Shih, JL. et al. (Eds.) (2023). Proceedings of the 31st International Conference on Computers in Education. Asia-Pacific Society for Computers in Education

Democratising AI education: Teaching autoencoders to out-of-school children from low-income backgrounds

Saumay GARG[a*]
[a]*Department of Computer Science and Information Technology, Jaypee Institute of Information Technology, India*
*saumaygarg@gmail.com

Abstract: Artificial Intelligence (AI) has become an integral part of personal computing devices and is gaining importance in school curriculums. Several studies and tools have been developed for teaching AI to school-aged children and it continues to be a growing area of research in computer science education. However, most of these interventions are targeted towards the more privileged and therefore cannot be directly used with children from socio-economically backward families. Children from such backgrounds are often unable to complete even basic elementary education due to several social, economic and systemic challenges. This work aims to help extend the reach of AI education to these less privileged children. A learning activity based on the five AI4K12 big ideas is presented, the activity follows a series of sessions aided by an easy to create tool for supporting tinkerability. The activity attempts to give a high-level overview of the functioning of autoencoders. The feasibility of this approach is verified with a case study involving out-of-school children from low-income families. The findings show that all of the AI4K12 five big ideas can be addressed and effectively taught to non-English speaking children with low numeric literacy. The proposed approach was also shown to increase children's curiosity and sense of agency while learning about AI.

Keywords: Computational Thinking, Inclusive Education, Artificial Intelligence

1. Introduction

Artificial Intelligence (AI) has been accepted to be an integral part of computational thinking (Denning and Tedre, 2019). This is primarily because AI has become omnipresent with most computing devices relying on its applications in one way or another. However, AI can also have inherent biases that lead to discrimination towards marginalized groups (Angwin et al, 2016; Buolamwini et al, 2018). This omnipresence and potential for bias requires that children are taught about AI in schools in both a technical and sociological context.

Artificial Intelligence for K-12 (AI4K12) is a globally recognized initiative that supplies practitioners with standardized guidelines for AI education in schools. This initiative was started as a joint project by the Computer Science Teachers Association (CSTA) and the Association for the Advancement of Artificial Intelligence (AAAI). The AI4K12 guidelines are developed around the five big ideas for AI education. These big ideas broadly describe the goals for K-12 AI education while also serving as guiding principles for developing related resources and best practices (Touretzky et al, 2022). The 5 big ideas help ensure that the curriculum and resources focus on teaching all significant areas encompassing the AI field and not just a few superficial concepts. These five ideas are represented in Figure 1.

Smart phones have become extremely common in recent years, so much so that even marginalized and low-income families own smartphones now. According to a recent report, up to 88% of Indian households own a smartphone (Bhattacharya, 2022). And surprisingly, even children from these socio-economically backward backgrounds are extremely proficient in the usage of these tiny mobile computers. But these kids are not educated in the more technical functioning of computers and AI. Such children are very likely to be out of school, have never attended school or have subpar resources in their poorly funded schools. To put the magnitude of this problem in context, around 59.1 million primary school-aged children don't attend schools (UNESCO, 2019), primarily due to financial and social constraints. Consequently, these families continue to be stuck in an endless cycle of poverty and are unfortunately devoid of the opportunities that education and computer skills can bring. Lack of research into teaching practices for such populations makes the educational disparity even more concerning.

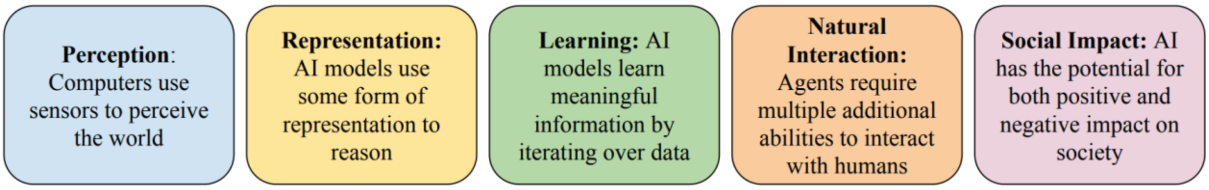

Figure 1. The five Big Ideas in AI4K12.

Research into teaching and learning of AI for K-12 has examined various best practices and studied the efficacy of teaching resources for AI. Some of these studies include works by Druga et al (2019), Long et al (2020) and Payne et al (2019). Though the progress in research for supporting AI literacy is promising, there are several unexplored yet significant research areas. The primary research gaps in AI education that this work addresses are as follows:

1) Though language as a barrier to computing education has been considered by Banerjee et al (2018), most of the research fails to take into account other socio-economic conditions relevant to marginalized children with little or no formal schooling.
2) Studies have shown that children develop a more nuanced understanding of AI when involved with actual code for AI agents and that scaffolding can lead to optimal learning (Druga et al, 2021; Hitron et al, 2019). However, what we don't know is how children from low-income backgrounds and lack of access to schools react to this exposure to code and what scaffolding might be required in such demographics.
3) In an in-depth review of resources for AI education by Druga et al (2022), it was found that most of these resources cover only a few of the 5 big ideas of AI4K12 and no existing resource covers all five.

2. Method

2.1 Participants

The participants for this study were out-of-school children in the age group of 9-13. These are children who have once attended school but have since left and don't plan to rejoin due to financial or social constraints. Though these children had attended schools till different grades, their numerical aptitudes were more or less similar. For example, an 11-year-old boy claimed to have attended school till fifth grade but was unable to perform single digit additions. This can be attributed to poor instruction in previous schools and frequent absenteeism on the student's part. Therefore, the children in this study were essentially at similar educational levels. The study participants were the children of migrant workers working low-paying jobs like domestic help and manual labour in India. They were selected primarily on the basis of their socioeconomic background. The total number of children who participated in the study were ten and all except two showed basic numeric literacy (the ability to count and recognize numbers from zero to ten) or had some understanding of smartphones (experience using YouTube and Google assistant features). None had any English proficiency and all had Hindi as their native language. Six of the ten children identified as male and the remaining as female. The proposed activity was conducted in a single sitting with all sessions taking place one after another, taking about 50 minutes to complete on average.

2.2 Supporting Learning Resource

The teaching resource used in this study is very simple to develop and can easily be recreated even without much technical proficiency. The autoencoder and the associated interface was developed in the Google Colab, a cloud hosted development environment for python and machine learning. An Autoencoder (Hinton et al, 1989) trained on the MNIST dataset to generate hand-written digits was programmed using TensorFlow. This code, however, is unlikely to be understood by the child participants because 1) it requires proficiency in English and the target demographic for this study is non-English speaking and 2) codes for neural networks are not trivial.

The exact architecture of the autoencoder is not introduced to the participants for the sake of simplicity. This study's focus is on teaching the broad concept of encoders, decoders and latent space representation. To allow the children to interact and tinker with the code, variables for the latent feature vector (सीखने की क्षमता) and number of epochs (इतनी बार दोहराए) in the code are parameterized with the forms feature in Colab. The value of the latent feature parameter can be adjusted by choosing a value from a dropdown list with three values: 0, 3, 6. Setting it to zero would render the model useless, this can be understood as the neural network not creating any representation of the input and thereby not learning anything. Higher dimensions of this vector means that it is retaining more information to reconstruct images from. Similarly, the parameter for the number of epochs is also presented as a dropdown menu and the training process is visualized with a progress bar. Finally, an option to supply the trained neural network with user input is provided. These elements are illustrated in Figure 2.

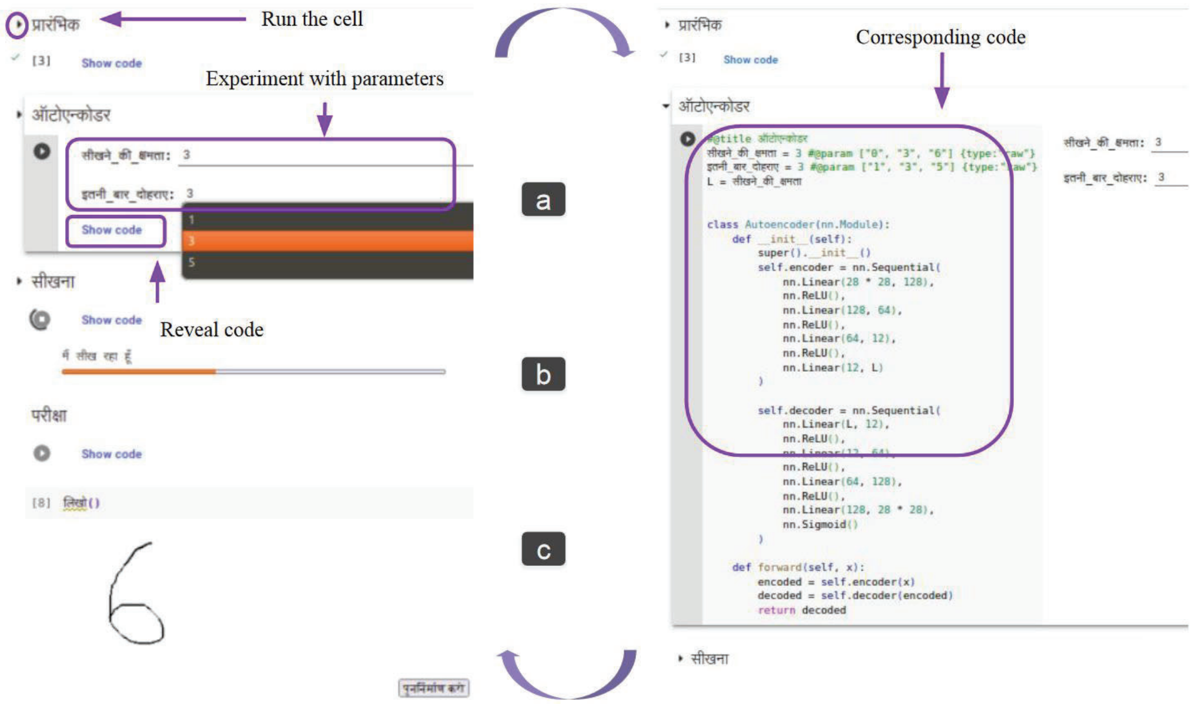

Figure 2. Users can run a cell, tinker with parameters and view the corresponding code a) dropdown menu for parameterized variables in the code b) progress bar to visualize the training process of the neural network c) drawing board for user input.

2.3 Study Procedure

The study introduced autoencoders to the children in a series of logically segregated sessions aligned with the AI4K12 big ideas (Figure 3). These sessions help structure the learning process, thereby simplifying both the learning and instruction. Also, they can be held over a course of multiple days which can be used as a scaffolding technique. However, in this study, these sessions were held consecutively on the same day. The 3 sessions are described below:

Session 1 (Develop analogies): The session starts with a general introduction of computing devices using the example of smartphones as these children are most likely to have access to these devices. This discussion serves two purposes, it helps the instructor gauge the aptitude of children with respect to computers and clarify any related misconceptions. Children are then introduced to Big Ideas 1 (perception), 2 (representation) and 3 (learning) using analogies like the ones described in Figure 3 (a). The experience of learning in a classroom setting was used because the participants found the scenario relatable.

Session 2 (Introducing code and interface): The personal analogies form the basis for developing an intuitive understanding of the functioning of an autoencoder. The first 3 big ideas are sequentially introduced to the children with support from the analogies established in session 1. This progression is apparent in Figure 3 (b). The idea is to help children see the neural network as a learning entity like themselves; this is pretty much like applying the theory of mind (Ensink et al, 2010) for the computer. This process takes place while discussing the Colab notebook containing the code and interface elements. Showing code helps distinguish the AI model as a distinctly human made entity, similar yet very different to humans. Exposure to code also aids the development of a more nuanced understanding of AI (Druga et al, 2021).

The general flow of interaction with the interface is illustrated in Figure 2. Initially the child uses the run button to run a cell containing a piece of code, completion of its execution is reflected by a green tick mark. Then, the child is encouraged to experiment with the model by altering the given parameters in the interface. Once a familiarity is established, the “show code” option is used to reveal the corresponding code to the child. This procedure can repeat itself multiple times depending upon the number of cells and exposed parameters in the notebook. In addition to this, the children are allowed to test the performance of the model on custom inputs for different latent space sizes.

Session 3 (Usability and impact): AI4K12’s fourth (natural interaction) and fifth (societal impact) big ideas are discussed in the final session. First goal of this session is to elicit a discussion on how the deep learning model is intelligent but unusable for any practical application involving other people. Second, a more open-ended discussion on the potential usage of such models for social good caters to the requirements of the fifth big idea. This discussion naturally segues into potential biases of such models.

3. Analysis and Findings

The evaluation is done based on perceptions observed during one-one sessions and interviews. Efficacy of instruction was gauged based on the child’s ability to understand and relate to the presented analogies, their understanding of the presented neural network and its visible parameters, comprehend the associated social implications, and suggest novel use cases for the AI model. A qualitative analysis led to the following key findings.

A simple form-like interface can improve tinkerability and inquiry of neural networks: Since the children recruited for this study were not fluent in English, they found it difficult to tinker with the program (like by changing the number of epochs) as they couldn’t read through the code to find the variable to alter. The supporting interface proposed above removes this need to understand the code and reliance on English for tinkering with the neural network. Children in this study could easily find the relevant parameters when provided as interface elements in their native language. Since the options for such parameters were limited by a dropdown menu, children found it easier to draw meaningful insights from their inquiry into the functioning of neural networks. This also prevented them from feeling overwhelmed as the number of options presented was carefully limited such that only valid permutations could be formed.

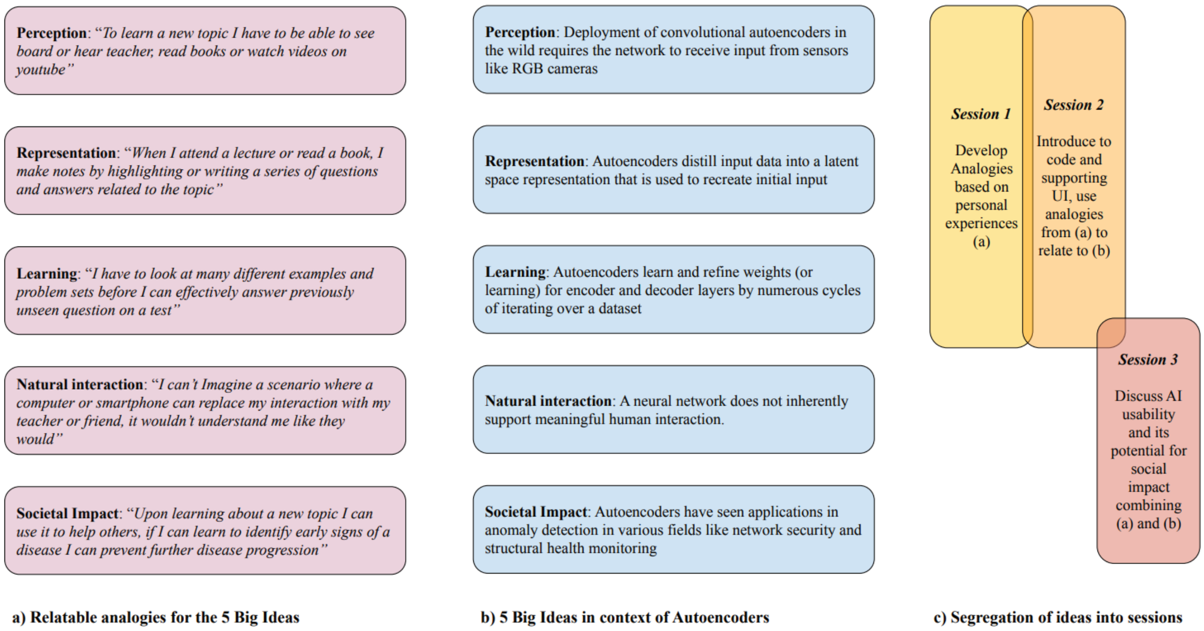

Figure 3. The overall progression of the proposed activity a) sample analogies developed around the AI4K12 big ideas b) equivalent explanations relating the analogies to the actual functioning of autoencoders c) segregation of the five big ideas into a series of sessions.

Showing code can elicit a sense of agency: Children could easily interact with the notebook using the interface elements but this alone was not sufficient to elicit ownership and a thorough understanding. The children were unable to think of the AI agent as a programmable entity when no code was shown. They simply saw it as another black box application which they merely use and not actively contribute to (despite them tinkering with the code parameters). However, when code was shown and it was pointed out how the code changes when they choose values from the dropdown menu, the kids felt more in control. This encouraged them to try new combinations of values and feel more confident in dealing with the deep learning notebook.

Analogies and discussion can help accommodate most AI education goals: Though the learning resource presented here did not support all five big ideas, the three sessions proved to be efficacious in incorporating all the five goals. When analogies related to learning were established before introducing the code and the interface, the participants' understanding of autoencoders seemed robust and well-informed. The clear progression of the three proposed sessions allowed for a smooth transition from one idea to another.

Instruction would have to take place individually: During the group instruction, significant challenges were faced that prevented the sessions from being successful. It was difficult to engage all the kids in even smaller groups as they tended to get distracted by their environments. For example, in many instances, the child was called by their parents to help with chores like fetching water or collecting donations. Also, unlike a regular classroom, the children in the group had significantly different aptitudes due to differences in prior experience. These differences made it difficult for some children to focus or effectively participate. Some such children also acted out and prevented others from engaging. However, they were able to comprehend and contribute effectively in individual instruction. Individual instruction also helped deal with the unique needs of these children.

Basic numerical literacy and experience with modern computers is the bottleneck: Children who did not possess basic numeracy skills were unable to make sense of the second session. Basic numeric literacy is required to tinker with parameters and make sense of the task performed by the neural network. These kids seemed disengaged even in individual interactions, but were able to understand and meaningfully engage in the first and third sessions. A similar response was observed with children who had no prior experience with computers, they couldn't grasp the idea of AI as they had no prior experience. Such children couldn't proceed past the first session. The participants who met these two criteria sufficiently understood the high-level functioning of autoencoders and were able to generalize their functionality to other generative tasks like reconstructing alphabets or sentences.

4. Conclusions, Limitations and Future Work

After a case study with ten out-of-school children with minimal numeric literacy and no English language proficiency, it was found that AI concepts can be meaningfully understood by such children irrespective of other educational differences. Use of analogy driven discussions was shown to effectively cover all five big ideas described by AI4K12 when instruction took place individually. Despite being unable to understand code, exposure to code made children feel more in control when this exposure was mediated with a simple form driven interface. These interfaces are extremely scalable and can be easily created with little technical expertise, thus allowing teachers or volunteers lacking in appropriate training to effectively be able to use such tools in their instruction. Nonetheless, this study was limited by its short duration and small sample size. Future research could include longer interventions and comparative studies with control groups to strengthen the generalisability of these findings. Methods for effectively discussing the AI4K12 big ideas in greater detail should be pursued while also adapting and refining the existing progression charts and policies to cater to out-of-school children. Perceptions and challenges in translating the proposed approach for teaching neural networks in greater depth and breadth also remain to be explored.

References

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2022). Machine bias. In Ethics of data and analytics (pp. 254-264). Auerbach Publications.

Banerjee, R., Liu, L., Sobel, K., Pitt, C., Lee, K. J., Wang, M., ... & Popovic, Z. (2018). Empowering families facing english literacy challenges to jointly engage in computer programming. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (pp. 1-13).

Bhattacharya, Pramit. 88% of households in India have a mobile phone (2016). https://tinyurl.com/3988y394. Accessed July 14, 2022.

Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In Conference on fairness, accountability and transparency (pp. 77-91).

Denning, P. J., & Tedre, M. (2019). Computational thinking. Mit Press.

Druga, S., & Ko, A. J. (2021). How do children's perceptions of machine intelligence change when training and coding smart programs?. In Interaction design and children (pp. 49-61).

Druga, S., Otero, N., & Ko, A. J. (2022, July). The landscape of teaching resources for ai education. In Proceedings of the 27th ACM Conference on on Innovation and Technology in Computer Science Education Vol. 1 (pp. 96-102).

Druga, S., Vu, S. T., Likhith, E., & Qiu, T. (2019). Inclusive AI literacy for kids around the world. In Proceedings of FabLearn 2019 (pp. 104-111).

Ensink, K., & Mayes, L. C. (2010). The development of mentalisation in children from a theory of mind perspective. Psychoanalytic Inquiry, 30(4), 301-337.

Hinton, G. E., & Sejnowski, T. J. (1986). Learning and relearning in Boltzmann machines. Parallel distributed processing: Explorations in the microstructure of cognition, 1(282-317), 2.

Hitron, T., Orlev, Y., Wald, I., Shamir, A., Erel, H., & Zuckerman, O. (2019, May). Can children understand machine learning concepts? The effect of uncovering black boxes. In Proceedings of the 2019 CHI conference on human factors in computing systems (pp. 1-11).

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In Proceedings of the 2020 CHI conference on human factors in computing systems (pp. 1-16).

Payne, B. H. (2019). An ethics of artificial intelligence curriculum for middle school students. MIT Media Lab Personal Robots Group. Retrieved Oct, 10, 2019.

Touretzky, D., Gardner-McCune, C., & Seehorn, D. (2022). Machine Learning and the Five Big Ideas in AI. International Journal of Artificial Intelligence in Education, 1-34.

UNESCO, New Methodology Shows that 258 Million Children, Adolescents and Youth Are Out of School (2019). UNESCO Institute for Statistics Database.