The Landscape of Teaching Resources for AI Education

Stefania Druga*
Information School, University of Washington
Seattle, Washington, United States
st3f@uw.edu

Nancy Otero*
Kitco
San Francisco, California, United States
nancy.otero.o@gmail.com

Amy J. Ko
The Information School, University of Washington
Seattle, Washington, United States
ajko@uw.edu

ABSTRACT

Artificial Intelligence (AI) educational resources such as training tools, interactive demos, and dedicated curriculum are increasingly popular among educators and learners. While prior work has examined pedagogies for promoting AI literacy, it has yet to examine how well technology resources support these pedagogies. To address this gap, we conducted a systematic analysis of existing online resources for AI education, investigating what learning and teaching affordances these resources have to support AI education. We used the Technological Pedagogical Content Knowledge (TPACK) framework to analyze a final corpus of 50 AI resources. We found that most resources support active learning, have digital or physical dependencies, do not include all the five big ideas defined by AI4K12 guidelines, and do not offer built-in support for assessment or feedback. Teaching guides are hard to find or require technical knowledge. Based on our findings, we propose that future AI curricula move from singular activities and demos to more holistic designs that include support, guidance, and flexibility for how AI technology, concepts, and pedagogy play out in the classroom.

CCS CONCEPTS

• **Applied computing → Interactive learning environments**; • **Social and professional topics → Children**.

KEYWORDS

AI education, K12, Teaching Support

ACM Reference Format:
Stefania Druga*, Nancy Otero*, and Amy J. Ko. 2022. The Landscape of Teaching Resources for AI Education. In *Proceedings of the 27th ACM Conf. on Innovation and Technology in Computer Science Education Vol. 1 (ITiCSE 2022), July 8–13, 2022, Dublin, Ireland*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3502718.3524782

1 INTRODUCTION

Modern computing is rapidly embracing artificial intelligence (AI) for it's great promise in improving our lives via advances in digital voice assistants, AI supported learning and increased accessibility [8, 17, 18]. However, AI systems can also amplify bias, sexism, racism, and other forms of discrimination, particularly for those in marginalized communities [1, 2]. In this context, promoting both technical and sociotechnical literacy of AI in primary and secondary education is critical [6, 13, 16, 24].

How to achieve this, however, is still an open question. Explorations of AI applications in education are challenging since the mechanisms and opportunities of AI are unfamiliar to most people outside of computer science. AI education is considered a vital part of computational thinking [4, 24], and there are arguments to include AI literacy in the primary and secondary education CS curricula [6, 14, 15]. Some works have begun to systematize competencies and skills for AI literacy [13].

One part of achieving AI literacy is the creation of technology resources to facilitate learning and teaching. For example, dedicated coding platforms such as Cognimates[1] and Machine Learning for Kids[2] have emerged to enable AI learning. Organizations like AI4All[3] have also created a free AI curriculum for secondary students. These technologies and their designs matter [12] as they shape and constrain what content knowledge can be taught. Educators must understand and appropriate AI resources to integrate them into their practice [20].

Despite the proliferation of AI education, prior work has only begun to examine its efficacy and appropriateness for primary and secondary teaching and learning. For example, studies have recently found that whether data is personal can influence student learning [16], that AI curriculum needs to be adapted to different cultural references and languages to become more inclusive [6, 26], that children become more skeptical of machine intelligence if they engage in active training and coding with AI [5], that carefully designed scaffolding is key to learning and transfer of knowledge [9], that gaps in access to technological resources and appropriate infrastructure, especially in the global south, can prevent learning from happening at all [26], and that teaching machine learning differs from teaching computer science as it is not "rule-based" [23].

While prior work has begun to reveal the pedagogies necessary for AI literacy, no prior work has examined the technological resources necessary to support these pedagogies. Prior studies have focused on more narrow aspects of machine learning learning ressources, either by analyzing visual tools for teaching machine learning in K-12 [27] or by doing a systematic review of research efforts on AI education[28]. For our analysis we choose to analyse how existing AI resources support pedagogical efforts and teachers. Therefore, we asked: *What learning and teaching affordances do existing AI resources have for supporting teaching AI?* To answer this, we conducted a systematic analysis of 50 AI resources curated from the most popular AI Education communities in North America: the

ITiCSE 2022, July 8–13, 2022, Dublin, Ireland.

ACM ISBN 978-1-4503-9201-3/22/07.
https://doi.org/10.1145/3502718.3524782

[1] http://www.cognimates.me
[2] https://machinelearningforkids.co.uk
[3] https://ai-4-all.org/

AI4K12 repository[4], the CSTA repository[5], the MIT AI Education repository[6].

Building on the Technological Pedagogical Content Knowledge (TPACK) framework [12], we formulated a series of questions and criteria to identify the extent to which current AI learning resources offer the support that educators might need. Overall, we found that AI resources broadly do not consider educators' needs to adapt and customize them for pedagogical use. In the rest of this paper, we elaborate on these findings in detail and discuss implications for design.

2 METHOD

To answer our question, we analyze a corpus of resources that could be used for AI learning. This mirrored prior corpus of studies of learning technologies, such as those considering coding tutorials [11] and programming environments for novice programmers more broadly [10]. Our focus is on resources that explicitly engage AI concepts relevant to AI literacy, including those not necessarily designed to be learning technologies.

2.1 Inclusion and Exclusion Criteria

To obtain a corpus of AI resources, we focused on curated lists of resources recommended for primary and secondary educators in North America: the AI4K12 repository[7], the CSTA repository[8], the MIT AI Education repository[9]. From these lists we considered only: curriculum materials, demos, list of links, online course, and software packages.

Based on these lists, the first two authors gathered an initial corpus of 100 resources. They then identified a subset of resources that were still available and functional and removed all duplicated entries, reducing the set to a total of 50 demos, interactive activities, tools, and curricula. The final corpus of 50 AI Education resources together with our final analysis is available here tinyurl.com/aiedk12.

2.2 Theoretical Framework

Since our research question focused specifically on teaching and learning concerns, we developed our framing based on theories that would make salient varying levels of support for teaching and learning. Our primary frame was the Technological Pedagogical Content Knowledge framework (TPACK) [12]. Building upon Shulman's Pedagogical Content Knowledge framework (PCK) [22], which posited the existence of knowledge of how to teach particular content knowledge, TPACK makes a similar claim. TPACK analyzes the existence of teacher knowledge of how to use technology (TK), how to use technology to teach (TPK), how technology and content influence and constrain each other (TCK), and how to use technology to teach particular content (TPACK).

We specifically used the TPACK definition proposed by Cox for our investigation, which synthesizes 89 other definitions. Her definition describes TPACK as five connected facets of teacher knowledge: "(1) the use of appropriate technology (2) in a particular content area (3) as part of a pedagogical strategy (4) within a given educational context (5) to develop students' knowledge of a particular topic or meet an educational objective or student need" (p.65) [3]. Each facet describes what a teacher needs to know about technology to use it for teaching and learning.

For the content knowledge dimension of our TPACK framework, we used the AI4K12 guidelines[10], which at the time of this writing defined five "big ideas" about artificial intelligence: 1) Perception: *computers perceive the world with sensors*, 2) Representation & Reasoning: *agents maintain representations of the world and use them for Reasoning*, 3) Learning: *computers can learn from data*, 4) Natural Interaction: *intelligent agents require many kinds of knowledge to interact naturally with humans*, and 5) Social Impact: *AI can impact society in both positive and negative ways*. These ideas provide structure for analyzing the kinds of content knowledge that resources can feasibly help students learn.

While the above TPACK framework is not necessarily theoretical, it derives from particular theoretical traditions that view teachers as pedagogical experts who develop content and technological knowledge to facilitate student learning [19]. While we acknowledge other more sociocultural [21] and sociopolitical teaching theories [7], our specific focus here is on educators' cognitive and pedagogical needs in their AI teaching practice.

2.3 Analysis

Our analysis built on the definition by Cox [3] by devising guiding analysis questions for each of its five facets, leading to 20 questions that structured our systematic evaluation of each resource. For example, one of our questions was "What types of pedagogical strategies does the tool support?" with fixed potential answers (i.e., "interactive learning", "direct instruction", and "hybrid between direct instruction and interactive learning"). The complete listing of these questions is available at tinyurl.com/aiedk12. Both first two authors collaborated on answering these 20 questions for each resource, resulting in a large spreadsheet with labels for each of the five facets of existing teacher support. Any disagreements in answering the questions were discussed until consensus was reached.

3 RESULTS

Overall, there were many distinct genres of resources by various creators: 39% were curriculum collections, 27% were single activities, 18% were demos, and 16% were tools. Only 20% were behind a paywall, though some of the more extended curricula offerings had a prohibitive price (i.e., ReadyAI charged more than USD 2.5k, TeensinAI charged more than 2.5k€). In this section, we evaluate the different genres of existing AI resources concerning how well they support teaching AI.

3.1 Communication of Intended Use

We considered the first facet of TPACK educators' need to know what technology is "appropriate" for a given student and learning goal. Therefore, we examined what kinds of information educators might need to judge the appropriateness of analyzing resources. A critical piece of information was the *intended use* of a resource, which illustrates the resource designers' assumptions about users'

[4]https://ai4k12.org/resources/list-of-resources/
[5]https://www.csteachers.org/page/resources-for-virtual-teaching
[6]https://raise.mit.edu
[7]https://ai4k12.org/resources/list-of-resources/
[8]https://www.csteachers.org/page/resources-for-virtual-teaching
[9]https://raise.mit.edu
[10]https://ai4k12.org

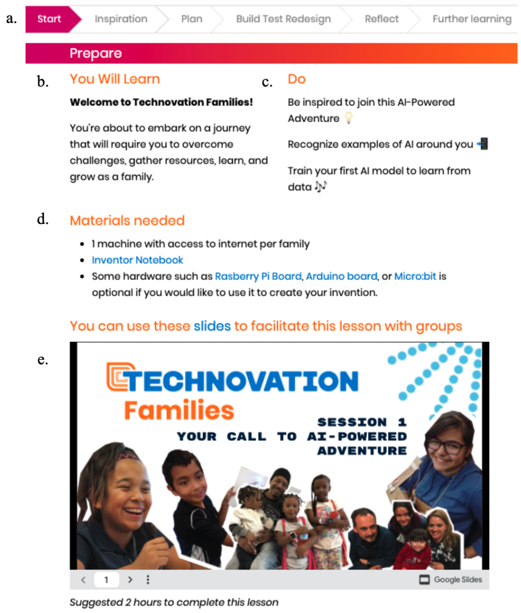

Figure 1: Curiosity Machine offered clear guidance to teachers about appropriate use, including: a) clear curriculum progression, b) learning goals, c) activity overview, d) materials description, e) and teaching materials.

prior knowledge and context. To analyze resources' intended use, we asked questions such as: "does the resource provide teaching guides?" and "does it provide explanations of the AI concepts it demonstrates?"

Teaching guides were one way to articulate intended use. Overall, we found that 59% of resources offered them. However, some teaching guides were minimal; for example, Zhorai[11] provided brief descriptions of "moderator" and "student" roles without grounding AI concepts and activities in existing curricular standards and practices. In contrast, platforms such as AI4ALL and Curiosity Machine[12] (shown in Figure 1) offered clear guidance for educators across several pedagogical dimensions, including learning objectives, pedagogical demonstrations, and materials required.

Another indicator of appropriate use was prior knowledge required to engage a resource. For example, 36% of the resources required users to perform an initial setup before testing or using the AI activity. Many of these setup requirements implicitly assumed particular content knowledge (i.e., terminal use, version control knowledge), with no guidance on how to acquire it. Similarly, while many resources were framed as learning materials—69% offered some written explanations of AI concepts—many explanations were not on the main page of the activity. Still, they were found in other locations like GitHub repositories, further obscuring whether the resource was intended for teaching and learning.

Trends in the clarity of intended use were primarily shaped by the genre of the resources. **Demos**, were often designed to emphasize one or more components of AI functionality, not to

[11]http://zhorai.csail.mit.edu
[12]https://www.curiositymachine.org

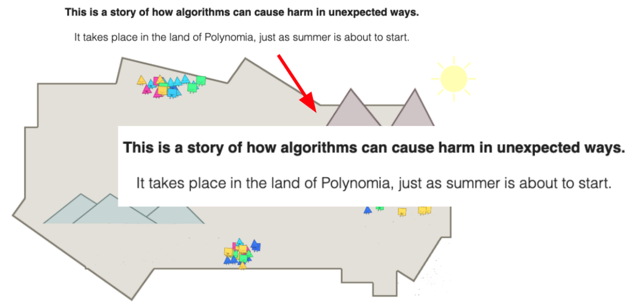

Figure 2: The Supervised Polygon activity creatively demonstrated unintended consequences of machine perception.

teach a comprehensive understanding of AI. None of the demos had teaching guides, only 50% of them explained the AI concepts they were addressing, and just 20% of them allowed participants to change the demonstration's output by modifying either the input data or the parameters. For example, TensorFlow Neural Network Playground[13] (Figure 1c) demonstrated how modifying different neural network parameters could lead to different outcomes. This resource offered a separate blog post explaining neural networks but did not integrate the explanation into the experience.

Activities were similar to demos, but often applied AI without a particular teaching goal. Only 30% of activities included teaching guides, only 50% of them explained how a part of AI works, and 62% allowed users to customize their creations. For example, Doodle Bot[14] was an activity for building a bot that uses speech commands to tell a bot what to draw. The activity listed instructions for building the bot and training the AI model with just one paragraph of AI explanation which mentions the pre-trained models used by the system (i.e., "ml5.soundClassifier()").

Tools gave even less direction for use. They offered platforms for creating new artifacts. Just two of the tools had teaching guides, but 75% included explanations of how AI works. Cognimates is an example of a tool that could be used to program interactive games using AI by training models to recognize specific images or text. It provided explanations of what algorithms were used to train the models.

Curricula were the clearest about their intended use offering explicit learning progressions for learners. All curricula included teaching guides, AI explanations, and 63% of them included a fixed progressive trajectory. For example, AppsForGood [15] had 14 sequential teaching sessions covering topics from what machine learning is to highlighting careers in machine learning. Of the curricula, 84% used both active learning and direct instruction. AI+Ethics curriculum included several activities that explored ethical questions and AI by doing projects as well as slides that teachers can use to explain AI concepts such as supervised machine learning.

[13]https://playground.tensorflow.org/
[14]https://mitmedialab.github.io/doodlebot
[15]https://www.appsforgood.org/courses/machine-learning

3.2 Big Ideas Coverage

The second facet of TPACK is content-specificity: teachers' knowledge of technology must be linked to the specific content knowledge they are teaching. Therefore, we examined the extent to which each resource covered the five AI4K12 big ideas [24].

Resources varied widely in their coverage. Most covered more than one big idea (88%), and most (72%) covered *Perception*. Some, typically curricula, covered all five (24%). The second most prevalent combination of coverage were resources that covered *Perception*, *Representation & Reason*, and *Learning* (18%). These resources were creative tools that typically allowed participants to input sound, images, or video, change the model's parameter, and get an output that showcases how a specific AI algorithm works. These resources typically covered supervised learning and training (28%), neural networks (20%), GANs (12%), image classification (8%), and word embeddings (4%). *Social Impact* was the least common, present only 2% of resources, typically in full curriculum or specialized activities on that topic.

AI big idea coverage varied by genre. For example, **demos** varied substantially in their coverage: 80% covered *Perception*, none covered *Social Impact*, half of the demos covered two ideas, and 30% had just one idea (*Perception* or *Learning*). One example was Pix2Pix which was a website that modifies a picture in real time based on drawings made by the learner[16]. Half of the demos covered two ideas, for example Scrooby, a website that enables participants train a cartoon based on movement perceived on the webcam. For one of the demos, Art Climate Change[17], it was not clear which AI big idea was present. Half of the demos had an explanation of the big ideas they covered.

Most **activities** focused on *Perception*. For example, Supervised Polygons[18], as shown in Figure 2, creatively used data on polygons' shapes (*Perception*) to illustrate AI concepts with unintended consequences (*Social Impact*). Most (84%) also focused on *Learning*; for example, PlushPal used data from the movement of a microbit to train a sound model. Half (53%) explained concepts; for instance, in FarmBeats, learners could use AI to optimize their farms and directly referenced AI4K12 big ideas.

Tools tended to cover at least three of the big ideas, most often *Perception*, *Learning*, and *Representation & Reasoning*. For example, the Personal Image Classification from App Inventor[19], where users could create, train, and test their image classifier and use it to create a game. Most tools (75%) had an explanation of the big idea; for example, Wekinator[20] offered detailed descriptions of algorithms used to train models.

Curricula such as AI4ALL and CuriosityMachine (Figure 1) were the most comprehension, with 63% covered all the "big ideas". Some curricula covered the ideas in narrow ways, focusing on a particular technology. For example, Embeducation[21] focused specifically on word embeddings. Nearly all (90%) curriculum had explanations of at least one of the five big ideas.

[16]https://www.tensorflow.org/tutorials/generative/pix2pix
[17]https://experiments.withgoogle.com/cold-flux
[18]https://supervised-polygons.github.io
[19]https://appinventor.mit.edu/explore/resources/ai/personal-image-classifier
[20]http://www.wekinator.org/
[21]https://embeducation.github.io

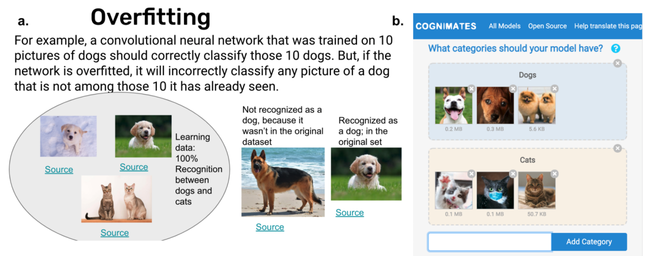

Figure 3: Examples of pedagogy integration from AI4All providing both direct instruction a) and active learning using Cognimates.

3.3 Pedagogical Strategies

The third facet of TPACK is how teacher knowledge of technologies is tied to particular pedagogical strategies. To examine these resources from this perspective, we analyzed the types of teaching methods resources engaged (active learning, direct instruction, or both) and the extent to which a resource accounted for learner prior knowledge.

Overall, we found that all the resources use either exclusively active learning or integrate active learning and direct instruction. Every resource had some interactive component, whether support for creating projects, training a model, or changing the model's parameters and seeing the outcome. We did not find any resources that were designed for purely direct instruction with no opportunity for practice or tinkering.

Despite this consistency in pedagogy, resource genres varied in their implementation. **Demos**, for example, primarily focused on self-contained interactive activities with limited opportunities for tinkering. Moreover, none offered any direct instruction, so it would be up to teachers to integrate them into a broader pedagogical strategy. InferKit[22], for example, was a demo that uses a neural network to generate text; it could support a range of pedagogical strategies involving active learning but offered no detailed guidance on how to do so.

Whereas demos offered unrestrained opportunities for tinkering, **activities** offered more structured active learning experiences with lightweight guidance. For example, Doodle Bot enabled participants to create a robot trained to draw based on speech commands, offering direct step-by-step instruction in tutorial form. About half of these resources offered multiple activities, with 27% giving learners the option to choose their activity and 28% offering fixed sequences of activities. For example, Code.org's AI for Oceans structured multiple activities around training a model to identify fish from garbage, unlocking activities as a learner makes progress.

Tools offered the most learner agency but also offered little scaffolding. Most (62%) gave learners the choice of what activity to do next. An example is RunwayML[23], a tool for creating a video with AI. Its environment offered several opportunities to build knowledge in arbitrary sequences of tutorials.

Whereas all of the other genres generally offered relatively little scaffolding, **curricula** offered the most structure and pedagogical

[22]https://app.inferkit.com/demo
[23]https://app.runwayml.com/

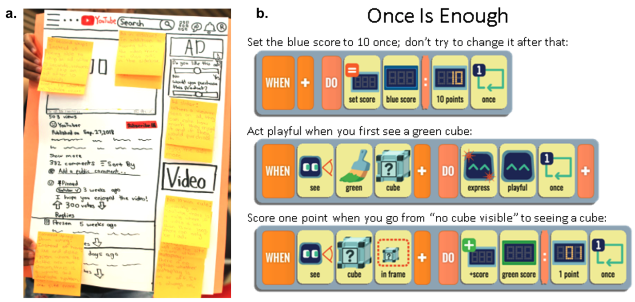

Figure 4: Some resources offered unplugged activities requiring no device, including AI Ethics and Calypso's activity sheet.

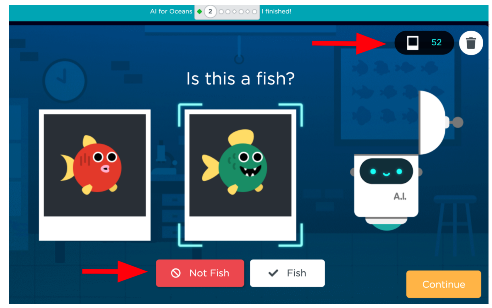

Figure 5: The AI Ocean Activity failed to provide any feedback, even when the learner mislabeled fish images.

support. The majority (63%) had a fixed sequence of activities. For example, STEM UK[24] was a curriculum with four sequential challenges, starting with an introduction of AI to later centering on the role of AI in making transportation safer, cleaner, and better connected. However, 26% allowed learners to make some choices in their progression. For example, Machine Learning for Kids[25] lets learners select activities based on project types, difficulty, and program environment. Most curricula (84%) used both direct instruction and active learning methods. For example, Figure 3 shows how AI4All combined direct instruction about overfitting with opportunities to tinker with overfitting a model by training a dog classifier.

3.4 Educational context

The fourth facet of TPACK is the particular educational context in which teacher knowledge is bound. To address this in our analysis, we considered the kinds of educational contexts the AI resources could support, asking: 1) what equipment they required, 2) if teachers might need to prepare a particular technical setup to use the resource, 3) if the resources were designed for a particular level, age, or grade, and 4) if the resources were accessible on a tangible or digital medium.

Overall, we found that 36% of the resources required some form of setup either because of their use of hardware, specific technical requirements such as libraries, or the creation of accounts. Most of the resources (62%) were digital-only, but 30% required a physical component, such as an unplugged learning activity or hardware integration. Only 8% of the resources were exclusively non-digital.

Only 59% of resources explicitly noted age or grade level. Of those 59%, most did not have implicit assumptions about either educator or students' prior AI and technical knowledge. For example, Scroobly[26], ModelZoo[27], and Ml5 Tool[28] required prior knowledge of both CS and AI, despite being framed as learning resources.

Each genre had distinct context assumptions. **Demos**, for example, all required computers with sufficient memory and compute power as some of the AI models they used were RAM intensive, but none required a technical setup beyond a web browser. Of the **activities**, 58% required some additional technical setup, and 25% had instructions for age and grade levels. Those that involved hardware, such as AIY kits for vision and sound [29], required significant familiarity with hardware components and technical setup. More than half (57%) of **tools** required a technical setup; all required computers or mobile apps. Fewer than half (43%) offered specific instructions regarding the age and grade levels of users. **Curricula** had the fewest technical requirements, with only 33% requiring configuration. However, all but two curriculum resources required the use of computers; the exceptions, shown in Figure 4, included AI Ethics[30] and Calypso[31], both of which involved activities that used paper and writing utensils instead of computers. Most curricula (77%) had age- and grade-based guidance, though several left the intended audience unstated.

3.5 Support for practice and assessment

The fifth and last facet of our TPACK analysis concerns how knowledge is deployed to develop students' knowledge. We, therefore, focused our analysis on how resources could support teachers in facilitating practice and assessment, analyzing if each resource: 1) provided support for practice and assessment, 2) provided opportunities for personalizing the learning experience, and 3) supported collaborative learning.

Overall, we found that 68% of the resources supported practice and assessment, 64% provided opportunities for customizing the learning experience by allowing teachers to either change the parameters of the resources or change the training data. In total, 40% of the resources supported collaborative learning.

Demos offered the fewest support for practice and assessment: only 33% supported repeated practice, only 22% allowed teachers to customize the configuration for learning, and only 11% allowed collaborative learning. None offered explicit support for assessment.

Activities tended to support practice (58%), often by allowing users to engage more in customizing either the input for the AI demo (i.e., record specific gestures like in the case of Plushpal[32]) or by customizing the output of the demo by changing how the demo output is displayed (i.e., Teachable Machine, allowing users

[24] https://www.stem.org.uk/resources/collection/447030/grand-challenges
[25] https://machinelearningforkids.co.uk
[26] https://www.scroobly.com/
[27] https://modelzoo.co/
[28] https://ml5.js
[29] https://aiyprojects.withgoogle.com/vision
[30] https://www.media.mit.edu/projects/ai-ethics-for-middle-school/
[31] https://calypso.software/
[32] https://www.plushpal.app/

to choose animations, text or sound). In total, 66% of activities supported the customization of the AI experience parameters, 58% had support for the practice, and 33% supported collaborative learning. Most activities offered no form of feedback on learners' actions; for example, the AI Oceans activity shown in Figure 5 allowed learners to label fish however they wanted and offered no explanation of how that might affect training.

Most **tools** (87%) offered substantial opportunities for practice. For example, iNaturalist[33] was a tool that used AI to support citizen scientists in classifying organisms. It had a path to practice adding IDs of an organism, comments, and observations before creating a project. On this platform, participants could post as many projects as they want. Most of the tools (75%) also allowed participants to personalize and customize their creations. One of the tools that do not allow it was Jukebox[34], a neural net that generated music. Jukebox let learners play with the creations of the model but unless participants could run the model on their computer they could not create their music. Another tool, AI Playground[35], allowed users to go more in-depth in modifying the AI parameters by controlling the number of training cycles (epochs). In some cases, tools tried to scaffold practice with activity sheets. Many sheets might be confusing because they introduced many new terms and references. For example, the activity sheet from Calypso (shown in Figure 4b) was meant to support users to learn how to program a robot but it could be difficult to grasp because it introduces a new programming language together with a series of new icons and terms.

All of the **curricula** we could access had activities for participants to practice AI concepts. For example, the AI and Machine Learning Module at Code.org[36] taught AI concepts at several different levels. Most curricula (89%) had the option to input customized data and personalized the outcome of the activities. Another example in this group is AI Ethics. The last module in this curricula is about YouTube re-design. Participants in this activity learn how YouTube uses AI, select what features they want to re-design, and have the option to present their mock-ups.

4 DISCUSSION

Overall, our analysis found the following:

- *Intended use.* Most resources, even those not designed for teacher use, had guidance that conveyed intended use. But the direction was often hard to find or required obscure technical knowledge to find and comprehend.
- *Content.* While most of the resources covered many of the AI4K12 big ideas, most did not cover all five, in most cases overlooking *Social Impact.* Curricula were the most likely to cover all five.
- *Pedagogy.* Most resources supported direct instruction and active learning combinations, though few were responsive to learners' prior knowledge.
- *Educational Context.* Most resources had some form of device dependency, constraining the learning and IT contexts in which they were compatible. Demo hardware requirements could be quite prohibitive for schools that do not have access to updated computers [26].
- *Student Learning.* While most resources offered substantial opportunities for individual and collaborative practice with AI concepts and skills, few offered assessment support or learner feedback.

In some ways, these findings reflect prior work on other classes of CS educational technologies. For example, Kim and Ko's evaluation of coding tutorials found a similar focus on active learning, a similar lack of communication about intended audience and context use, a lack of responsiveness to the student prior knowledge, and a disregard for formative and summative assessment [11]. Our results also mirror Kelleher's review of novice programming environments, showing a bias toward tinkering over direct instruction [10]. Our results also mirror the experience of educators who are currently designing their AI curriculum and directly expressed the need for support to combine the various AI resources and create a friendly learners interface [20, 26].

Our evaluation adds to these prior works in two ways. First, our results suggest that AI learning resources repeat some of the same mistakes of non-AI CS educational resources. Second, our results expand upon this, showing that many of the needs educators might have in developing TPACK to use AI resources aren't yet supported. Most resources do not clarify their assumptions about learner prior knowledge, required classroom resources context, alignment with pedagogical strategies, or even intended use. Even many of the curricula we analyzed were vague on these points. Some of the resources were consistent with implications from recent studies (e.g., leveraging personal data to an extent [16], embracing emerging student skepticism about AI [6], and leveraging embodiment [25]). But most resources did not meet basic pedagogical design principles, let alone offer the information teachers need to develop TPACK appropriate for successfully using the resources.

These findings have several implications for research. Future work might explore creating design principles for CS educational technology designers and understanding the barriers designers face in meeting those principles. In some cases, research is needed to achieve these principles. We see an opportunity for educators and designers to develop a common language based on a common set of guidelines, similar to the five big ideas [24]. For example, "features" could be described as "observable detail of object", "training" as "machines learning from data", and "model" as "application of what the machine has learned".

In terms of practice, our results suggest that until resource designers are more explicit about the various dimensions of TPACK in resource content, metadata, and design, teachers will have to make complex judgment regarding what resources might be appropriate for their students' learning. The curricula in our corpus generally fared best from a TPACK perspective (though not all were equal), with only two at the time of this writing—Curiosity Machine and AI4ALL—offering a clear path to adoption for teachers. Perhaps with time, resource designers and educators will find better ways of partnering, ensuring that all AI education resources can empower teachers to better facilitate AI education for all.

[33]https://www.inaturalist.org/
[34]https://openai.com/blog/jukebox/
[35]https://theaiplayground.com/
[36]https://studio.code.org/s/aiml-2021

REFERENCES

[1] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. *ProPublica, May* 23 (2016), 2016.

[2] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.

[3] Suzy Cox. 2008. *A conceptual analysis of technological pedagogical content knowledge*. Brigham Young University.

[4] Peter J Denning and Matti Tedre. 2019. *Computational thinking*. MIT Press.

[5] Stefania Druga and Amy J Ko. 2021. How do children's perceptions of machine intelligence change when training and coding smart programs?. In *Interaction Design and Children*. 49–61.

[6] Stefania Druga, Sarah T Vu, Eesh Likhith, and Tammy Qiu. 2019. Inclusive AI literacy for kids around the world. In *Proceedings of FabLearn 2019*. ACM, 104–111.

[7] Paolo Freire. 1996. Pedagogy of the oppressed (revised). *New York: Continuum* (1996).

[8] Joshua Grossman, Zhiyuan Lin, Hao Sheng, Johnny T-Z Wei, Joseph J Williams, and Sharad Goel. 2019. MathBot: Transforming Online Resources for Learning Math into Conversational Interactions. (2019).

[9] Tom Hitron, Yoav Orlev, Iddo Wald, Ariel Shamir, Hadas Erel, and Oren Zuckerman. 2019. 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*. 1–11.

[10] Caitlin Kelleher and Randy Pausch. 2005. Lowering the barriers to programming: A taxonomy of programming environments and languages for novice programmers. *ACM Computing Surveys (CSUR)* 37, 2 (2005), 83–137.

[11] Ada S Kim and Amy J Ko. 2017. A pedagogical analysis of online coding tutorials. In *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education*. 321–326.

[12] Matthew Koehler and Punya Mishra. 2009. What is technological pedagogical content knowledge (TPACK)? *Contemporary issues in technology and teacher education* 9, 1 (2009), 60–70.

[13] Duri Long and Brian Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–16.

[14] Radu Mariescu-Istodor and Ilkka Jormanainen. 2019. Machine learning for high school students. In *Proceedings of the 19th Koli Calling International Conference on Computing Education Research*. 1–9.

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

[16] Yim Register and Amy J. Ko. 2020. Learning Machine Learning with Personal Data Helps Stakeholders Ground Advocacy Arguments in Model Mechanics. In *Proceedings of the 2020 ACM Conference on International Computing Education Research* (Virtual Event, New Zealand) *(ICER '20)*. Association for Computing Machinery, New York, NY, USA, 67–78. https://doi.org/10.1145/3372782.3406252

[17] Sherry Ruan, Jiayu He, Rui Ying, Jonathan Burkle, Dunia Hakim, Anna Wang, Yufeng Yin, Lily Zhou, Qianyao Xu, Abdallah AbuHashem, et al. 2020. Supporting children's math learning with feedback-augmented narrative technology. In *Proceedings of the Interaction Design and Children Conference*. 567–580.

[18] Sherry Ruan, Liwei Jiang, Justin Xu, Bryce Joe-Kun Tham, Zhengneng Qiu, Yeshuang Zhu, Elizabeth L Murnane, Emma Brunskill, and James A Landay. 2019. Quizbot: A dialogue-based adaptive learning system for factual knowledge. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.

[19] Rosemary S Russ, Bruce L Sherin, and Miriam Gamoran Sherin. 2016. What constitutes teacher learning. *Handbook of research on teaching* (2016), 391–438.

[20] Alpay Sabuncuoglu. 2020. Designing one year curriculum to teach artificial intelligence for middle school. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education*. 96–102.

[21] Donald P Sanders and Gail McCutcheon. 1986. The development of practical theories of teaching. *Journal of Curriculum and Supervision* 2, 1 (1986), 50–67.

[22] Lee S Shulman. 2015. PCK: Its genesis and exodus. In *Re-examining pedagogical content knowledge in science education*. Routledge, 13–23.

[23] Matti Tedre, Tapani Toivonen, Juho Kaihila, Henriikka Vartiainen, Teemu Valtonen, Ilkka Jormanainen, and Arnold Pears. 2021. Teaching Machine Learning in K-12 Computing Education: Potential and Pitfalls. *arXiv preprint arXiv:2106.11034* (2021).

[24] David Touretzky, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn. 2019. Envisioning AI for K-12: What Should Every Child Know about AI?. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 9795–9799.

[25] Henriikka Vartiainen, Matti Tedre, and Teemu Valtonen. 2020. Learning machine learning with very young children: Who is teaching whom? *International journal of child-computer interaction* 25 (2020), 100182.

[26] Anu Vazhayil, Radhika Shetty, Rao R Bhavani, and Nagarajan Akshay. 2019. Focusing on teacher education to introduce AI in schools: Perspectives and illustrative findings. In *2019 IEEE Tenth International Conference on Technology for Education (T4E)*. IEEE, 71–77.

[27] Christiane Gresse von Wangenheim, Jean CR Hauck, Fernando S Pacheco, and Matheus F Bertonceli Bueno. 2021. Visual tools for teaching machine learning in K-12: A ten-year systematic mapping. *Education and Information Technologies* (2021), 1–46.

[28] Xiaofei Zhou, Jessica Van Brummelen, and Phoebe Lin. 2020. Designing AI Learning Experiences for K-12: Emerging Works, Future Opportunities and a Design Framework. *arXiv preprint arXiv:2009.10228* (2020).