PROJECT REPORTS



Learning by Enhancing Half-Baked AI Projects

Ken Kahn¹ · Niall Winters¹

Received: 27 July 2020 / Accepted: 14 June 2021 / Published online: 24 June 2021 © The Author(s) 2021

Abstract

We have developed thirty sample artificial intelligence (AI) programs in a form suitable for enhancement by non-expert programmers. The projects are implemented in the Snap! blocks language and can be run in modern web browsers. These projects have been designed to be modifiable by school students and have been iteratively developed with over 100 students. The projects involve speech synthesis, speech and image recognition, natural language processing, and deep machine learning. They illustrate a variety of AI capabilities, concepts, and techniques. The intent is to provide students with hands-on experience with AI programming so they come to understand the possibilities, problems, strengths, and weaknesses of AI today.

Keywords Project-based learning · Blocks programming · Snap! · Artificial intelligence · Machine learning

1 Introduction

One very effective way of acquiring a deep understanding and appreciation of AI is by building AI programs. However, building AI programs can be difficult and time-consuming. These shortcomings can be significantly reduced, however, by providing learners with high-level building blocks and associated guides [6]. Here we present an additional way to enable learners to experience the construction of AI programs despite a lack of experience and a limited amount of time.

The idea is to provide a range of half-baked or minimal AI programs designed to be enhanced by learners. This builds upon the work of using half-baked projects as resources in learning to program computers [7]. A related effort is the Machine Learning for Kids website [8] where tutorials are provided for constructing simple AI projects.

As part of the eCraft2Learn project [5] and subsequently, we developed several small AI projects. All the projects were built upon the Snap! [3] programming blocks designed to support AI programming [6]. The programs illustrate several different AI concepts, techniques, and capabilities. While

 Ken Kahn toontalk@gmail.com
Niall Winters niall.winters@education.ox.ac.uk each one is able to perform a simple task or two, they can be enhanced by non-expert programmers to be more capable. In doing so, we expect that by engaging in how the programs work to improve them, the students will learn about the abilities, strengths, and weaknesses of AI today. At least one or two layers of technology underlying AI programs will stop being magical black boxes to the students. We are currently at an early stage of evaluating these expectations with empirical studies.

Some of the projects illustrate ways of creatively using speech synthesis, speech and image recognition, natural language processing, and deep machine learning. Others explore innovative ways of applying pre-trained deep learning models to tasks. The variety of projects maximises the chance that a project will fit with a student's interests and goals.

2 Some Half-Baked Projects

The complete list of thirty projects can be found at the eCraft2Learn AI home page [4]. Descriptions of a sample of these projects follows.

2.1 Speech Commands

We have developed a variety of projects that illustrate using speech to control an application. Typically these commands control the movement of graphical sprites. The simplest

¹ Department of Education, University of Oxford, 15 Norham Gardens, Oxford OX2 6PY, UK

projects respond to single words or fixed phrases. One project creates the illusion that one can give full sentence commands. It works by searching for a keyword and a number somewhere in the sentence. E.g. "Could you please go forward 25 steps" is interpreted as the command "FORWARD 25". Several of these projects complement speech recognition with speech synthesis to support a voice-only interface.

The mechanism of speech recognition varies between projects. One project uses the Web Speech API [10] supported by several browsers. It is the most reliable and flexible method for recognising speech. This reliance upon a web service, however, prevents the recognition functionality from being more than a black box to students. Furthermore, it requires a fast reliable Internet connection and raises privacy issues. Another project instead relies upon Snap! blocks that we provide for training the system to recognise a handful of words. Yet another takes advantage of Google's Teachable Machine [2] where students can create and train audio recognition models that are then imported into their Snap! projects. Both of these approaches rely only upon the student's computer to train and run the model and hence protects their privacy.

When students train their own speech recognition models they are quickly exposed to the imperfections of recognition. If trained with samples only from one person it may not be very reliable with other speakers. They can discover that with more and varied samples the recogniser becomes more robust. If they increase the number of words their models can recognise they will discover the system's accuracy drops. Working with their trained model they become exposed to the notion of "confidence scores"that indicate the relative certainty of correctly classifying audio input. Finally, they need not limit their projects to speech but can explore the recognition of different sounds.

2.2 Interactive Sentence and Story Generators

Two projects use speech input and output to interactively fill-out a sentence or story template. Experience with older textual versions of this is that students discover that without careful attention to parts of speech and grammar rules ungrammatical sentences result. These projects give students first-hand experience with the construction of voice-only interfaces.

2.3 Training a System to Distinguish Between Different Categories of Images

Projects include a rock, paper, and scissors game and a drawing program that is driven by gestures. As with the speech recognition projects, the students encounter the need for a sufficient number of varied inputs for trained models to be accurate enough for their projects. They may acquire some degree of understanding how popular apps that recognise people, objects, and gestures work internally. And reasons why these apps make mistakes or exhibit biases.

2.4 Using Body Pose and Segmentation to Provide Augmented Reality Apps

This includes a project that detects if someone is touching his or her face and another where virtual balloons are popped with real video hands. These projects may contribute to an understanding of how gesture-based games such as those that use Microsoft's Kinect [12] work and how filters are added to images in apps such as Instagram and Snapchat.

2.5 An App that Generalises Mathematical Relationships Based on Examples

This is the simplest of the projects that rely upon creating, training and evaluating deep neural networks. Unlike projects that rely upon real-world data the goal here is to build a trained model that can approximate a mathematical relationship as simple as doubling or the more challenging square root. While enhancing this project students may learn the need for larger and deeper models for certain kinds of relationships. And how the number of samples and training cycles affects the quality of the predictions. Students may be surprised to discover that neural networks can quickly learn to approximate very well a relationship like doubling and yet fail to produce exact outputs.

2.6 A Program that Learns to Play Better Tic Tac Toe

This app is an example of how a neural network can be created, trained, and then used to guide game play. Learning to play a game typically involves more advanced machine learning methods such as deep reinforcement learning. As a much simpler alternative we provide a framework where logs of previous games are used to train a model to predict the probabilities of winning given the state of a game. The program makes a move based on the probabilities of winning from the boards that result from each possible move from the current state. One can engage with fundamental concepts such as the trade-off between exploiting one's current knowledge and exploring new things to increase one's knowledge. This project also illustrates issues in how to encode something like a game board into a list of numbers a neural network can process. One can learn the strengths and weaknesses of learning from self-play.

Note that by providing students with a functioning Tic Tac Toe implementation students can focus their energies and attention on getting the computer to play it well.

2.7 An App to Predict How One Rates Abstract Art

This app uses deep machine learning to predict how a user will rate randomly generated images. Each image is generated from a dozen random numbers and the program learns to associate lists of image generating numbers with user provided ratings. The trained model can then be used to "recommend" new images similar to how many popular recommender systems work.

2.8 Three Very Different Ways to Create a Question Answering App

One project fetches answers from web services, another extracts answers from passages of text, and the third one is trained to recognise paraphrases of questions with known answers. Experience enhancing any of them should help demystify conversational agents such as Alexa and Siri. Exploring three approaches to question answering should reveal their different strengths and weaknesses.

2.9 An App that Learns to Determine How Much Confidence is Revealed in a Text Passage

This is an example of emotion detection that relies upon sentence encoding. This simple example of sentiment analysis can be revised by students to categorise other sentiments such as worry, fear, optimism, etc..

2.10 An App that Learns to give Names to Random Colours

This app provides a simple example of a categorical deep learning classifier. Given three numbers corresponding to red, green, and blue it tries to match it to one of the colour names it was trained with. Students can learn how the number of colour names, number of samples, model architecture, and training regime influence accuracy.

2.11 An App that Transfers the Styles of Famous Artists to New Images

This app is a fun example of combining the learned style of an artist with a new image. It provides an example of AI generated art by mashing up two different kinds of inputs in a high-level manner.

2.12 Word Guessing Games that Rely upon Word Embeddings

Word embeddings map words into a high-dimensional space (300 in our implementation). Words that are closely related are close in this space and unrelated words are far apart.

nbeddings to provi

203

Here a very simple game uses word embeddings to provide clues to a player trying to guess a word. This and other uses of word and sentence embeddings expose students to the idea of distributed meanings. Other projects could be created to illustrate very different uses of word embeddings such as finding word analogies or doing translations.

3 Learning Resources

The eCraft2Learn website [4] includes programming guides, tutorials, and other resources for learning to use the Snap! AI blocks that underlie all the half-baked projects. There are programming guides for blocks for speaking, listening, seeing, natural language processing, neural networks, and accessing a variety of pre-trained deep neural networks. Each of these guides are interactive web pages containing several exercises implemented as embedded instances of Snap!. Each instance is loaded with relevant blocks and comments. There are also several interactive tutorials consisting of sample uses of the AI blocks interleaved with instructions for their use and exploration suggestions. Both the guides and tutorials cover all of the AI Snap! blocks used in the half-baked projects, as well as additional blocks students may use to enhance their projects.

4 Preliminary Results from Trials

We have introduced our half-baked projects to over a hundred students in several workshops. However, with only three hours for introducing the tools, libraries, and sample projects, little time is left for deep engagement or significant enhancements.

Students were proud of what they created even though they had made only a few enhancements to projects we provided them. In some cases, with young students (8–12 years old) in a 3 hour workshop, it is hard to imagine how they could complete a project in any other manner. A few of the projects were multiple-day efforts by older students (14–16 years old) who became very engaged in changing the seed project to make something that was authentically theirs.

Our half-baked projects have also been introduced to children from 6 to 16 in India and the Middle East in one-month long AI programming online "boot camps" offered by Clevered, an Indian company that offers professional programming and AI training. To date they have offered three boot camps with six to eight students each time. For each boot camp we provided feedback to the students when they proposed their projects and when they made final presentations.

While most of our half-baked projects have yet to be enhanced by students in our studies, a few that have include:

- 1. Taking a spoken command app and a pose detection app and creating a drawing program that responds to spoken commands to change the pen colour and line thickness. Drawing is controlled by moving one's hand in the air.
- 2. Another student changed the spoken command app into a spoken version of the classical textual adventure games.
- 3. Starting with an app that controlled the movement of a sprite to the left or right by pointing, a young student added two more directions.
- 4. A student enhanced the "name a colour" app so that when the app incorrectly names a colour it adds that instance to its training data to reduce the chance of similar mistakes in the future.
- 5. A student created an app based on the project that judges the level of confidence expressed in a sentence. Her app determined how bad someone was feeling based on what they wrote and then opened a meditation video depending upon the degree of unhappiness detected.
- 6. Several of the Clevered boot camp students created chatbots that enhanced projects that provided speech interfaces for Wikipedia or weather anywhere in the world.
- 7. Two groups of Clevered students used custom image classification to create an app that categorises items as recyclable or as trash.
- 8. Other Clevered students introduced simple gesture recognition into computer games that had developed building upon the *Rock Paper Scissors* sample project.

4.1 An Online Synchronous Course to Enhance the Don't Touch Your Face Project

We recently completed teaching a course consisting of 13 one-hour Skype sessions. The participating three young women from non-traditional academic backgrounds are being supported by the Go Girl Project [11]. They all had some experience with Scratch [9].

Each session consisted of the girls presenting their homework followed by one of the authors (Kahn) introducing them to some new technology and concepts followed by the students doing an exercise. The first half of the course involved a combination of a general introduction to AI and an introduction to the new AI blocks we have added to Snap!.

The Don't Touch Your Face app was inspired by donottouchyourface.com which uses a webcam to determine if the user is touching his or her face and if so issues a warning. Our project relies upon Posenet [1] which quickly reports the locations of up to 17 face and body parts. The half-baked version issues a warning when either wrist (hands are not tracked) comes close to either eye or the nose (the mouth isn't tracked). All the students fine-tuned the threshold for deciding what "close" means as well as the threshold for the minimum confidence score that a body part's location is known by the model. They changed the artwork for the hands, eyes, and nose. They changed the message when touching to include how many seconds since the start of the touch. They incorporated sounds that they recorded into the app. One student added a mouth, ears, and the rest of the face despite the fact that the system did not report their locations. To do so she made the mouth's location be an offset from where the nose is.

The students discovered that the app had been tuned to work well only when someone was the appropriate distance from the camera. Because the threshold was defined in terms of pixel distance it produced false positives when one is far from the camera and false negatives when close. The students followed a suggestion that they compute the threshold as the distance between the images of the eyes times a tunable parameter.

A surprising outcome happened after they enhanced the app to display the camera feed in the background. The sprites for the nose and eyes were displayed on top of the video. They drew funny eye glasses, earrings, and noses that became "filters" (Fig. 1).

The students were regularly able to creatively enhance the *Don't Touch your Face* project but only after guidance that usually consisted of demos of making a variety of enhancements. Rarely did they invent a project enhancement, instead



Fig. 1 The Don't Touch Your Face app turned into a "filters" app

they personalised example enhancements that we presented to them. Typically, their presentations of their enhancements indicated that they understood their programs beyond a superficial level.

5 Conclusion

We have designed thirty half-baked AI projects in the Snap! programming system. They were designed to be explored and enhanced by students. We hope to learn more about the pedagogical effectiveness of half-baked AI projects as we continue our research. We encourage others to make use of our AI learning resources in their teaching and research. It is open source and Creative Commons license. Over the coming years we hope to see many creative remixes of our projects. And we expect that the students who created these enhancements will, in the process, have acquired an appreciation and understanding of AI.

Funding The eCraft2Learn project is funded by the European Union's Horizon 2020 Coordination and Research and Innovation Action under Grant Agreement No 731345. The Go Girl project is funded by the University of Oxford IT Innovation Seed Fund and has received subsequent funding from Goldman Sachs Gives.

Availability of data and material https://github.com/ecraft2learn/ai

Declarations

Conflict of interest The authors declare that they have no confluct of interest.

Code availability https://github.com/ecraft2learn/ai

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source,

provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

- Google: Posenet (2021) https://github.com/tensorflow/tfjs-models/ tree/master/posenet. Last accessed 5 March, 2021
- Google: Teachable Machine (2021) https://teachablemachine. withgoogle.com/. Last accessed 5 March, 2021
- Harvey B, Garcia DD, Barnes T, Titterton N, Armendariz D, Segars L, Lemon E, Morris S, Paley, J (2013) Snap! (build your own blocks). In: Proceeding of the 44th ACM technical symposium on computer science education, pp. 759
- Kahn K (2021) eCraft2Learn AI home page (2021). https://ecraf t2learn.github.io/ai/. Last accessed 5 March
- Kahn K, Montero CS, Voigt C (2018) Steam learning in formal and informal settings via craft and maker projects. In: Proceedings of the 17th ACM conference on interaction Design and Children, pp. 728–733
- Kahn K, Winters N (2018) AI programming by children. In: Proceedings of constructionism 2018 conference. Constructionism 2018
- Kynigos C (2007) Half-baked logo microworlds as boundary objects in integrated design. Inform Educ-An Int J 6(2):335–359
- Lane D (2021) Machine learning for Kids (2021). https://machi nelearningforkids.co.uk/. Last accessed 5 March
- Maloney J, Resnick M, Rusk N, Silverman B, Eastmond E (2010) The Scratch programming language and environment. ACM Transac Comput Educ (TOCE) 10(4):1–15
- Mozilla: Web Speech API (2021) https://developer.mozilla.org/en-US/docs/Web/API/Web_Speech_API/Using_the_Web_Speech_ API. Last accessed 5 March, 2021
- University of Oxford: Go Girl (2021) https://gogirloxford.org/ about-the-project/. Last accessed 5 March, 2021
- 12. Wikipedia: Kinect (2021) https://en.wikipedia.org/wiki/Kinect. Last accessed 5 March, 2021