



# AI Learning Priorities for All K-12 Students





## Authors and Leadership

The project is primarily planned, facilitated, and coordinated by the Principal Investigators:



Bryan Twarek  
Jake Koressel



Dr. Christina Gardner-McCune  
Dr. David Touretzky



## Project Support

Additional team members played substantive roles in the success of the workshop and producing high-quality outputs.

### DATA SYNTHESIS AND REPORT WRITING:

Dr. Julie M. Smith, Senior Education Researcher, Institute for Advancing Computing Education

### EVENT PLANNING AND COORDINATION:

Oliver Moss, Administrative Assistant, Carnegie Mellon University

## Funding Support

This project is supported by the National Science Foundation (NSF) under Grant No. 2444214. 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 NSF.



## Suggested Citation

CSTA & AI4K12 (2025). AI Learning Priorities for All K-12 Students. New York, NY: Computer Science Teachers Association. Retrieved from <https://csteachers.org/ai-priorities>.

This work is licensed under [CC BY-NC-SA 4.0](#).

## Acknowledgements

Thank you to the Carnegie Mellon School of Computer Science for hosting the workshop.

We also thank the following reviewers and CSTA K-12 Standards Advisory Board members for providing invaluable feedback:

Jared Amalong, Katie Bosch-Wilson, Josh Caldwell, Dr. Alexis Cobo, Dr. Jamila Cocchiola, Michael Corcoran, Christy Crawford, Dr. Melissa DeLaurentis, Amanda Dykes, Dr. Amy Eguchi, Sara Frey, Dr. Michelle Friend, Katherine Goyette, Laura Gray, Dr. Yasemin Gulbahar, Eli Hamrick, Alecia Henderson, Elissa Hozore, Dr. Christine Liebe, Jeremy Maker, Bill Marsland, Dr. Jaci McCune, Dr. Kelly Mills, Dr. Cynthia Mitchell, Paula Moore, Adam Musto, Brandon Nicholson, Katie O'Shea, Jen Rosato, Dr. Carolyn Rose, Dr. Pati Ruiz, Dr. Laurie Salvail, Spruha Satavlekar, Sherri Smith, Brett Tanaka, Dr. Ryan Torbey, Dr. Jane Waite, Shana V. White, John Wiseman, Kristina Yamada, Shaun Young, and Wonjin Yu.



# Table of Contents

<b>Convening Participants</b>	<b>05</b>
<b>Executive Summary</b>	<b>06</b>
<b>01. Introduction</b>	<b>08</b>
<b>02. The Process</b>	<b>13</b>
<b>03. Foundational K-12 AI Learning Outcomes</b>	<b>18</b>
A. Humans and AI	20
B. Representation and Reasoning	21
C. Machine Learning	22
D. Ethical AI System Design and Programming	24
E. Societal Impacts of AI	25
<b>04. Promising Practices for Teaching K-12 AI</b>	<b>26</b>
<b>05. AI Curriculum Alignment</b>	<b>28</b>
<b>06. Recommendations for AI4K12 Guidelines</b>	<b>30</b>
<b>07. Priorities for Future K-12 AI Education Research</b>	<b>32</b>
<b>08. Tensions and Challenges</b>	<b>34</b>
<b>Conclusion</b>	<b>36</b>
<b>Appendices</b>	<b>37</b>
A. Related Resources	37
B. Promising Practices Presentations	38
C. Initial List of Priorities	40
D. Categories of AI Bias	42
E. Glossary of Key Terms	43
F. Prioritized Foundational K-12 AI Learning Outcomes	47





Photo: Convening participants.

# Convening Participants

**Leah Aiwohi**, Teacher, Kauai High School, Lihue, HI

**Quiana Bannerman**, Director, Maryland Center for Computing Education, Adelphi, MD

**Dr. Tiffany Barnes**, Distinguished Professor, North Carolina State University, Raleigh, NC

**Kris Beck**, Director of Computer Science, Chicago Public Schools, Chicago, IL

**Dr. Nancye Blair Black**, CEO, The Block Uncarved, Lakeland, FL

**Angela Chavez**, Teacher, Computer Science Virtual Academy, LAUSD, Los Angeles, CA

**Beverly Clarke MBE**, Education Consultant, United Kingdom

**Sofia De Jesús**, Associate Program Manager, CMU CS Academy, Pittsburgh, PA

**Dr. Leigh Ann DeLysier**, Director, SRI, New York, NY

**Dr. Kayla DesPortes**, Assistant Professor, NYU, New York, NY

**Charlotte Dungan**, Chief Learning Officer, Mark Cuban Foundation, Durham, NC

**Veronica Ellis**, Content Development and Research Manager, Code.org/TeachAI, Boston, MA

**Mark Emry**, Austin Program Manager, Institute for Computing in Research, Austin, TX

**Rudy Escobar**, STEM and Computer Science Coordinator, Stanislaus County Office of Education, Modesto, CA

**Charity Freeman**, CS Curriculum and PD Specialist, Illinois Math and Science Academy, Aurora, IL

**Laura Gray**, AI and CS Instructional Coach, Gwinnett County Public Schools, Athens, GA

**Dr. Shuchi Grover**, Director of AI and Education Research, Looking Glass Ventures, Austin, TX

**Mary Cate Gustafson-Quiett**, RAISE Lead for K-12 Curriculum Development, MIT, Cambridge, MA

**Mahmoud Harding**, Instructional Design Director, Data Science 4 Everyone, Cary, NC

**Katie Henry**, Head of North America, Micro:bit Educational Foundation, Pittsburgh, PA

**Sallie Holloway**, Director of AI and CS, Gwinnett County Public Schools, Suwanee, GA

**Dr. Maya Israel**, Associate Professor, University of Florida, Gainesville, FL

**Dr. Sean Jackson**, K-12 CS Lead, Kentucky Department of Education, Tollesboro, KY

**Amber Jones**, Computer Science Curriculum and PD Consultant, Amber Sparks Education, Lithonia, GA

**Sarah Eli Judd**, Design and Technology Instructor, Franklin School, Jersey City, NJ

**Dr. Sonia Koshy**, Chief Research Officer, Kapor Foundation, Oakland, CA

**Dr. Victor Lee**, Associate Professor, Stanford University, Stanford, CA

**David Lockett**, K-16 Outreach, Meharry School of Applied Computational Sciences, Nashville, TN

**Kate Lockwood**, Director of CS and Engineering, St. Paul Academy, St. Paul, MN

**Dr. Fred Martin**, Professor and Chair, University of Texas at San Antonio, San Antonio, TX

**Dr. Cynthia Mitchell**, Director, AP Computer Science Principles, College Board, Satellite Beach, FL

**Christian Pinedo**, Chief of Staff, aiEDU, San Francisco, CA

**Kelly Powers**, CS Teacher, Hackley School, Tarrytown, NY

**Dan Schneider**, Principal Content Developer, Code.org, Tucson, AZ

**Vicky Sedgwick**, President, CSTA Greater Los Angeles, Canoga Park, CA

**Keisha Tennessee**, Computer Science Coordinator, Virginia Department of Education, Richmond, VA

**Dr. Emily Thomforde**, Computer Sciences Department Chair, Reach University, Asheville, NC

**Matthew Winters**, AI Education Specialist, Utah State Board of Education, Salt Lake City, UT

**Dr. Helen Zhang**, Research Associate Professor, Boston College, Chestnut Hill, MA





## Executive Summary

Rapid advancements in artificial intelligence (AI) necessitate changes in what AI content is taught to K-12 students. These changes will ensure that students are prepared to be smart consumers and competent creators of AI, as well as informed citizens. To meet this need, CSTA, in partnership with AI4K12, spearheaded the *Identifying AI Priorities for All K-12 Students* project. The project gathered experts – including teachers, researchers, administrators, and curriculum developers – to articulate priorities for AI education. This report summarizes the result of that effort. The project had four goals:

### 1. Identify priorities for AI learning across each K-12 grade band.

As a result of a collaborative, iterative process, the project articulated five categories for AI learning:



**HUMANS AND AI:** Younger students will benefit from comparing and contrasting human intelligence and AI systems and discussing how to use AI appropriately. Older students will need to critically explore whether and how to use AI tools for various purposes.



**REPRESENTATION AND REASONING:** The priorities in this area reflect the need for students to understand that, in order to be useful to humans, AI must first represent data about the world, and then use these representations to reason.



**MACHINE LEARNING:** This category includes sensing, data, how computers learn, and building and using AI models. It forms the technical core of understanding AI.



**ETHICAL AI SYSTEM DESIGN AND PROGRAMMING:** The power of modern AI systems suggests their potential for profound impacts, both positive and negative. A key to preparing students to be thoughtful creators and users of AI is experience with ethical considerations regarding AI systems.



**SOCIETAL IMPACTS OF AI:** AI has impacted many facets of daily life. Exploration of the societal impacts of AI will enable students to be thoughtfully informed in a society where AI systems play an ever larger role.



## 2. Suggest updates to the AI4K12 Guidelines.

Advances in generative AI necessitate updates to the AI4K12 Guidelines. This is especially true for Big Idea #4: Natural Interaction, since generative AI represents a substantial advance in the ability of AI to interact with humans. Similarly, generative AI raises many ethical questions relevant to Big Idea #5: Societal Impacts.



## 3. Advance the research agenda for K-12 AI education.

Priorities for research in AI education include the importance of supporting teachers, promoting inclusive and student-centered pedagogies, developing appropriate tools, gaining a better understanding of AI's impact on learning, and ensuring equity in AI education.

## 4. Share promising practices across the AI and CS education communities.


Participants shared their work in AI education. While the practices described varied, there were some common themes. Concerns about ethics and responsible AI were foregrounded, and hands-on learning activities were featured prominently. Meeting the needs of all children was a key concern, with approaches and tools that are widely accessible as well as engaging for all students. As a result of these common themes, we offer related recommendations for AI curriculum and instruction.


The report also includes an exploration of the tensions and challenges that emerged from the project, such as the difficulty of categorizing, organizing, and prioritizing learning content.

Preparing students to succeed personally and professionally in a world powered by computing will require rigorous, high-quality, and equitable learning opportunities in AI education. This project sought to determine priorities for AI education for all students to learn as part of a robust foundation in computer science, as well as options for more comprehensive study of AI. Within and across these priorities, two themes stand out:

- All students need to explore the personal, societal, and environmental impacts – both positive and negative – of AI.
- Students need to develop a broad conceptual understanding of how AI works: a frequent refrain from the project's participants was that students need to understand that "AI isn't magic."


While implementing high quality AI education, at scale, for all students will be challenging, the work already undertaken by convening participants demonstrates that there are elements of a foundation in place, one that can be built upon to ensure that all students are prepared to flourish in a world powered by computing.





Over the last few years, we have witnessed an unprecedented expansion in artificial intelligence (AI) capabilities, with ChatGPT experiencing the most rapid adoption of any consumer app in history.<sup>1</sup> Generative AI tools (that is, AI capable of generating text, audio, images, and/or video) have significant economic, legal, societal, and national security implications. Increasing public awareness of these impacts – ranging from environmental consequences<sup>2</sup> to debates about intellectual property<sup>3</sup> to concerns over potential job losses due to automation<sup>4</sup> – have made AI a staple topic of the news cycle in the last few years.

The goal of the *Reimagining CS Pathways* project was to establish a community consensus on what foundational CS experiences are needed to prepare students for personal and professional success in a world increasingly powered by computing.



<sup>1</sup> Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base. Reuters.

<sup>2</sup> Berthelot, A., Caron, E., Jay, M., & Lefèvre, L. (2024). Estimating the environmental impact of Generative-AI services using an LCA-based methodology. *Procedia CIRP*, 122, 707–712.

<sup>3</sup> Lucchi, N. (2024). ChatGPT: A Case Study on Copyright Challenges for Generative Artificial Intelligence Systems. *European Journal of Risk Regulation*, 15(3), 602–624.


<sup>4</sup> Gmyrek, P., Berg, J., & Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality. International Labour Organization.



Education has also been affected, as students rapidly adopted generative AI tools and school systems scrambled to articulate policies governing their use. Similarly, AI significantly impacts computer science (CS) education, as AI is a specialized field within the broader discipline of CS. This is evident in *Reimagining CS Pathways*, where the CS education community identified AI as a top priority for foundational CS content that positions “every student to be prepared for a world powered by computing.”<sup>5</sup>

It is thus no surprise that eight out of ten CS teachers believe that learning about and using AI should be part of a foundational CS learning experience and that CS standards should include AI.<sup>6</sup> In fact, when asked to identify the topics they teach, over two-thirds of CS teachers stated they covered AI specifically, despite the lack of explicit definition in most CS standards.<sup>7</sup>

**Figure 1: CS teacher beliefs about AI**



Currently, CSTA is in the process of revising its [K-12 Standards](#), which serve as a source for most states’ CS standards. These standards delineate learning objectives for a complete CS curriculum at the PK–12 level and are intended to define foundational learning for all students. The current CSTA standards are nearly a decade old, and research on how best to teach CS has evolved since these standards were written. And so has the field of CS, especially regarding cybersecurity, quantum computing, data science and, of course, AI. Additionally, implementation and scaling of CS has expanded: when the current CSTA standards were written, only 7 states had their own CS standards – now, 43 states do,<sup>8</sup> and some states – such as [Colorado](#), [Florida](#), [Ohio](#), and [Virginia](#) – have recently adopted CS standards focused on AI.

The [AI4K12](#) initiative is focused on three goals: advancing national guidelines for K-12 AI education, disseminating resources to support AI K-12 education, and developing the community of practitioners, tool and resource developers, and researchers who are focused on K-12 AI education.


As part of its standards revision effort, CSTA, in partnership with AI4K12, spearheaded the *Identifying AI Priorities for All K-12 Students* project. The project gathered experts – including teachers, researchers, administrators, and curriculum developers – to articulate priorities for AI education. This project was designed to support the revision of CSTA standards so that the new standards will be positioned to incorporate appropriate AI learning outcomes.

<sup>5</sup> CSTA, IACE, ACM, Code.org, College Board, CSforALL, & ECEP Alliance. (2024). *Reimagining CS Pathways: Every student prepared for a world powered by computing*.

<sup>6</sup> TeachAI & CSTA. (2024). *Guidance on the Future of Computer Science Education in an Age of AI*.

<sup>7</sup> CSTA & Kapoor Foundation. (2025) [Forthcoming]. *Landscape Study of PreK-12 CS Teachers in the United States*.

<sup>8</sup> CSTA & IACE. (2024). *K-12 Computer Science Standards Comparison Report: Examining the Similarities and Differences in State-Adopted K-12 Computer Science Standards and the CSTA K-12 Standards, Revised 2017*. New York, NY: Association for Computing Machinery.



This project had four goals:

## 1. Identify priorities for AI learning across each K-12 grade band.


These priorities will guide the new CSTA standards in their coverage of AI-related topics. The new CSTA standards are intended to reflect what all students – not just those who have access to and choose to study CS or who plan on pursuing a career in CS – should know before they graduate from high school in order to be prepared to thrive in a world powered by computing. The focus of these standards is on what is feasible for every school to implement. Thus, these are priorities for AI learning for all K-12 students.

## 2. Suggest updates to the AI4K12 Guidelines.

In contrast to the CSTA standards, the [AI4K12 Guidelines](#) (see Figure 2) articulate a deeper, more comprehensive study of AI – one that not all students will necessarily experience. These guidelines shape instruction in AI electives and CTE pathways. They also provide guidance to curriculum developers, school districts, nonprofits, after school programs, and teachers.

The advances in AI described above have made revisions to these guidelines desirable. For example, large language models (LLMs, a category which includes ChatGPT) are capable of markedly improved interaction via natural language than was previously possible. As a result of these advances, the societal issues raised by AI have also expanded, and these issues include trustworthiness, the effects of biased/inaccurate training data, the use of intellectual property in training AI systems, future regulation, and existential risks posed by AI agents.

Figure 2: The AI4K12 Big Ideas Wheel





### 3. Advance the research agenda for K-12 AI education.

Rapid advances in AI – combined with the relative newness of K-12 AI education – mean that K-12 AI education is under-studied and little is known about how best to teach it. Findings from CS education research – such as best practices for teaching students how to code – may or may not need to be updated as students increasingly use AI tools to support these tasks. As AI education scales nationally, it is important to understand what questions are most in need of additional research in order to best serve students.


### 4. Share promising practices across the AI and CS education communities.

Most convening participants had extensive experience with resources and practices for teaching AI, and some of these were shared at the convening. These practices represented a broad range of contexts – from experiences with how professional learning for AI education is implemented to AI-focused summer camps for students – and they consistently emphasized providing high quality, engaging, and equitable learning opportunities.

*Photo: At the convening, a group of participants discuss which topics related to AI representation and reasoning should be prioritized in K-12 AI education.*



In pursuit of these goals, the project was guided by the following values, which are adapted from the *Reimagining CS Pathways: High School and Beyond*<sup>9</sup> project:



#### EQUITY-CENTERED

Promotes broad and equitable access, participation, and experiences in CS education among all students.



#### COMMUNITY-GENERATED

Meets the needs of the community, including K-12 educators, postsecondary institutions, students, parents, and industry.



#### FUTURE-ORIENTED

Anticipates future needs of current learners, and prepares them for a future that is increasingly reliant on computing.



#### GROUNDED IN RESEARCH

Reflects the evolving body of knowledge of how students learn CS.



#### FLEXIBLE IN IMPLEMENTATION

Considers multiple pathways for meeting individual needs of learners, including regional, cultural, ability, social, and economic factors.

*Photo: Convening participants.*







## The Process



The project team selected participants to be purposefully diverse across numerous dimensions, including experience (see Figure 3), primary professional role (see Figure 4), grade band expertise (see Figure 5), gender (72% women, 23% men, and 2% non-binary), race (16% Asian, 21% Black, 12% Latine, 60% White, and 7% Multiracial), disability (14% have a disability or chronic condition), and geography (with 18 US states and the UK represented).<sup>10</sup>

Figure 3: Participants' experience



<sup>10</sup> These percentages do not total 100 since participants may be included in more than one category.



**Figure 4: Participants' primary professional role****Figure 5: Participants' grade band expertise**

Before the convening, participants were asked to identify:


- a list of what AI-related knowledge and skills should be prioritized for all K-12 students (see Appendix C and Figure 6)
  - how AI learning outcomes should progress across grade bands
  - what AI content advanced high school students should learn if they are interested in continuing their study of AI beyond the foundation articulated for all students
  - a short description of any K-12 AI education-related projects that they are working on (e.g., curriculum development, teacher professional development), including example activities, instructional strategies, and how the work aligns with the AI4K12 Guidelines
  - what AI content they find most challenging to teach



**Figure 6: A word cloud summarizing participants' initial prioritization of AI content**



In preparation for the convening, the project team inductively coded participants' lists of prioritized AI knowledge and skills in order to organize these priorities into themes and topics. The themes that initially emerged were:



Each theme had several topics and subtopics, which were developed iteratively. For example, consumer AI was divided into awareness of AI applications and their uses, practical usage of AI tools, and ethical use of AI tools, among others. Within these topics, all items mentioned by participants were categorized. Then, this content was mapped by the project team to a set of extant categories (some of which were aligned to the AI4K12 big ideas, see [Appendix C](#)).

This iteration of the categorization served as a starting point for participants as they worked together at the convening to consider what AI knowledge and skills should be prioritized.

The project team and participants attended a two-day convening at Carnegie Mellon University in February 2025, where a variety of collaborative activities were used to refine the content and organization of AI learning priorities, to be reflected in the updated CSTA standards and/or in the updated AI4K12 Guidelines.

As a part of the work of the convening, the categories synthesized in the pre-work were refined. This refinement was necessary due to the limitations on instructional time at each grade band. Several sessions of the convening were devoted to work that involved discussing and then combining, consolidating, and/or eliminating items from the preliminary content in order to generate the prioritized learning outcomes. This work involved several rounds of elaboration, prioritization, and discussion, as well as digital and analog polling. As shown in Table 1, some changes were also made to the initial categories of content, including reducing the number of categories from twelve to five. These revised categories serve as the basis for organizing the AI learning outcomes for all K-12 students, regardless of their career plans.




**Table 1: Initial and final categories of prioritized learning outcomes**


<b>Initial Categories</b>	<b>Changes</b>	<b>Final Categories</b>
Humans and AI	<i>None</i>	Humans and AI
Representation and Reasoners	<i>Slight name change</i>	Representation and Reasoning
AI Data: Input Data and Training Data	Combined with Machine Learning Fundamentals	Machine Learning
Machine Learning Fundamentals	Combined with AI Data: Input Data and Training Data	
Ethical Design Considerations and Evaluation	Combined with AI System Design and Programming	Ethical AI System Design and Programming
AI System Design and Programming	Combined with Ethical Design Considerations and Evaluation	
Societal Impacts of AI	<i>None</i>	Societal Impacts of AI
Sensation and Perception	<i>These categories were not included in the list of final categories, but some content from these categories is included in other final categories.</i>	
Generative AI		
Foundations of Neural Networks		
Past, Present, and Future AI Innovations		
AI Career Exploration		

Once these categories were determined, participants iteratively developed subtopics and learning outcomes.

Interspersed throughout the convening were brief presentations from participants highlighting their experiences with the practice of K-12 AI education (see [Appendix B](#)). The convening also featured time dedicated to articulating priorities for future AI education research.

After the convening, the project team analyzed and synthesized the products of the convening – which included notes, polls, surveys, and other artifacts – in order to generate this report. Then, the draft report was shared with the project’s advisors, convening participants, and other expert reviewers for feedback. The feedback was used to improve the report, including identifying the most important content to highlight.





## Foundational K-12 AI Learning Outcomes

Foundational K-12 AI education will support students in developing positive identities as informed citizens, critical consumers, and responsible creators of AI. It is the content that is most important for *all* K-12 students to learn. As a result of learning it, students should be able to:

### 1

Understand how AI technologies work and recognize where AI might be used

### 2



Use and critically evaluate AI systems, including their societal impacts and ethical considerations

### 3

Create – and not just consume – AI technologies responsibly

### 4

Be innovative and persistent in solving problems with AI




*Photo: Convening participants.*



This section delineates the specific prioritized content to teach within K-12 computer science courses. See Figure 7 for an overview of the categories and subtopics containing the prioritized content.


**Figure 7: Categories and subtopics of prioritized content.**



The following subsections include charts with the prioritized AI learning outcomes:



## HUMANS AND AI



Computing systems that include AI are increasingly able to perform tasks that have traditionally been done by humans. These tasks include writing, creating images, finding patterns in data, and using novel approaches to solving problems. Since these systems are easily accessible to the general public, students – particularly the youngest students – will benefit from comparing and contrasting human intelligence and AI systems and discussing how to use AI appropriately. Older students will need to critically explore whether and how to use AI tools for various purposes.

**Table 2: Humans and AI prioritized learning outcomes**

Subtopic	Grades K-2	Grades 3-5	Grades 6-8	Grades 9-12
<b>The Nature of Humans and of AI</b>	Compare and contrast the nature of humans versus the nature of AI (e.g., living versus nonliving).	Compare and contrast the ability of humans and of AI to perform various tasks and serve in various roles (e.g., create art, recognize emotions, be a friend, serve as a tutor).	Identify the assumptions inherent in the operation and output of an AI model and how these assumptions might have different implications for different people.	Debate what differences do or should exist between human and artificial intelligence, sentience, consciousness, rights, and responsibilities.
<b>The Human Role in Creating AI</b>	Understand that AI is a tool created by humans to make decisions or to generate something (e.g., an image).	Describe the role of humans in the creation of AI.	Describe the roles that humans play (including in data curation and labeling) in creating and refining AI models.	Evaluate and analyze the roles of humans and human decision-making in the creation of AI.
<b>The Choice to Use AI</b>	N/A	Evaluate when AI is or is not a helpful resource to carry out a task.	Debate when humans should or should not use AI to perform a specific task.	Analyze the risks, benefits, and effectiveness of using AI for specific tasks (e.g., coding, brainstorming), including when AI is used to fully automate a process or is used with a human-in-the-loop approach.

Given the wide scope of computer science and already full curriculum, the project team emphasized the need to define levels of priority. As a result, **rows within each table are marked with a star icon and highlighted in light green to indicate the most important of the foundational AI content**. This determination was made based on artifacts of the convening and detailed feedback during the review process for this report. Prioritized content (i.e., highlighted rows) across all five categories is compiled into a singular table in Appendix F.






## REPRESENTATION AND REASONING

Regardless of the specific type of AI system, two core concepts are involved in its operation: representation and reasoning. For AI to approach any problem, it requires data about the world to be represented in some way. For example, processing an image requires the image to first be represented digitally (e.g., as a series of pixels where each color is represented by a number). But representation isn't just about the input data (e.g., pixels in an image); it also encompasses knowledge (e.g., road maps, animal taxonomies) and data structures that support and are created in the reasoning process (e.g., search trees constructed by the reasoning algorithm, such as the example for a game of tic tac toe shown in Figure 8).

Representations are not just the input to the reasoner; they can be constructed as part of the reasoning process. Reasoning is accomplished via an algorithm that uses the data to reach a result in some systematic manner; that result might be a classification, prediction, inference, or the generation of media (e.g., an image). The AI priorities in this area reflect the need for students to understand that AI must sense, represent, and then reason in order to be useful to humans.

**Figure 8: Part of a search tree depicting a game of tic tac toe.**

Image source: Traced by User:Stannered, original by en:User:Gdr, CC BY-SA 3.0, via Wikimedia Commons



**Table 3: Representation and Reasoning prioritized learning outcomes**

Subtopic	Grades K-2	Grades 3-5	Grades 6-8	Grades 9-12
<b>Understanding Representation</b>	N/A	Understand how a representation is an abstraction that focuses on some features and leaves others out.	Understand that representation includes modalities (text, speech, audio, image, video) and symbolic mappings (text, graphs).	Describe how current AI models (e.g., LLMs) use data representation.
<b>Creating a Representation</b>	Create a representation of a physical object (e.g., line art drawing).	Create an abstract representation of a physical system that can be used to solve a problem (e.g., a map).	Create and evaluate different abstract representations (e.g., subway map).	Choose and use an appropriate representation of complex data for processing by an AI algorithm.
<b>Reasoning</b>	Explain how binary choices (e.g., up/down, on/off, under/over) can be used to make decisions that lead to a specific goal by either a human or a machine.	Train a model that can make decisions based on defined criteria (e.g., a dichotomous key to determine which movie to see).	Identify the kinds of AI models (e.g., classifier, predictor, recommender) people interact with in their daily lives.	Describe different types of AI algorithms and models, and compare and contrast the strengths and limitations of their reasoning.






## MACHINE LEARNING

While the study of data is likely to occur across most subjects, data has a special role in AI education given the need for AI systems to use data to reason. No AI system is perfect, which means that students will need to understand how various types of bias are manifest in data sets. Under this topic, students explore the roles of sensors, data, and how computers learn, as they work toward building their own AI models. Note that, in AI, the term 'bias' is defined differently than it is in common use; see [Appendix D](#).

*Whether to say that computers "learn" (or to use other terminology) is a difficult issue. On the one hand, "learn" may be confusing – especially to younger students – who may not grasp the differences between human and machine learning. On the other hand, the word is commonly used in the field of AI. Convening participants had differing opinions on its appropriateness. We have used it in this report, but we encourage care to ensure that AI systems are not anthropomorphized.*

- 
- 
- 
- 
- 



*Photo: Convening participants use 'dot voting' to indicate which topics they prioritize in the category of AI Data (which was later combined with the Machine Learning category).*



**Table 4: Machine Learning prioritized learning outcomes**

<b>Subtopic</b>	<b>Grades K-2</b>	<b>Grades 3-5</b>	<b>Grades 6-8</b>	<b>Grades 9-12</b>
<b>Sensing</b>	Compare and contrast human sensing with computer sensors.	Describe various ways that a human might interact with an AI system (e.g., through voice, text, or gestures).	Use sensors to collect data, and then train an AI model using the sensor data.	Using sensor data (e.g., from autonomous vehicles), train an AI model.
<b>Data</b>	Explore how AI models learn from data.	Explore the relationship between the properties of training data (e.g., size, features, biases) and an AI model's output.	Describe the ways that bias can be introduced and mitigated in an AI model.	Evaluate the data used to solve a problem, including its source(s) and whether privacy is protected, if/how the data has been processed, data quality (e.g., accuracy, reliability, validity), what the data represents, and biases.
<b>How Computers Learn</b>	Understand how computers learn from data and patterns.	Investigate how AI models learn by using data (including why examples and non-examples are required in training sets) and algorithms to find patterns and generate output.	Create and evaluate an appropriate AI algorithm (e.g., a decision tree classifier) to accomplish a task.	Select and use an appropriate AI algorithm for a classification task (e.g., KNN, decision tree).
<b>Building and Using AI Models</b>	Use data to construct a model for making decisions (e.g., a decision tree to determine what to wear based on the weather).	Using a dataset, develop an AI model to classify inputs.	Using a dataset and a machine learning pipeline, develop an AI model, and consider the impact of the model on various users.	Using a dataset and a systematic process, develop an AI model to generate for classification or prediction, and articulate the assumptions made at each of these steps: (1) develop a question solvable with AI, (2) collect or curate data, (3) evaluate the data, (4) train an AI model on the data, (5), evaluate the model, and (6) iteratively improve the model.






## ETHICAL AI SYSTEM DESIGN AND PROGRAMMING

The power of modern AI systems suggests their potential for profound impacts, both positive and negative. A key to preparing students to be thoughtful creators and users of AI is experience with ethical considerations regarding AI systems. Integrating ethics and design within one category implies that ethics is not an add-on or a separate topic but rather that ethical considerations should be integrated into every step of the design process. Key to ethical concerns is the recognition that AI systems will often have different impacts on different people and communities.

**Table 5: Ethical AI System Design and Programming prioritized learning outcomes**

<b>Subtopic</b>	<b>Grades K-2</b>	<b>Grades 3-5</b>	<b>Grades 6-8</b>	<b>Grades 9-12</b>
<b>Ethical Design Criteria</b>	N/A	Investigate an example of AI decision making, considering if it is fair – as well as what it means to be fair.	Explore strategies to turn ethical considerations into actions, such as mitigating bias in datasets.	Evaluate an AI model (e.g., using a model card) to determine the model's features as well as its biases, explainability, fairness, privacy, accuracy, and transparency.
<b>Ethical Evaluation of AI Systems</b>	Explore how an AI system can help and harm different groups at the same time.	Investigate examples of AI, considering differences in experience by different people in different contexts.	Describe the properties, biases, and assumptions of various kinds of AI models (e.g., classifier, predictor, recommender).	Evaluate the design, motivation, outcomes, and potential impacts of AI systems using ethical design criteria and/or ethical frameworks.
<b>Ethical Creation of AI Systems</b>	N/A	Describe an AI design process that considers the impact on end users and others who are impacted by the AI system.	Create a program using available AI tools, AI plugins, APIs, and/or AI models, with the following ethical considerations for the model's end users as well as others who are impacted by the model: fairness, bias, and accuracy, and then create a model card.	Train, iteratively improve, and then develop a model card for an AI model with the following ethical considerations for the model's end users as well as others who are impacted by the model: fairness, bias, safety, security, intellectual property, privacy, robustness, explainability, accuracy, transparency, and accountability.





## SOCIETAL IMPACTS OF AI

In recent years, AI has impacted many facets of daily life, moved economic markets, and presented new challenges and opportunities in fields ranging from environmental protection to medical discoveries. Exploration of the societal impacts of AI will enable students to be thoughtfully informed in a society where AI systems play an ever larger role.




Table 6: Societal Impacts of AI prioritized learning outcomes

Subtopic	Grades K-2	Grades 3-5	Grades 6-8	Grades 9-12
<b>Individual Impacts</b>	Identify where AI is being used in daily life.	Explore how one's actions may result in the collection of data.	Explore the tradeoffs related to human agency (including privacy, safety, creativity, autonomy, and intellectual property) when AI is used.	Evaluate how AI use impacts an individual's decision making and other behavior.
<b>Societal Impacts</b>	Explore how some people use AI in their jobs and in their communities.	Explore ways in which some jobs involve the creation and/or use of AI.	Identify the intended and unintended impacts of AI on society — including government, education, entertainment, culture, careers, and national security — while considering how these impacts may differ among diverse communities.	Evaluate the intended and unintended impacts of AI on society (e.g., deep fakes, job loss) — including government, education, entertainment, culture, careers, and national security — while considering how these impacts may differ among diverse communities.
<b>Environmental Impacts</b>	N/A	Explore the impact of AI on the environment.	Investigate the positive and negative environmental impacts of AI (e.g., minimizing deforestation via application of AI, energy use by AI).	Design ways to minimize negative environmental impacts of AI and communicate those ways to others.

### Implementing the Foundational Priorities



Note that the learning objectives in the charts above may be covered in many different courses (i.e., in an interdisciplinary way) and not necessarily in a CS course. Additionally, standards are normally written to be broad and general to permit implementation in a wide variety of contexts. This is especially important for AI-related learning standards, given the rapid technological changes that are anticipated over the next decade. (Plus, the approach of this project was to avoid learning standards tied to one particular tool, method, or approach.) However, broad standards can be difficult for teachers to implement – particularly when they are teaching a topic that is new to them, as is often the case for AI. Therefore, it is crucial to develop bridges between AI standards and their implementation that scaffolds the process for teachers. Examples of appropriate activities and tools will be helpful in this effort.





To meet the project's goal of sharing promising practices across the AI and CS education communities, the convening featured brief presentations from participants that highlighted their work in this field. See [Appendix B](#) for brief summaries of each presentation.

*Photo: Dr. Emily Thomforde presents a flash talk describing her promising practices.*



While the practices described covered a wide range of target audiences, topics, and approaches, there were some common themes across them as well. First, concerns about ethics and responsible AI were foregrounded as a key to AI education. Curricular materials related to AI ethics involved topics ranging from biases in datasets to the impact of generative AI on writers' careers.



For example, one presentation focused on engaging students' interest in sports by exploring the use of AI-generated game summaries for fantasy football leagues and then comparing and contrasting that practice with the use of AI-generated content in a major sports magazine, with an emphasis on the impacts on fantasy league participants, magazine readers, and professional writers.


Second, hands-on learning activities were featured prominently, emphasizing the need to engage students with experiences of creating and using AI, often via an unplugged approach for the youngest students. For example, the youngest students might develop a 'small language model' based on a popular children's book to understand conceptually how a large language model works – and why it sometimes delivers inaccurate output.

Finally, meeting the needs of all children was a key concern, with approaches and tools that are widely accessible as well as engaging for all students. One example of this was the application of a framework focused on ensuring that instruction is appropriate for all students; another example showcased how creating a programming environment that permits students to easily access online datasets and AI models can enable them to pursue personal interests.

Based on these presentations, we offer the following recommendations for AI curriculum:

### BEGIN IN EARLY ELEMENTARY SCHOOL

While studying some aspects of AI requires advanced mathematics, many aspects do not. Even the youngest elementary school students can learn – ideally from engaging, hands-on experiences – that, for example, a decision tree can be used to show the process by which decisions are made. This type of learning creates a foundation that permits older students to study more advanced topics.




### USE SUPPORTIVE TOOLS

We note that more tools need to be developed to support AI education. However, extant tools can help make complex topics accessible to students, and they can also provide interactive learning experiences that do not require programming to implement. For example, middle school students can train classifiers using Google's [Teachable Machine](#) and [MIT RAISE Playground](#), and high school students can experiment with neural networks using [Neuron Sandbox](#) or [TensorFlow Playground](#).

### PREPARE STUDENTS TO BE CRITICAL CONSUMERS, RESPONSIBLE CREATORS, AND INFORMED CITIZENS

It is important to avoid the tendency to think of AI education as primarily preparing those who will work as AI specialists. Only a vanishingly small portion of all students will pursue that career path. On the other hand, all students – regardless of career choice – will need to be critical consumers of AI as they determine whether, for example, they believe that the benefits of using an AI tool outweigh its environmental cost or whether the potential for historical biases in the tool's training data imply that the tool's output should not be trusted. Similarly, as citizens, they will benefit from developing well-informed and thoughtful positions about topics such as the advisability of regulating AI.





## AI Curriculum Alignment

Since the release of the AI4K12 Five Big Ideas in AI and the corresponding grade band progression charts, a number of national and international organizations have adopted them and used them to support their work. However, prior to the convening, we had little data about the extent to which curricula and standards are currently covering the guidelines.

Before the convening, participants were asked to share K-12 AI materials (e.g., lessons, curricula, professional development materials) developed by their organizations, as well as how that work aligns to the current AI4K12 Guidelines. Fifteen alignment documents were received. These alignments were combined and can be viewed as a series of heatmaps, a simplified version of which is found in Figure 9.

Overall we found that within our small sample, a large number of guidelines across the Five Big Ideas and grade bands were being covered. As the table shows, no subconcept is covered in more than six different curricula, and coverage in one, two, or three curricula is much more common. We observed large variability in coverage for individual subconcepts across grade bands. However, for vertical alignment within a grade band, we saw higher levels of coverage for the Big Ideas: Perception (as high as 78-88%), Representation & Reasoning (63-100%), Learning (67-91%), Natural Interaction (43-71%), and Societal Impact (60-90%). Of the five big ideas, Natural Interaction had the least coverage across grade levels. The guidelines within Natural Interaction include natural language, commonsense reasoning, and theory of mind, and these are currently the focus of revision given the advancement of generative AI, including LLMs.

This data suggests that even within a small sample, there are AI lessons, curricula, and standards that cover a large proportion of the AI4K12 Guidelines. The most covered subconcepts across multiple grade bands include sensors (computer sensors and sensing versus perception), AI in daily life, training AI models, bias in datasets, and using AI to solve societal problems. We also observed that many subconcepts have more coverage at the middle school level than at lower or higher grade bands, but this may be due at least in part to the fact that there are currently more curricula at that level and they were the most represented in our sample.



**Figure 9: Map of coverage of the AI4K12 Guidelines in AI curriculum. The numbers indicate how many curricula – of the fifteen that were mapped – cover each subconcept.**



Big Idea	Concept	Subconcept	K-2	3-5	6-8	9-12
<b>Perception</b>	Sensing	Living Things	1	0	0	0
		Computer Sensors	4	2	4	3
		Digital Encoding	0	2	1	2
	Processing	Sensing vs. Perception	3	3	5	4
		Feature Extraction	2	1	1	2
		Abstraction Pipeline: Language	0	1	1	1
		Abstraction Pipeline: Vision	1	0	0	1
	Domain Knowledge	Types of Domain Knowledge	1	1	2	0
		Inclusivity	1	1	3	2
<b>Representation and Reasoning</b>	Representation	Abstraction	1	2	1	0
		Symbolic Representations	3	1	1	0
		Data Structures	1	2	3	1
		Feature Vectors	2	2	3	1
		State Spaces and Operators	0	1	3	2
	Search	Combinatorial Search	0	0	2	0
		Types of Reasoning Problems	1	1	2	2
	Reasoning	Reasoning Algorithms	1	1	4	1
		Humans vs. Machines	0	0	5	1
		Finding Patterns in Data	3	4	1	2
<b>Learning</b>	Nature of Learning	Training a Model	2	5	5	1
		Constructing vs. Using a Reasoner	0	2	3	2
		Adjusting Internal Representations	0	1	3	2
		Learning from Experience	0	0	1	0
		Structure of a Neural Network	0	1	4	0
	Neural Networks	Weight Adjustment	0	1	2	0
		Feature Sets	3	2	5	1
		Large Datasets	0	2	0	0
	Datasets	Bias	1	0	6	1
		Structure of Language	1	0	0	3
<b>Natural Interaction</b>	Natural Language	Ambiguity of Language	1	0	0	2
		Reasoning about Text	0	1	1	1
		Applications	0	1	2	1
		Commonsense Reasoning	0	0	0	1
	Understanding Emotion	Understanding Emotion	0	2	1	0
		Philosophy of Mind	2	0	0	0
<b>Societal Impact</b>	Ethical AI	Diversity of Interests and Disparate Impacts	1	1	3	2
		Ethical Design Criteria	1	0	0	3
		Practicing Ethical Design	1	0	3	0
		Generative AI	0	1	4	1
	AI and Culture	AI in Daily Life	3	1	4	5
		Trust and Responsibility	0	0	3	1
	AI and the Economy	Impacts of AI on Sectors of Society	2	0	1	2
		Effects on Employment	0	1	2	2
	AI for Social Good	Democratization of AI Technology	0	1	2	2
		Using AI to Solve Societal Problems	1	2	2	2

## LIMITATIONS AND CAVEATS

There are several possible reasons why any given subconcept might not be covered in a particular curriculum, including that the curriculum was designed for a specific context (such as a topical summer camp) where the subconcept was not relevant. Or, a subconcept may be covered in a different subject, such as large datasets being embedded in a CS, math, or data science course. Further, coverage in a relatively high number of curricula does not necessarily indicate that a subconcept has been deemed especially important; rather, it may be that it is considered easy to teach and/or engaging for students. Topics may also vary in AI curricula that were not represented in this sample.

This mapping is the first attempt at aligning AI curricula to the AI4K12 Guidelines, and this effort will lay the groundwork for rigorous data collection in the future, which will enable better understanding of how the guidelines are covered.





## Recommendations for the AI4K12 Guidelines

The CSTA K-12 Standards and the AI4K12 Guidelines serve two different purposes. The CSTA K-12 Standards describe what all students – not just those who choose to study CS – should know. In contrast, the AI4K12 Guidelines articulate AI learning outcomes for each grade band when time permits for a more in-depth exploration of AI. In the elementary grades, the guidelines provide students with opportunities to explore foundational AI concepts. In middle and high school, the guidelines are designed to spark deeper interest in AI and the development of AI skills that can translate into a wide range of AI use from critical consumer, AI-enabled professional (e.g., doctor, business owner, teacher), to responsible and innovative AI creator (e.g., machine learning engineer, software engineer, AI researcher). Currently, the AI4K12 Guidelines are being used at the middle and high school level to inform the development of standalone AI electives, Career and Technical Education (CTE) pathways, and state standards that enable students to explore a more comprehensive study of AI.

Recommendations that emerged from participants' consideration of updates to the AI4K12 Guidelines are described below.

### SOME RECOMMENDATIONS SPANNED MULTIPLE BIG IDEAS:

The 9th-12th grade learning outcomes may be too advanced for high school students, and the 6th-8th grade learning outcomes may be a better fit for high school.

Teachers themselves who do not have any experience with AI may be at the level of the 6th - 8th grade learning outcomes, which would be appropriate starting points for professional development.

Learning outcomes could be framed more generally (e.g., "What are the different types of ways that AI makes decisions?") as opposed to more specific framing that may be too focused on the concerns of the present moment.

Now that AI systems are more widely available, the guidelines should include more content on whether and how students should use AI for a given task.



As described in the introduction, advances in generative AI – particularly LLMs – necessitate updates to the AI4K12 Guidelines (see Figure 2 for the current AI4K12 framework). This is especially true for Big Idea #4: Natural Interaction, since LLMs represent a substantial advance in the ability of AI to interact with humans using natural language; issues related to sentience and personhood (such as rights and responsibilities) may also be germane. Similarly, generative AI raises many ethical questions relevant to Big Idea #5: Societal Impact.

**BIG IDEA  
#2****REPRESENTATION AND REASONING**

Some parts of this idea, such as search, may be too specialized and/or too abstract for most students, especially at the younger grade levels. Also, learning objectives should explore what kind of reasoning LLMs use. Reasoning in LLMs is rapidly evolving, which adds to the challenge of teaching this topic.

**BIG IDEA  
#3****LEARNING**

Broaden the current focus on machine learning to include learning by AI systems, ensuring that all forms of AI are included. Provide a definition of AI systems that helps students understand that AI systems can combine several types of reasoners within architectures and algorithms.

**BIG IDEA  
#4****NATURAL INTERACTION**

This topic is the most in need of updating in the wake of recent advances in generative AI, with many proposed additions:

- ways to interact with AI, including principles of human-computer interaction and multimodality (i.e., audio) – not just natural language (i.e., text)
- physical computing, especially for younger students
- LLMs, including implications of their architecture for their output, n-grams, cultural context of natural interaction, and guardrails

**BIG IDEA  
#5****SOCIETAL IMPACT**

This topic should cover positive as well as negative impacts of AI on society, and it should include discussion of intellectual property, bias, regulations, and environmental concerns. The current learning objectives related to AI's impact on employment could be re-organized to be more developmentally appropriate and have more emphasis on developing the skills to succeed in a workplace that uses AI tools. This topic should also explore recognizing when a system is or is not using AI, as well as knowing that it may not always be possible to determine which is the case.


**Key recommendation:** Artificial intelligence has existed for decades, but it hasn't yet been a significant part of K-12 education, especially at scale. Scaling AI education will require that pre-service and existing teachers are prepared to teach prioritized AI content as part of foundational computer science.





## Priorities for Future K-12 AI Education Research

07



Participants identified priorities for future K-12 AI education research, which were then grouped into categories (see Figure 10). (Note that there is some overlap between the categories, and some priorities, such as inclusive pedagogy, might fit easily into more than one category.) The categories of research priorities are:

### SUPPORTING TEACHERS

By far, the highest priority for future K-12 AI education research is centered around the question of how best to support teachers, particularly by providing high-quality professional development at scale to prepare teachers to implement AI education. A recent survey found that 88% of CS teachers desire support for teaching about and with AI.<sup>11</sup> Similarly, less than half of middle and high school CS teachers (44% and 46%) feel equipped to teach AI, a percentage that is even lower for elementary school teachers (34%).<sup>12</sup> While offering high-quality professional development at scale is a challenge across all subjects and grade levels, it is particularly acute in AI education for two reasons. First, the rapid changes in AI mean that professional development will require an approach that can accommodate the evolution of the field. Secondly, most educators will have had little to no exposure to AI in their own education or in prior professional development. This lack of prior experience presents a challenge to be sure, but it also points to the importance of researching all aspects of AI education so that the field can scale according to best practices before less effective approaches become entrenched. While standards and guidelines themselves will not necessarily focus on teacher learning, that learning should nonetheless be foregrounded in designing and implementing these materials. Otherwise, they may be overwhelming in terms of breadth of content, vocabulary, and so forth for those without AI education experience. And, of course, very few at any role or level in the K-12 education system have substantial AI education experience. Thus, scaffolds – including glossaries, examples, and explanations – would likely aid effective implementation.

### PEDAGOGY

A vast array of issues related to pedagogy require more research, which is not surprising given the newness of AI education. Specific research questions related to pedagogy include how to ensure that pedagogies for teaching AI are inclusive, which tools work best for teaching about AI, how to teach about AI in an interdisciplinary manner, and what types of learning activities related to AI are most effective for students. We note that, while pedagogy was one of the top research priorities, curriculum per se was not.

<sup>11</sup> TeachAI & CSTA. (2024). *Guidance on the Future of Computer Science Education in an Age of AI*.

<sup>12</sup> CSTA & Kapor Foundation. (2025) [Forthcoming]. Landscape Study of PreK-12 CS Teachers in the United States.



## TOOL DEVELOPMENT

The importance of developing new, grade-appropriate tools emerged not only in discussions of research priorities but also as an observation about AI4K12's Big Idea #4: Natural Interaction; specifically, many of the subconcepts that are or could be in Big Idea #4 would be easier to teach if new tools were developed. However, the terms of use for LLM-based AI tools generally require users to be at least 13 years old, which presents obstacles to their use by elementary and middle school students. Part of the research agenda related to tool development will require work at the intersection of law, ethics, and tool development to ensure that younger students have appropriate access to AI tools. Plus, these tools need to be accessible to a broad range of learners and contexts.

## AI'S IMPACT ON LEARNING

There is very little research on how students' thinking and learning change when they use LLMs. Similarly, it is possible that optimally effective approaches to instruction will change in the wake of the integration of LLMs and other AI systems into the learning process.


## EQUITY ISSUES

Research focused on barriers to AI education for all students and how to implement high-quality AI education in underserved communities is needed. Further, all research questions should examine the impact of AI education on different groups of learners, including students with disabilities, students from different socioeconomic groups, and so forth. How best to mitigate disparities in access to AI systems – especially robotics – is also a concern. Given current political and cultural trends, scaling equitable AI education will be particularly challenging.

## OTHER PRIORITIES

Other research topics prioritized by participants – although to a lesser extent than those mentioned above – include the impacts of AI education on students, how to assess AI education, systemic analysis of AI education implementation, how to align AI education to the needs of industry, how to implement AI education in ethical ways (e.g., given the potential harms of AI systems to individuals, communities, and the environment), and the impact of AI on teaching practice.

Figure 10: Priorities for AI Education Research



08

# Tensions and Challenges



We note several tensions and challenges that surfaced as convening participants grappled with the question of what AI learning outcomes should be prioritized:

## TERMS AND THEIR DEFINITIONS

Unfortunately, words related to CS and AI are not used with perfect consistency or universal agreement about their meaning. (See the glossary [Appendix E](#).) Thus, deciding how to phrase priorities was sometimes challenging, especially with the goal of producing recommendations that are accessible to teachers, curriculum developers, and others who may not be experts in AI. For example, participants debated whether it was appropriate to say that AI ‘learns’ – a framing common in CS but one that may not be clear to younger students who may not appreciate the differences between how humans learn and how AI learns. Similarly, sometimes parents and administrators are not clear that AI is a subset of CS, or that AI education is not limited to the use of AI tools.

## CATEGORIES AND ORGANIZATION

Before and during the convening, the project team and participants had to organize AI learning priorities into categories in order to produce a coherent outcome. However, there is no perfect way to classify learning activities. For example, two categories of priorities – (1) Machine Learning and (2) Ethical AI System Design and Programming – will necessarily have some overlap, given that machine learning topics will generally be part of AI programming. However, combining these categories would have led to disadvantages as well, since there are distinct constructs within each.

## DISCIPLINARY BOUNDARIES

It is not always clear what topics should be considered part of AI learning, part of CS more generally, or even part of another school subject. For example, ethical issues related to the use of data (e.g., identifying limitations in a data set, securing permission to use data, maintaining a data source’s privacy) are crucial to learning about AI – but are also part of CS more generally and may also be covered in, for example, a social studies or a science class. Given the twin desires to neither have gaps in student learning nor to duplicate instruction, it is sometimes difficult to determine what topics should be considered a priority for AI learning. Additionally, there is a distinction between AI education as a subset of CS (i.e., creating AI) and AI literacy (i.e., using AI), and it is not clear what school subject(s) would be the most appropriate home for AI literacy.



## GRANULARITY

It can be difficult to determine the best granularity level for AI priorities. For example, in the preparation for the convening, some participants identified relatively narrow learning priorities (e.g., adjusting weights and parameters of an AI system), while other priorities were defined more broadly (e.g., developing an AI model). While a consistent level of granularity might be most useful, determining the appropriate level is difficult.

## BALANCING THE SOCIAL AND THE TECHNICAL

Some participants were concerned that the priorities for AI education included too much emphasis on technical aspects of AI – which are likely to change rapidly – at the expense of emphasizing the personal and societal impacts of AI. On the other hand, other participants expressed the opposite concern: too much emphasis on societal and ethical issues and not enough on technical aspects of AI, given the context of CS instruction (as opposed to, for example, the context of a social studies, ethics, or philosophy course).

## FUTURE-PROOFING

Obviously, AI is advancing at a rapid pace. With the goal of articulating AI priorities that are likely to be in use (e.g., as a basis for new state CS standards) for a decade or more, it is difficult to determine what topics and approaches will be relevant in the future. For example, will LLMs command the same focus in AI in a decade that they do today?

## TIME CONSTRAINTS

A guiding assumption of the CSTA standards revision project is that instructional time for CS (including but not limited to AI) is likely to be as follows: (1) 20 - 40 hours per year (or 30 - 60 minutes per week) throughout elementary school, (2) one year-long course (or equivalent) during middle school, and (3) one year-long course (or equivalent) during high school. Thus, with an already overfilled CS curriculum, instructional time for AI will be quite limited, and many desirable topics will simply not be able to be included. It is difficult to prioritize learning outcomes within these time constraints. Time constraints were a primary consideration in the decision to exclude the K-2 grade band in some of the progressions defined in Tables 2-6.

## CHALLENGING CONTENT

Participants were asked what AI content they found most challenging to teach. The most common responses were (1) ethics and societal impacts and (2) any AI topic that relied on a strong math background, which students often have not been given sufficient support to develop. There is also a lack of appropriate tools to teach various AI topics.

We note that many of these tensions and challenges might be at least partially mitigated by additional AI education research.



Photo: Convening participants.



# Conclusion



Preparing students to succeed personally and professionally in a world powered by computing will require rigorous, high-quality, and equitable learning opportunities in AI education across all grade bands, for all students. While some AI related topics (such as data preparation) may fit into other school subjects, AI will generally be part of CS due to the need for a technical understanding of how AI systems work. Deciding what foundational AI content for all students should be prioritized is a challenging question across several dimensions, given an already full curriculum, rapid technological change, and the need to support teachers who in most cases will be teaching content that is new to them.

This project sought to determine priorities for AI education for all students, as well as options for more comprehensive study of AI. Within and across these priorities, two themes stand out. First, all students need to explore the personal, societal, and environmental impacts – both positive and negative – of AI. While rapid technological advances are quite likely to change much about the technical nature of AI over the next decade, students will nonetheless need to appreciate AI's impacts in order to be informed and responsible consumers and creators of AI products.

Second, students need to develop a broad conceptual understanding of how AI works: a frequent refrain from the convening was that students need to understand that "AI isn't magic" and that demystification is important for students at all levels. The conceptual understanding that comes from demystification will vary depending on the AI system and future advances in AI, but that understanding will equip students to better understand the AI systems that they encounter. For example, understanding the architecture of an LLM prepares students to critically evaluate the model's output in light of its capabilities and limitations. In other words, students should be thoughtful evaluators – not passive consumers – of AI.

There was also an overarching concern for teachers. There is a clear desire to rapidly scale high-quality AI education, but with the concomitant realization that it requires comprehensive support of teachers who are likely to have had little to no academic study of AI themselves. High-quality professional development designed for teachers whose subject matter expertise is in areas other than AI will be key to supporting these teachers and their students.

While implementing high-quality AI education, at scale, for all students is obviously very challenging, the broad and diverse work already undertaken by convening participants demonstrates that there are elements of a foundation already in place, one that can be built upon to ensure that all students are prepared to flourish in a world powered by computing.



## APPENDIX A: Related Resources

- These items were shared with participants as preparatory for the convening and provide an overview of issues related to AI priorities:
  - Draft TeachAI Brief: [What AI Experiences are Foundational for Every CS Student?](#)
  - [Reimagining CS Pathways: Every Student Prepared for a World Powered by Computing](#) (particularly sections 2, 3.3, and 7)
  - [AI Design Brief for Standards Writers](#)
- TeachAI and CSTA conducted a [survey](#) of CS teacher views about AI in 2024
- A [glossary](#) of AI-related terms
- AI4K12 current [guidelines](#) for AI in K12, including grade band progression charts
- [Heatmaps](#) of alignment of selection K-12 AI resources to the current AI4K12 Guidelines
- Other resources shared by participants
  - [The Daily-AI workshop](#) from MIT RAISE
  - [AI Literacy](#) from Digital Promise
  - [AI for K-12](#) from the AI-CARING Institute
  - [Everyday ai](#) from EdAI
  - [AlforCA](#) from CSforCA
  - [Self-driving Car Starter](#) Scratch program from Emily Thomforde
  - [Privacy and Safety Checks](#) for AI tools from Sarah Wood
  - [Gradual Release of Responsibility](#), based on work by Doug Fisher and Nancy Frey
  - [Resources](#) from AI4All
  - [How AI Works, Generative AI for Humanities, Foresight in AI chart, Model Cards, and Classification Models](#) from Code.org
  - [CS Fusion across the Curriculum](#) from Angela Chavez
  - [Framework for Administrators and Decision-makers on AI Implementation in Schools](#) from Sofía De Jesús and Kip Glazer
  - [Technology Integration Matrix](#) from the Florida Center for Instructional Technology
  - [Picture This](#) materials from MIT RAICA
  - [AI Hub](#) from the Utah Education Network
  - [AI Summit Materials](#) from the Utah State Board of Education
  - [AI Introduction & Implementation](#) from the Utah STEM Action Center
  - [AI in the Classroom: Strategies and Activities to Enrich Student Learning](#) from Nancye Blair Black
  - [Applied Computational Thinking with Python](#) from Sofía de Jesús and Dayrene Martinez
  - [How does FacelID classify images?](#) from CRAFT
  - [Next Level AI Skills for Educators](#) from ISTE



## APPENDIX B: Promising Practices Presentations

This appendix presents brief summaries and links to resources shared by participants during the convening.

**Nancye Blair Black** described how learning about AI can adopt a discovery-based approach and, when integrated into other subjects, can ensure that AI education is available to all students and that it can fit into an already full curriculum.

- [Artificial Intelligence in Education](#)
- [AI is for Everyone, Everywhere](#)

**Charlotte Dungan** described the use of AI model cards (similar to a nutrition facts label; it encourages transparency and explainability by describing an AI model's creation and performance, including ethics and safety) and constructionist learning methods.

- [Give P's a Change: Projects, Peers, Passion, and Play](#)
- [Projects, Passion, Peers, and Play](#)
- [Cultivating Communities of Creativity and Caring](#)
- [All I Really Need to Know \(About Creative Thinking\) I Learned \(By Studying How Children Learn\) in Kindergarten](#)
- [Importing Models in App Lab and Model Cards in App Lab](#)
- [Google Model Cards](#)
- Model cards for [Gemma](#), [MediaPipe FaceMesh](#), [Object Detection](#), and a [notebook](#) for generating model cards

**Charity Freeman** shared an example of an AI module where students analyze social media posts, develop a classification framework, and then explore both supervised and unsupervised machine learning, including a discussion of biases.

- [Find the AI Approach That Fits the Problem You're Trying to Solve](#)

**Laura Gray** shared lessons learned from delivering professional development about AI to teachers, emphasizing the role of coaching and mentorship for teachers and 'byte size' professional learning opportunities.

**Shuchi Grover** explored the use of NetsBlox, an extension of Snap! (a block-based programming language) that permits accessing online data sets and LLMs and thus supports AI instruction.

- [CS Frontiers](#)
- [AI & Cybersecurity for Teens](#)
- [Block-based Abstractions and Expansive Services to Make Advanced Computing Concepts Accessible to Novices](#)
- [Design of Tools and Learning Environments for Equitable Computer Science + Data Science Education](#)
- [Cybersecurity Education in the Age of AI: Integrating AI Learning into Cybersecurity High School Curricula](#)

**Mary Cate Gustafson-Quiett** and **Kelly Powers** shared standards-aligned modules focused on responsible AI.

- [Responsible AI for Computational Action](#)

**Mahmoud Harding** shared an example of a lesson that could be used to engage students in ethical issues related to AI. It explored the benefits and drawbacks of AI-generated articles about sports for both fantasy leagues and for major sports magazines.

- [AI, Human, or HumAln lesson plan](#)

**Sallie Holloway** shared how AI learning opportunities have been implemented in Gwinnett County Public Schools, including through an AI pathway with three courses: Foundation of AI, AI Concepts, and AI Applications, all of which emphasize programming, data science, ethics, machine learning, and problem solving. Their overall AI learning model is taught through a framework that's embedded into all K-12 classes (not a separate class) as well as the pathway.

- [GCPS AI Shared Resources](#)



**Maya Israel** described the integration of the principles of Universal Design for Learning into AI education in order to ensure that instruction is appropriate for all students, and how this approach was used in an AI education summer camp.

- [Universal design for learning guidelines for computer science and computational thinking](#)
- [A framework for inclusive AI learning design for diverse learners](#)

**Amber Jones** described a 9-week middle school AI curriculum that helps students explore AI systems.

- [AI4GA](#)
- [AI Career Cards](#)

**Sonia Koshy** presented the Kapor Foundation's *Responsible AI and Tech Justice: A Guide for K-12 Education*, which uses the lenses of racial and social justice to prepare students to interrogate AI, responding to the harms of AI systems. The guide includes interrogation questions, such as "What are the potential positive and negative consequences of using AI technologies for personal convenience?"

- [Responsible AI and Tech Justice: A Guide for K-12 Education](#)

**David Lockett** described modules developed to teach K-12 educators about trustworthy AI and how these modules are necessary due to the way AI is reshaping many industries.

- [AI Institutes Virtual Organization](#)
- [Trustworthy AI teaching modules for K-12 Educators](#)

**Kate Lockwood** explained how math is often a gatekeeper preventing high school students from data science, CS, and AI courses – but how these subjects can instead be a gateway to math, recommending an approach of "curate, don't create" for instructional materials.

**Daniel Schneider** reported how Code.org promotes the LEADERS framework to support teachers in leading discussions about the societal effects of AI: Link to your subject area, Errors occur for many reasons, Assume someone in the room is your datapoint, Data sources can be biased, Empathy for multiple stakeholders,

Real-world cases are history, not opinion, and Show solutions. This resource, and several others, were adapted from the work of Dr. Yim Register.

- [Foundations of Generative AI course](#)
- [More Than "If Time Allows": The Role of Ethics in AI Education](#)
- [Teachers Guide: Ethical Opportunities in Exploring Generative AI](#)
- [The Future of AI Can Be Kind: Strategies for Embedded Ethics in AI Education](#)
- [Societal Impact of Generative AI Lesson](#)
- [LEADERS framework for AI impact and bias lessons](#)

**Vicky Sedgwick** shared examples of AI instruction appropriate for kindergarten through second grade students, focused on the idea that AI is not magic but rather makes predictions and can be biased.

- [Peeking Under The AI Hood: Unplugged Activities For Grades K-2](#)
- [Teaching AI Starting in \(Pre\)K-2nd Grade](#)

**Keisha Tennessee** described the *Virginia 2024 Computer Science Standards of Learning* review and revision process, which included a focus on CS concepts and skills that students need as consumers, creators, and citizens (the three "Cs") in order to understand CS and AI through these three roles.

- [Computer Science Standards of Learning for Virginia Public Schools](#)

**Emily Thomforde** described how Reach University has incorporated AI topics into the preparation of all preservice teachers. This approach focuses on what all teachers need to know about AI, which they organize into three categories: using AI, the impacts of AI, and building AI.

- [CS 300: Computer Science course portal, syllabus, and AI touchpoints](#)
- [CS 300b: Computer Science and Society syllabus](#)
- [CS 333: Pedagogy and Andragogy Fall 2024 course portal and Spring 2025 course portal](#)



## APPENDIX C: Initial List of Priorities

Before the convening, participants were asked to share their priorities for AI experiences and skills. The project team organized and synthesized those priorities into categories, as shown below.



### Sensation & Perception

*Perception is the extraction of meaning from sensory information (raw data) using knowledge.*

- Explain the various ways computers get data (sensors, datasets collected by people, social media/technologies, others)
- Use tools to explore and explain how computers represent and understand sound (audio abstraction pipeline)
- Use tools to explore and explain how computers see and understand the visual world (computer vision pipeline)
- Model how AI systems extract meaning from sensory information using knowledge and multiple levels of abstraction



### Representations & Reasoners

*"Knowing" something means the ability to both represent it and reason with it.*

- Identify the kinds of reasoners (classifier, predictors, recommenders, planning, scheduling) students interact with in their daily lives
- Identify algorithms for different AI reasoning types
- Demonstrate that representations are data structures (graphs, trees, vectors, embeddings) and reasoning methods are algorithms
  - Create and evaluate a decision tree classifier using a decision tree learning algorithm
  - Create and evaluate a search tree using a search algorithm (e.g., explore a game space)
  - Create a feature vector to encode key characteristics of an object so that a ML algorithm can learn patterns from the data.
  - Describe the relationship between word embeddings, ML, and LLMs



### AI Data: Input Data & Training Data

*Garbage in. Garbage out.  
Selection and quality of data are important for reliable reasoning.*

- Identify kinds of questions that can and cannot be answered with a dataset
- Find/collect/curate dataset, select features, find patterns in data, train a model
- Explain why data science practices (like curating, cleaning, investigating data) are key to successful AI systems
- Evaluate the quality of data and potential sources of bias (e.g., missing data, errors, size, representation, encoding, proxies)
- Evaluate how the choice of training data shapes the behavior of the AI system and how bias can be introduced in the training dataset
- Explain concerns and issues with training data (e.g., cost, ownership, privacy) and common solutions

### Machine Learning Fundamentals

*Machine learning allows a computer to acquire behaviors without people explicitly programming those behaviors*

- Describe how data representation and probabilities inform AI & ML (AI isn't magic)

- Describe the different types of machine learning and example learning algorithms
- Use the machine learning pipeline to build and evaluate a model for classification or prediction
- Use a model for inference (classification or prediction).
- Explain the difference between training and inference
- Describe the ways the bias can be introduced and mitigated into the ML process. (understand that bias is inevitable despite best efforts).

### Foundations of Neural Networks

*Modern AI is based on neural networks, statistical models.*

- Describe the structure of a neural network
- Identify common applications of neural networks to solve problems (e.g., CNN - Computer Vision, RNN - time series predictions (stock/housing market prices), GAN-Image generation, Transformers-LLMs)
- Experiment with a single neuron (linear threshold unit) to understand how neurons compute their values
- Create a 1-3 layer neural network (input, 1-3 hidden layers, and output)



## Initial List of Priorities (Continued)



### Generative AI

*Intelligent agents require many kinds of knowledge to interact naturally with humans.*

- Describe the kinds and amounts of data used to train Generative AI and Foundational Models
- Identify AI agents and explain their capabilities
- Use tools to explore and explain how computers understand the meaning of a prompt, generate text, and images
- Describe the limitations of generative AI and foundational models



### AI System Design & Programming

*Empowering and cultivating AI designers not just consumers*

- Create a program using available AI tools, AI plugins, APIs, models
- Create a program that incorporates a model created by the student
- Evaluate AI Systems and their outputs
- Design and evaluate AI systems using ethical design criteria
- Program and debug with AI Assistance or AI tools for programming
- Explain the relationship between coding and AI
- Responsible AI development practices e.g., user-centered development, model cards, and iterative testing

### AI Career Exploration

*Equipping the next generation of AI-enabled professionals*

- Investigate AI Careers
- Explain the positive and negative impacts of AI on other careers

### Ethical Design Considerations & Evaluation

*Equipping responsible and ethically AI designers and evaluators*

- Identify potential sources of bias and positive/negative impacts of AI
- Design and evaluate AI systems using ethical design criteria

- Explain the ethical and societal issues and concerns about generative AI and foundational models

### Humans & AI

*Exploring the role of humans in the creation and usage of AI*

- Describe of the roles of humans in the creation of AI (esp. labeling, data curation, design process, and system usage.)
- Creating & Innovating using AI
- Investigate how copyright, creativity, and data ownership inform how AI systems are built and their impacts
- Nature of Humans vs AI

- Evaluate AI systems that impact people and society for fairness, transparency, adherence to intellectual property, and privacy

### Societal Impacts of AI

*AI can impact society in positive and negative ways*

- Evaluate the AI impacts on systems and tools that students use every day
- Evaluate societal impacts of AI on culture, different groups of people, industries, and sectors of society
- Identify intended & unintended Impact on how AI affects users
- Describe the environmental impacts of AI
- Investigate how AI uses data and where/how that data is generated and discuss the potential implications: who/what is represented, copyright, environmental impact, etc
- Identify AI errors and misuses of AI and their impacts on people (trust, privacy, rights)
- Investigate current regulations for the design and usage of AI systems

### Past, Present, Future AI Innovations


*Understanding AI Innovation*

- History of AI
  - History of AI: From Turing to GPT
  - History/Evolution of AI
- Current AI Innovations & Uses
  - Applications and Uses of AI
  - Understand the factors contributing to the advancement of AI Innovation
  - Beauty in Discovery: How AI & ML are Changing Research and Science
- Future
  - Anticipate the future AI Innovations



## APPENDIX D: Categories of AI Bias

Image source: Schwartz, R., Vassilev, A., Greene, K., Perine, L., Burt, A., & Hall, P. (2022). Towards a standard for identifying and managing bias in artificial intelligence (Vol. 3, p. 00). US Department of Commerce, National Institute of Standards and Technology.



## APPENDIX E: Glossary of Key Terms

This glossary is based on the AI4K12 [Glossary](#).

<b>Abstraction</b>	Simplification by the elimination of unimportant details; one of the core components of computational thinking.
<b>Agent</b>	Any artificial intelligence program that interacts with the world via a <i>sense-deliberate-act</i> cycle (sometimes called “sense-think-act”). Agents may be physical devices such as robots or exist purely as software, such as automated stock-trading programs.
<b>Artificial Intelligence (AI)</b>	A body of techniques that allow computers to do things that, when humans do them, are considered evidence of intelligence, including natural language processing, recognizing images, and generating content. AI systems use training data and, based on defined or learned patterns, produce outputs like predictions, recommendations, or decisions. AI4K12 defines Five Big Ideas for K-12 AI instruction: Perception, Representation and Reasoning, Learning, Natural Interaction, and Societal Impacts.
<b>Bias</b>	Systematic error in the behavior of a model due to incorrect assumptions or a mismatch between the data the model was trained on and the real world. See <a href="#">Appendix D</a> .
<b>Chatbot</b>	A conversational agent that communicates with humans via a text or voice interface. The simplest chatbots just search for keywords and return scripted replies, but more sophisticated chatbots may fall in the category of intelligent assistants.
<b>Classifier</b>	An algorithm that examines the features of an input pattern and assigns that input to a category (or class). Classifiers have many practical applications, such as spam filtering (using the classes “spam” and “not spam”), or sorting loan applications, where the classes might be low, medium, or high risk. Classifiers can be programmed by hand, but it is more common to train them using machine learning techniques and examples of each class.
<b>Decision Tree</b>	A method of classification based on a series of tests of features of the input. One advantage of decision trees is that their decisions are explainable by looking at the tests that were performed.
<b>Dichotomous Key</b>	An alternative name for a decision tree classifier, used in biology.
<b>Explainable</b>	Explainable AI provides a justification for its decisions by referring to features of the input that led to the decision. For example, if a loan application is categorized as high risk, this decision might be accompanied by an explanation such as a high debt to income ratio or an insufficient employment history.



<b>Fair/Fairness</b>	The decisions of a fair algorithm result in equal outcomes for people who should be treated equally. However, there are multiple, mutually incompatible definitions of fairness (e.g., equal outcomes versus equal risks), so no algorithm can be fair in every sense.
<b>Features</b>	The attributes of an input which are used to reason about it. Features can be numerical values (e.g., age) or categorical values (e.g., state of residence).
<b>Foundational Models</b>	Foundational Models AI models that are trained on massive datasets to learn a broadly useful set of features that allows them to be applied (with additional training) to many different specialized tasks.
<b>Generative AI</b>	Technologies that can generate new text, images, video, and/or audio in response to a prompt. Modern generative AI systems are created using machine learning algorithms run on massive amounts of training data.
<b>Human-in-the-Loop</b>	A system in which humans interact with an automated decision making system to guide its behavior and ensure that its actions are reasonable.
<b>K-Nearest Neighbors (KNN)</b>	A classification algorithm that finds the k-nearest training examples to an input pattern, and assigns that input the same class as the majority of the k neighbors. The definition of "nearest" depends on the features used to represent the patterns.
<b>Large Language Model (LLM)</b>	A neural network trained on massive amounts of text that can be used in a variety of language tasks such as sentence completion, question answering, machine translation, and chatbot functions. Large language models are one of the technologies that make up generative AI.
<b>Logic</b>	A system of formal reasoning using symbolic representations and rules of inference. Much of early AI was based on logic, whereas modern AI relies more on statistical reasoning.
<b>Machine Learning</b>	A subfield of artificial intelligence focusing on combining learning algorithms with training data to create models that can be used to complete tasks. Three important types of machine learning are <i>supervised learning</i> , <i>unsupervised learning</i> , and <i>reinforcement learning</i> . Modern AI applications, such as speech recognition systems, are often constructed using machine learning techniques applied to huge training sets.
<b>Model</b>	A tool for understanding a relationship or phenomenon, used to make predictions or in conjunction with algorithms to process input into output. Models may be based on equations, statistics, rules, agents, neural networks, or other representations. They may be created by humans making sense of data or by computers processing data with algorithms.
<b>Model Card</b>	A model card provides information about an AI model, such as its funding, training data, and performance on benchmarks.
<b>Multimodality</b>	The ability to process multiple types of input, such as text, images, video, and sound.




<b>Natural Language</b>	The language people use to communicate with each other in everyday life. This is distinct from computer languages used for programming or for giving commands to a computer. Computer languages have a rigid vocabulary and syntax, while natural language is fluid and ambiguous, and it includes phenomena such as imagery, humor, sarcasm, and alliteration. Natural language understanding is difficult for computers.
<b>Neural Network</b>	An approach to computing in which many simple processing units are organized into a network to collectively solve a complex problem. Neural networks draw inspiration from theories about how the brain might work, but their components and organizational structure are far simpler than real neural tissue.
<b>Perception</b>	The extraction of meaning from sensory signals. A microphone senses sound and converts it to an electrical signal. Understanding the sound (e.g., recognizing the words being spoken or the music being played) constitutes perception.
<b>Predictor</b>	A reasoner that produces a continuous value as its output, such as estimating the market value of a house based on its features. Predication is also known as regression. Compare classification, which outputs one of a finite set of class labels instead of a continuous value.
<b>Reasoning</b>	The process of making decisions or solving problems. There are many types of reasoning, including classification, prediction, recommendation, planning, and sampling.
<b>Recommender</b>	A reasoner that suggests items the user might like based on what other users who like similar things have liked. Recommendation can be used to suggest items to purchase, ads to show, or stories to include in a news feed.
<b>Reinforcement Learning</b>	A machine learning technique where the training data is labeled with a "reward" signal telling the computer how good its outputs have been. This differs from <i>supervised learning</i> where each training example's label tells the computer exactly what output it should have produced. Reinforcement learning is typically applied to sequential tasks where the reward signal comes only at the end. An example is game playing, where the reward signal comes after the final move and tells the computer whether it won or lost the game. Using reinforcement learning, computers have become experts in domains such as backgammon by playing against themselves for thousands of games.
<b>Representation</b>	An encoding of information in a way that is useful for reasoning. The concept of representation in AI is analogous to that of data structures in computer science.
<b>Search</b>	A reasoning method that involves the systematic exploration of possibilities until a solution is found.
<b>Sensing</b>	Translating a physical phenomenon into an electrical signal that can be measured and acted upon.



<b>Supervised Learning</b>	A machine learning technique where the training data consists of input examples plus the desired output for each input. The desired outputs are called labels, and the data is called labeled data. For example, the training data could be pictures, some of which contain cats, and the labels could be 1 for cat or 0 for no cat. The learning algorithm goes through the training set repeatedly, and it slowly adjusts the model's parameters to make it more likely to produce the correct output for each input.
<b>Training Data</b>	A collection of examples that can be used by a machine learning algorithm to construct a reasoner. Training data is labeled if we are told the correct class of each example. Data can also be unlabeled, in which case it may be used for clustering.
<b>Transparency</b>	Disclosing the details of a system's design or operation. For example, in an automated decision making system, transparency may require disclosure of the training data used to build the system, so users can assess whether the data is likely to be biased.
<b>Unsupervised Learning</b>	A machine learning technique that finds structure in unlabeled data. One example is <i>clustering</i> , where the computer examines a collection of examples and groups them into categories based on perceived similarity. In a sense, the computer "discovers" the categories; they are not given to it. (If each example came labeled with its correct category, this would be <i>supervised learning</i> ).

Photo: Convening participants discuss trends and patterns in AI curriculum.



**APPENDIX F:****Prioritized Foundational K-12 AI Learning Outcomes**

Due to the breadth of computer science and existing curriculum demands, the project team identified a prioritized set of foundational AI learning outcomes for all K-12 students. These priorities are indicated as highlighted rows in Section 3 and summarized in the following table.

<b>Subtopic</b>	<b>Grades K-2</b>	<b>Grades 3-5</b>	<b>Grades 6-8</b>	<b>Grades 9-12</b>
<b>The Human Role in Creating AI</b>	Understand that AI is a tool created by humans to make decisions or to generate something (e.g., an image).	Describe the roles of humans in the creation of AI.	Describe the roles that humans play (including in data curation and labeling) in creating and refining AI models.	Evaluate and analyze the roles of humans and human decision-making in the creation of AI.
<b>Reasoning</b>	Explain how binary choices (e.g., up/down, on/off, under/over) can be used to make decisions that lead to a specific goal by either a human or a machine.	Train a model that can make decisions based on defined criteria (e.g., a dichotomous key to determine which movie to see).	Identify the kinds of AI models (e.g., classifier, predictor, recommender) people interact with in their daily lives.	Describe different types of AI algorithms and models, and compare and contrast the strengths and limitations of their reasoning.
<b>Data (in Machine Learning)</b>	Explore how AI models learn from data.	Explore the relationship between the properties of training data (e.g., size, features, biases) and an AI model's output.	Describe the ways that bias can be introduced and mitigated in an AI model.	Evaluate the data used to solve a problem, including its source(s) and whether privacy is protected, if/how the data has been processed, data quality (e.g., accuracy, reliability, validity), what the data represents, and biases.
<b>Building and Using AI Models</b>	Use data to construct a model for making decisions (e.g., a decision tree to determine what to wear based on the weather).	Using a dataset, develop an AI model to classify inputs.	Using a dataset and a machine learning pipeline, develop an AI model, and consider the impact of the model on various users.	Using a dataset and a systematic process, develop an AI model to generate for classification or prediction, and articulate the assumptions made at each of these steps: (1) develop a question solvable with AI, (2) collect or curate data, (3) evaluate the data, (4) train an AI model on the data, (5), evaluate the model, and (6) iteratively improve the model.
<b>Ethical Evaluation of AI Systems</b>	Explore how an AI system can help and harm different groups at the same time.	Investigate examples of AI, considering differences in experience by different people in different contexts.	Describe the properties, biases, and assumptions of various kinds of AI models (e.g., classifier, predictor, recommender).	Evaluate the design, motivation, outcomes, and potential impacts of AI systems using ethical design criteria and/or ethical frameworks.
<b>Societal Impacts</b>	Explore how some people use AI in their jobs and in their communities.	Explore ways in which some jobs involve the creation and/or use of AI.	Identify the intended and unintended impacts of AI on society — including government, education, entertainment, culture, careers, and national security — while considering how these impacts may differ among diverse communities.	Evaluate the intended and unintended impacts of AI on society (e.g., deep fakes, job loss) — including government, education, entertainment, culture, careers, and national security — while considering how these impacts may differ among diverse communities.





[csteachers.org](http://csteachers.org)



[ai4k12.org](http://ai4k12.org)