (accessed December 17, 2020), they already have familiarity with simulating a network’s behavior. If they’ve played with Teachable Machine they will have prior experience with training a model to classify images. 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 by the last person in line is compared with the original message. Often the message gets garbled along the way. Students may also have familiarity with networks through exposure to food webs and classification trees in a science class, paths connecting locations in board games, and diagrams in math class.

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. A 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 (or links) 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.

Players and Roles

The game was designed to be played by 6 to 11 participants and led by 1 teacher. On the opening screen, the game asks for the number of players as an input from the teacher, and then generates unique web links for all participants to take on different roles. If there are more than 11 students who want to play, the teacher can run multiple instances of the game and assign a subset of students to each of the instances. Based on the number of players entered by the teacher, the game generates the appropriate number of nodes and associated web links per layer. Between 2 and 5 web links are created for the nodes of the input layer; between 2 and 4 web links are created for the nodes in the hidden layer 1; 1 web link is created for the node in hidden layer 2; and 1 web link is created for the node in the output layer.

After the teacher has initiated a game, students log in, and select a role or node to play (picking from the ones remaining). The teacher, through a dashboard view, can select an image to be classified (or upload a new one), monitor the progress of the game, and move the game ahead if a player is not responding. In the following sections we describe the game as played by 10 players with 4 players as nodes in the input layer, 4 players as nodes in hidden layer 1, 1 player as a node in hidden layer 2, and 1 player as the output layer node.

The Interface and Game Play

Each player gets their own unique view of the game interface as each layer has access to different information.

Input Layer Node View and Play: All 4 input nodes see an input image (Figure 4). Each player can draw 4 contours with a maximum of 30 pixels each on top of the image. The contours appear on blank canvases (lower portion of Figure 4) that show the contours in isolation. These contours get distributed to the hidden layer 1 nodes.

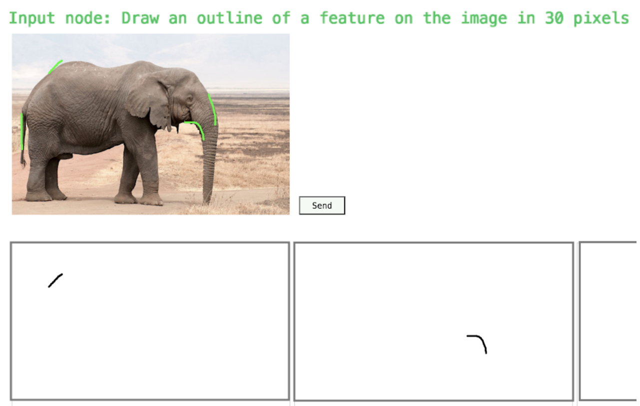

Figure 4: The Contour to Classification game input layer node view (cropped)

Hidden Layer 1 Node View and Play: Next, each of the Hidden layer 1 nodes 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, as shown in Figure 5. 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 the contours to the next hidden layer.

Hidden Layer 2 Node View and Play: The player in the hidden layer 2 receives 8 contours from the previous layer (2 from each node). The hidden layer 2 node can composite 4 of these contours to form a composite shape (Figure 6).

Output Layer Node View and Play: The output layer node receives a composite of the 4 contours and has to now classify what kind of animal is represented. The output layer

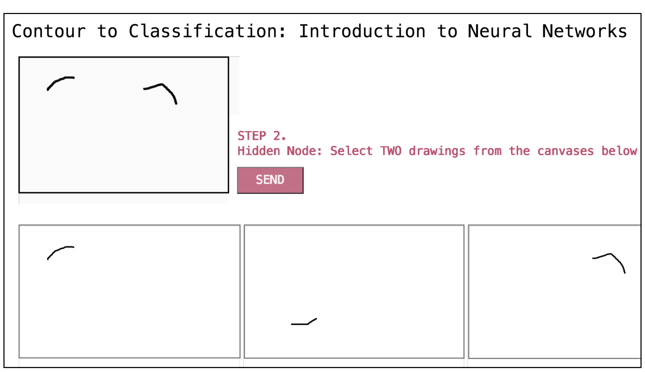

Figure 5: Hidden layer 1 forms a composite of 2 contours to send to hidden layer 2

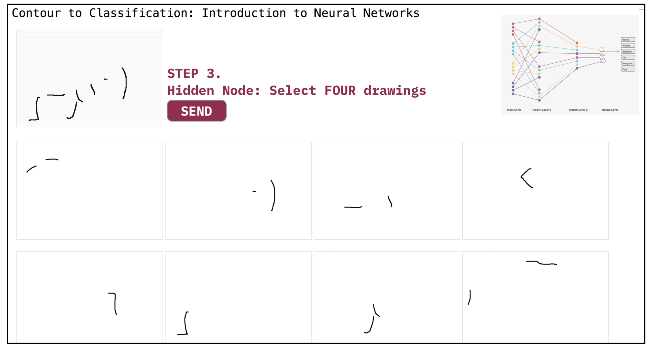

Figure 6: Hidden layer 2 received 8 contours and forms a composite of 4 contours to send to the output layer

node chooses from the 6 animal options provided. (The animal selected is called the prediction.)

Evaluation Phase: After the output node selects one of the 6 animal options, the prediction is evaluated for accuracy. If the selected animal matches the image label, the prediction is correct, otherwise the prediction is incorrect.

Back Propagation Phase: In the back propagation phase the hidden nodes and input nodes receive feedback based on the evaluation of the prediction. Information from each layer that reached the last layer is marked red (incorrect) or green (correct) based on whether the prediction was correct or incorrect (Figure 7). If the prediction was incorrect, 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 to mark the ears or trunk since they are more characteristic of elephants. The hidden nodes may strategize to select contours that are closer together.

Multiple Sub-round Play: The players can play additional sub-rounds to test their new strategies by drawing new contours (at the input layer) or sending new combinations of contours (at the hidden layers) to the output node. They keep repeating this process until the output node makes a prediction correctly (Figure 8).

Subsequent Full Round Play: Students can play additional rounds with new input images to test if their strategies work well for classifying other types of animals. The modification of strategies is described as analogous to the process of gradient descent. Through playing multiple rounds, input and hidden layer nodes can refine their strategies for sending information to the output node to help classify the image.

Discussion

The evolution of the design from the Artificial Neural Network game to the Contour to Classification game addresses barriers that hinder students from reaching our learning objectives and scalability. Drawing of contours to indicate salient features of the input images may be more accessible to young students than coming up with descriptive words. We removed the need for students to manually copy their descriptors (contours) and paste them into specific correct

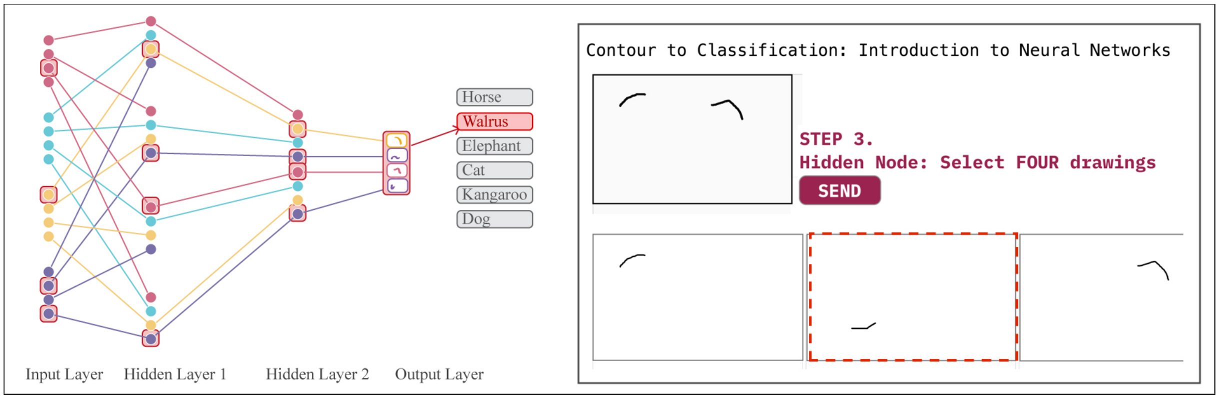

Figure 7: Evaluation and back propagation. The output layer makes an incorrect classification (left). Hidden layer 1 receives feedback indicating the contours that contributed to an incorrect prediction (right).

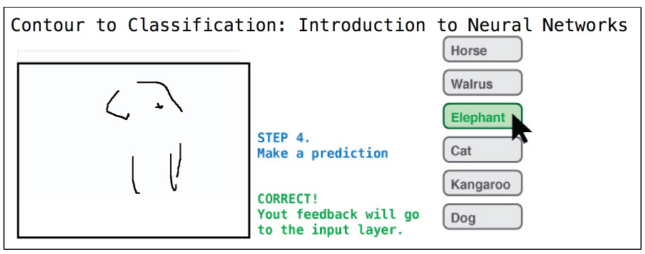

Figure 8: Players keep playing multiple sub-rounds until the output node classifies correctly

slots for delivery to subsequent nodes (as was required in the previous version of the game), 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 within a 60 minute session. We hope this advance will enable students to experience 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 training 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 with no sophisticated software setup. The new game is also more flexible than the previous version - new input images are readily available and can be easily added. Further, the online Contour to Classification game is accessible to students around the world due to its low bandwidth overhead thus it has greater potential to be used at scale.

In terms of the game's potential to support student learning about neural networks and supervised learning, the visual nature of the compositing of contour information into layers of abstraction is both readily accessible to students (they are familiar with overlays on images) and a better analogy to how convolutional neural networks are used in image classification. The game's goal of classifying images provides an example of the applications 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 or links 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 through trial and error.

A limitation of this design is the requirement of internet connection to play the game synchronously with other participants. While students use their knowledge of the world (latent dataset) to draw contours and predict classifications, this game does not explicitly address how large datasets are used to develop machine learning models. Whether or not we've created additional barriers through our redesign will be determined through field testing. Our next steps are to play-test the Contour to Classification game with students in the spring of 2021. We'd also like to gather students' suggestions of images or themes that they find relevant and interesting to classify. 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 about neural networks through the experience of playing the Contour to Classification game.

Conclusion

In this paper, we explore interactive and accessible ways to teach the concept of neural networks to middle school students. We describe our findings from play-testing a previous version of the Neural Network game and the game's evolution to the Contour to Classification game. We describe the design and gameplay of the Contour to Classification game. While the current version of the application is specifically designed to teach middle school students about neural networks, the contour drawing mechanism and multi-layer feedback simulation can serve as inspiration for educators and programmers for developing tools and interactions to teach learners about other machine learning algorithms such as Recurrent Neural Networks (or RNNs) (Schuster and Paliwal 1997). In future work, we will analyze the efficacy of the Contour to Classification game for its usability, students' engagement, and students' learning gains.

Acknowledgments

*W*e are grateful to the teachers and students of the STEM programs for participating in our activities. We are grateful to Daniel Smilkov for sharing his expertise and D. DiPaola, H. Zhang, Y. Cheng, and G. Kim for helping us collect data. This material is based upon work supported by the U.S. National Science Foundation under EAGER ITEST Grant Number 2022502. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Google. 2019 (accessed December 17, 2020). Teachable Machines. URL https://teachablemachine.withgoogle.com/.

Hecht-Nielsen, R. 1992. Theory of the backpropagation neural network. In *Neural networks for perception*, 65–93. Elsevier.

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".

Lee, I.; Ali, S.; and DiPaola, D. 2020 (accessed December 17, 2020). PastaLand: Introduction to Decision Trees. URL https://aieducation.mit.edu/daily/decisionTrees.html/.

Lee, I.; Ali, S.; Zhang, H.; DiPaola, D.; and Breazeal, C. 2021. Developing Middle School Students' AI Literacy. *52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21)* "In press".

National Academies of Sciences, E.; and Medicine. 2018. *How People Learn II: Learners, Contexts, and Cultures*. Washington, DC: The National Academies Press. ISBN 978-0-309-45964-8. URL https://www.nap.edu/catalog/24783/how-people-learn-ii-learners-contexts-and-cultures.

Schuster, M.; and Paliwal, K. K. 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing* 45(11): 2673–2681. doi:10.1109/78.650093.

TechGirlz. 2018 (accessed December 17, 2020). 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/.