Teaching Artificial Intelligence in K-12

David S. Touretzky, Computer Science Department, Carnegie Mellon University
Fred Martin, Associate Dean, University of Massachusetts Lowell

Raspberry Pi Research Seminar
January 11, 2022

Supported by NSF DRL-1846073

AI4K12.org Mission

- Develop national guidelines for teaching AI in K-12
  - Modeled after the CSTA standards for computing education.
  - Four grade bands: K-2, 3-5, 6-8, and 9-12
  - What should students know?
  - What should students be able to do?

- Develop a curated AI resource directory for K-12 teachers
- Foster a community of K-12 AI resource developers

AI4K12.org Steering Committee

Dave Touretzky
Carnegie Mellon
AI for K-12 Working Group Chair

Christina Gardner-McCune
University of Florida
AI For K-12 Working Group Co-Chair

Fred Martin
UMass Lowell
CSTA Chair of Board of Directors

Deborah Seehorn
Co-Chair of CSTA Standards Committee

AI4K12.org

K–12 Teacher Working Group Members

Grades K-2
Vicky Sedgwick (Lead)
Susan Amsler-Akacem
Dr. April DeGennaro

Grades 3-5
Kelly Powers (Lead)
Dr. Marlo Barnett
Dr. Phillip Eaglin
Brian Stamford
Alexis Cobo

Grades 6-8
Padmaja Bandaru (Lead)
Charlotte Dungan
Minsoo Park
Juan Palomares
Josh Caldwell
Sheena Vaidyanathan

Grades 9-12
Dianne O'Grady-Cunniff (Lead)
Jared Amalong
Dr. Smadar Bergman
Kate Lockwood
John Chapin

Five Big Ideas in AI

- Organizing framework for the K-12 guidelines.
- 5 Big Ideas are enough to cover the richness of the field, but small enough to be manageable by teachers.
- CSTA experience shows 5 is a good number.
- Not necessarily the way AI practitioners view their field, but appropriate for the needs of the K-12 audience.

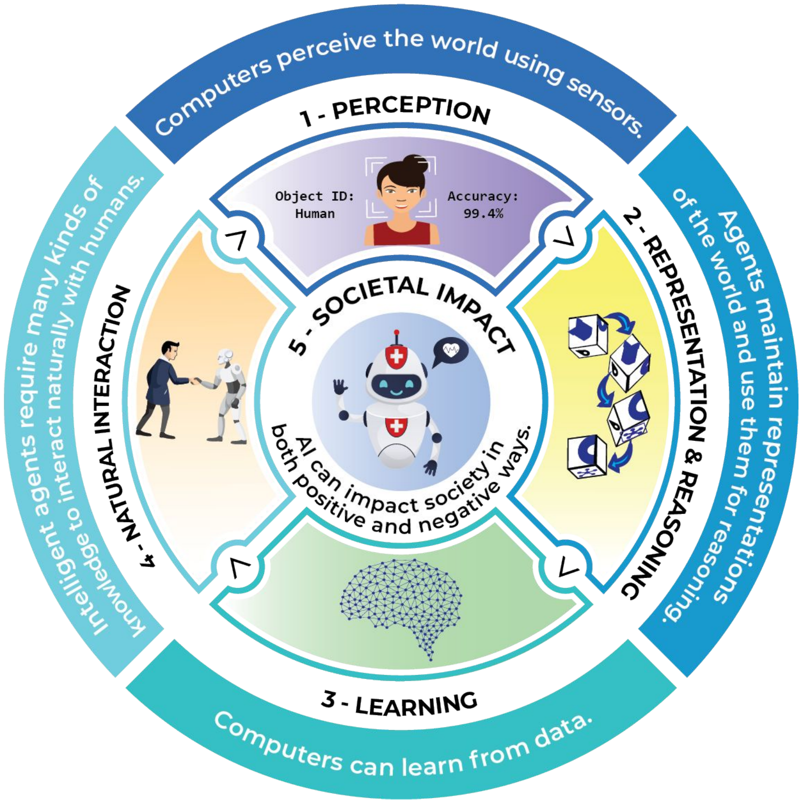

Five Big Ideas in Artificial Intelligence

1. Perception
Computers perceive the world using sensors. Perception is the process of extracting meaning from sensory signals. Making computers "see" and "hear" well enough for practical use is one of the most significant achievements of AI to date.

2. Representation & Reasoning
Agents maintain representations of the world and use them for reasoning. Representation is one of the fundamental problems of intelligence, both natural and artificial. Computers construct representations using data structures, and these representations support reasoning algorithms that derive new information from what is already known. While AI agents can reason about very complex problems, they do not think the way a human does.

3. Learning
Computers can learn from data. Machine learning is a kind of statistical inference that finds patterns in data. Many areas of AI have progressed significantly in recent years thanks to learning algorithms that create new representations. For the approach to succeed, tremendous amounts of data are required. This "training data" must usually be supplied by people, but is sometimes acquired by the machine itself.

4. Natural Interaction
Intelligent agents require many kinds of knowlege to interact naturally with humans. Agents must be able to converse in human languages, recognize facial expressions and emotions, and draw upon knowledge of culture and social conventions to infer intentions from observed behavior. All of these are difficult problems. Today's AI systems can use language to a limited extent, but lack the general reasoning and conversational capabilities of even a child.

5. Societal Impact
AI can impact society in both positive and negative ways. AI technologies are changing the ways we work, travel, communicate, and care for each other. But we must be mindful of the harms that can potentially occur. For example, biases in the data used to train an AI system could lead to some people being less well served than others. Thus, it is important to discuss the impacts that AI is having on our society and develop criteria for the ethical design and deployment of AI-based systems.

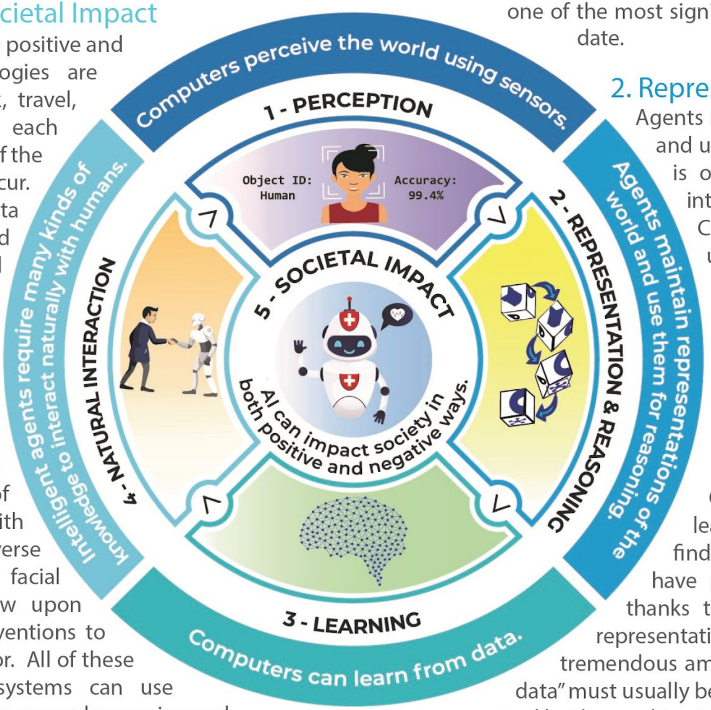

The AI for K-12 Initiative is a joint project of the Association for the Advancement of Artificial Intelligence (AAAI) and the Computer Science Teachers Association (CSTA), funded by National Science Foundation award DRL-1846073

AI4K12.org

11x14 Poster available for free download at AI4K12.org

Widespread Adoption of the Five Big Ideas

- Now being referenced by multiple curriculum developers in the US and elsewhere.
- Big ideas poster is available in 17 languages.

Arabic, Chinese, English, French, German, Hebrew, Hindi, Italian, Japanese, Korean, Portugese, Spanish, Slovenian, Tamil, Thai, Turkish

Chinese

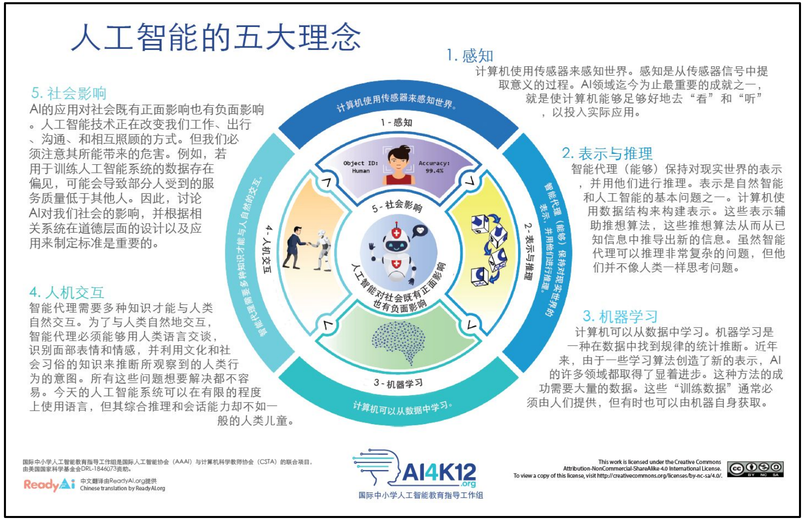

Korean

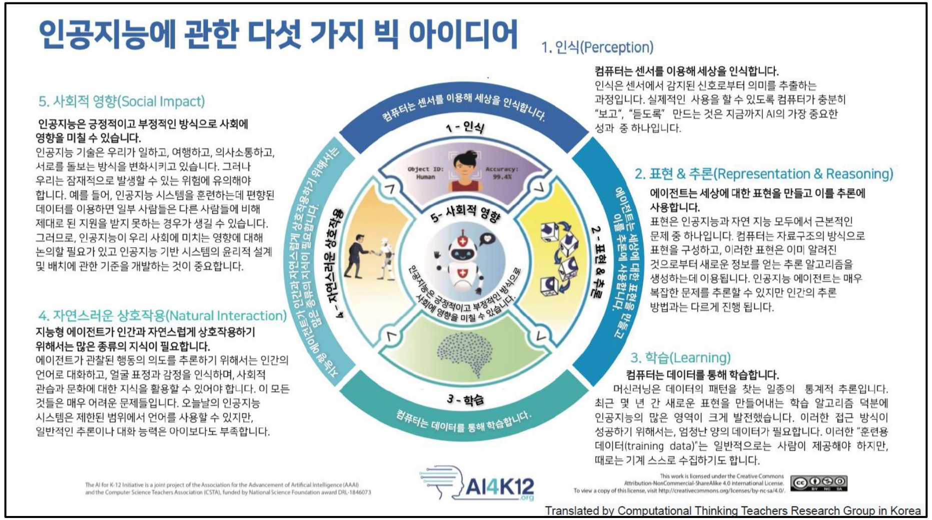

Big Idea #1: Perception

Computers perceive the world using sensors.

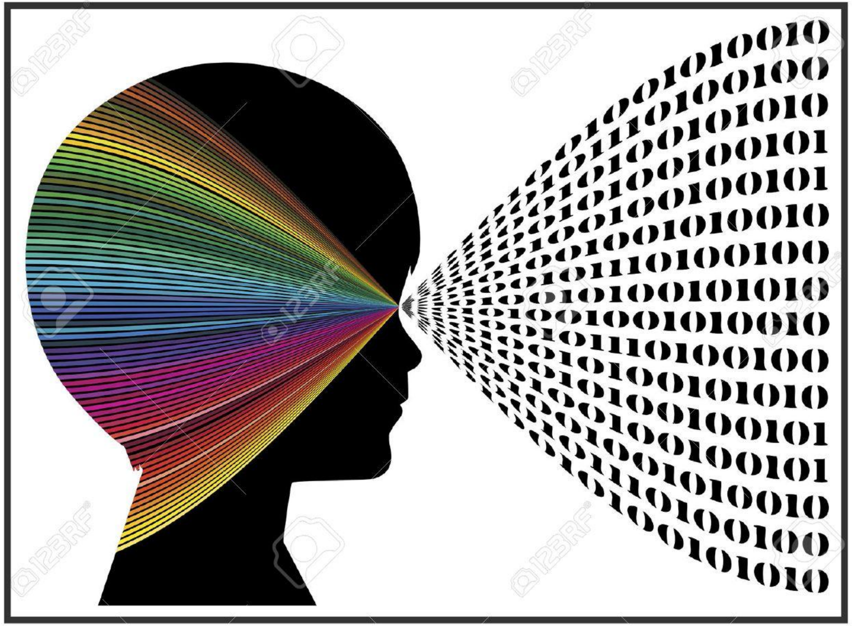

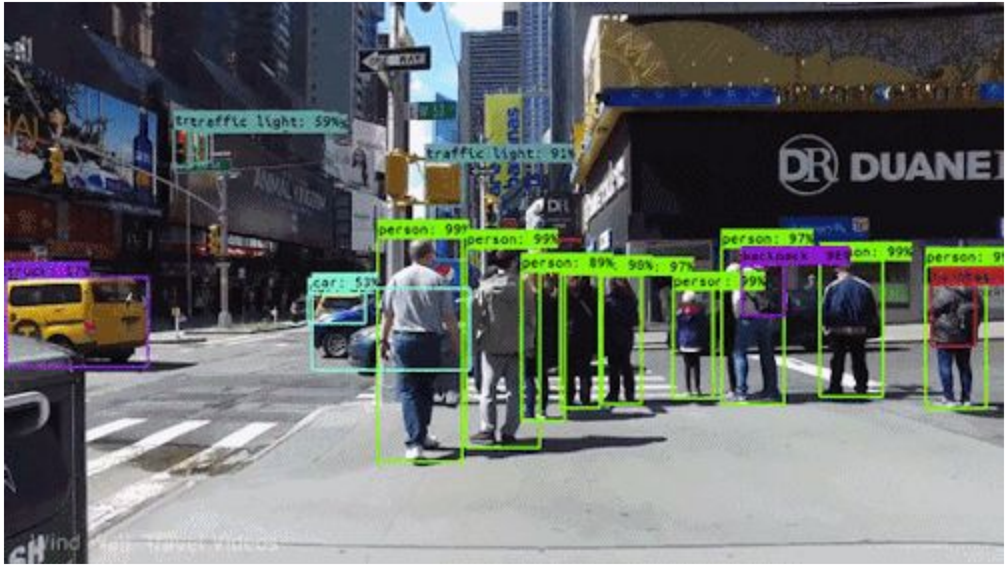

Perception is the extraction of *meaning* from sensory signals.

Big Idea #1 – What should students be able to do?

Grades K–2 (5 to 8 years-old)

- **Identify sensors on computers, robots, and intelligent appliances.**
- Interact with intelligent agents such as Alexa or Siri.

Grades 3–5 (9 to 11 yo)

- Describe how sensor inputs are used in perception.
- **Build an application using perception (e.g., with Scratch plugins or Calypso).**

Grades 6–8 (12 to 14 yo)

- Explain how sensor limitations affect computer perception.
- **Explain that perception systems may draw on multiple algorithms as well as multiple sensors.**
- Build an application using multiple sensors and types of perception (e.g., with Scratch plugins or Calypso).

Grades 9–12 (15 to 18 yo)

- Describe the domain knowledge underlying different forms of computer perception.
- **Demonstrate speech recognition difficulty in handling homophones and other types of ambiguity.**

Big Idea #2: Representation and Reasoning

Agents maintain representations of the world, and use them for reasoning.

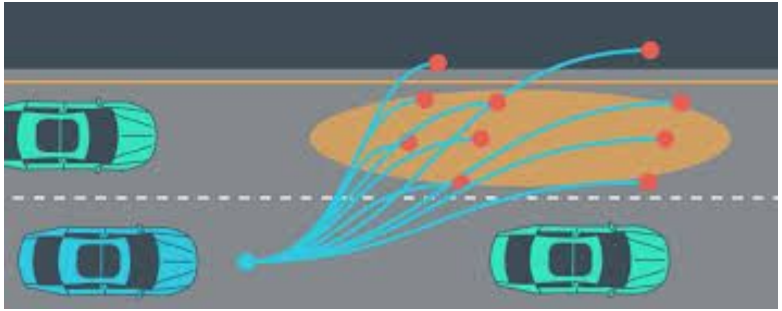

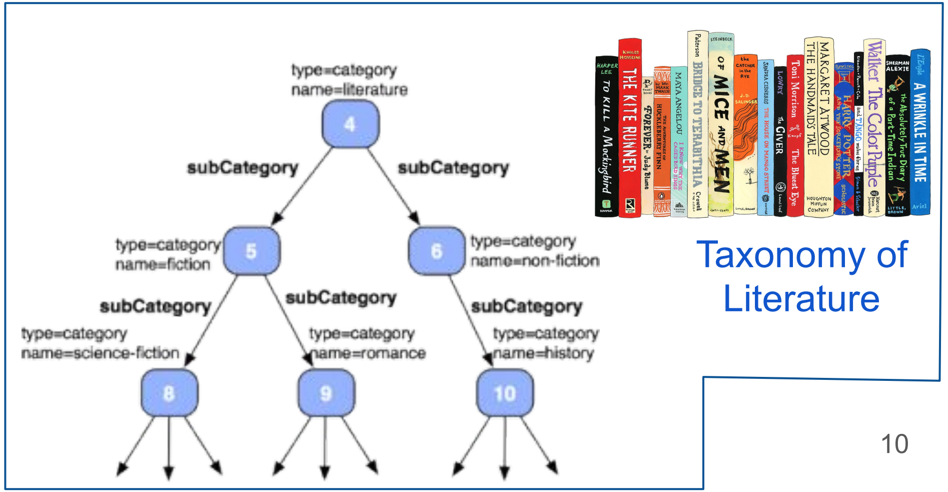

Taxonomy of Literature

Big Idea #2 – What should students be able to do?

Grades K–2 (5 to 8 years-old)

- **Draw a map of the classroom or school and compare the map to the actual room or school building and grounds.**
- Use a decision tree to make a decision

Grades 3–5 (9 to 11 yo)

- **Create/design a representation of an (animal) classification system using a tree structure.**
- Describe how AI representations support reasoning to answer questions

Grades 6–8 (12 to 14 yo)

- **Design a graph model of their home or locations in their community and apply reasoning to determine the shortest path to key locations on their map**

Grades 9–12 (15 to 18 yo)

- **Draw a search tree for tic-tac-toe**
- Describe the differences between types of search algorithms

Big Idea #3: Learning

Computers can learn from data.

Big Idea #3 – What should students be able to do?

Grades K–2 (5 to 8 years-old)

- **Learn from patterns in data with “unplugged” activities**
- Use a classifier that recognizes drawings.
- Use Google Autodraw or Cognimates Train Doodle to investigate how training sets work to identify images and discuss how the program knows what they are drawing

Grades 3–5 (9 to 11 yo)

- **Modify an interactive machine learning project by training its model.**
- Describe how algorithms and machine learning can exhibit biases.

Grades 6–8 (12 to 14 yo)

- **Identify bias in a training data set and extend the training set to address the bias**
- Simulate the training of a simple neural network

Grades 9–12 (15 to 18 yo)

- **Train a neural net (1-3 layers) *TensorFlow Playground***
- Trace and experiment with a simple ML algorithm

Big Idea #4: Natural Interaction

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

Humans are among the hardest things for AI agents to understand.

Big Idea #4 – What should students be able to do?

Grades K–2 (5 to 8 years-old)

- Identify words in stories that have positive and negative connotations.
- **Recognize and label facial expressions into appropriate emotions (happiness, sadness, anger) and explain why they are labeled the way they are.**

Grades 3–5 (9 to 11 yo)

- **Identify how humans combine multiple inputs (tone, facial expressions, posture, etc) in order to understand communication.**
- Describe some tasks where AI outperforms humans and tasks where it does not.

Grades 6–8 (12 to 14 yo)

- **Construct a simple chatbot.**
- Explain and give examples of how language can be ambiguous.
- Reason about the nature of intelligence, and identify approaches to determining whether an agent is or is not intelligent.

Grades 9–12 (15 to 18 yo)

- Demonstrate how sentence parsers handle ambiguity.
- Explore the Google Knowledge Graph
- **Identify and debate the issues of AI and consciousness.**

Big Idea #5: Societal Impact

“Artificial Intelligence can impact society in both positive and negative ways.”

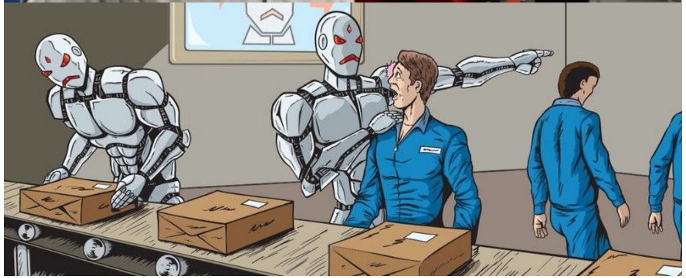

Big Idea #5 – What should students be able to do?

Grades K–2 (5 to 8 yo)

- **Identify common AI applications encountered in their daily lives**
- Discuss whether common uses of AI technology are a good or bad thing

Grades 3–5 (9 to 11 yo)

- Explore how behavior is influenced by bias and how it affects decision making
- **Describe ways that AI systems can be designed for inclusivity**

Grades 6–8 (12 to 14 yo)

- Explain potential sources of bias in AI decision making
- **Understand tradeoffs in the design of AI systems and how decisions can have unintended consequences in the function of a system**

Grades 9–12 (15 to 18 yo)

- Critically explore the positive and negative impacts of an AI system
- **Design an AI system to address social issues (or explain how AI could be used to address a social issue)**

Guidelines for supporting K-12 students

1. **Use transparent AI demonstrations** *that help students see what is going on inside the black box: it's not magic!*
2. **Help students build mental models** of *what is happening under the hood in AI applications.*
3. **Encourage students to develop AI applications** *using AI services.*

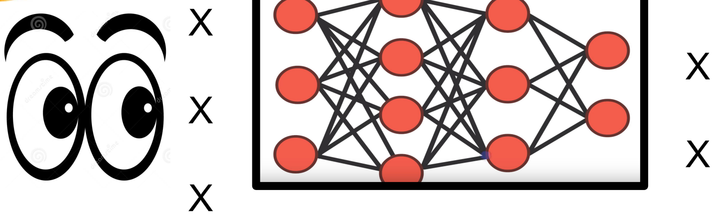

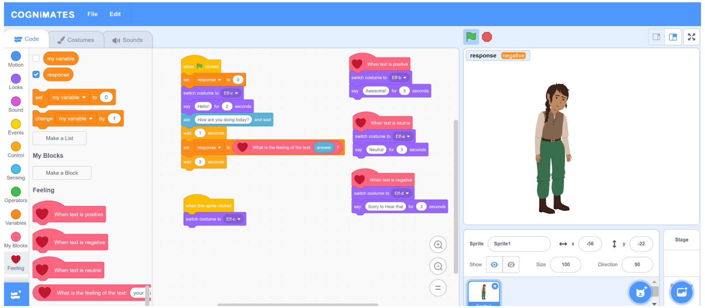

Student Activity Considerations

- **Experiment with AI agents** to investigate their behavior
- Encourage students to **build their own AI applications**
- **Explore case studies of AI-related societal issues** from multiple perspectives

These activities promote understanding of:

- How AI works
- Limitations of AI
- Systems thinking (AI systems are built from smaller components)
- Sources of bias in AI
- Societal impacts of AI systems

What Does AI Thinking Look like in K-12?

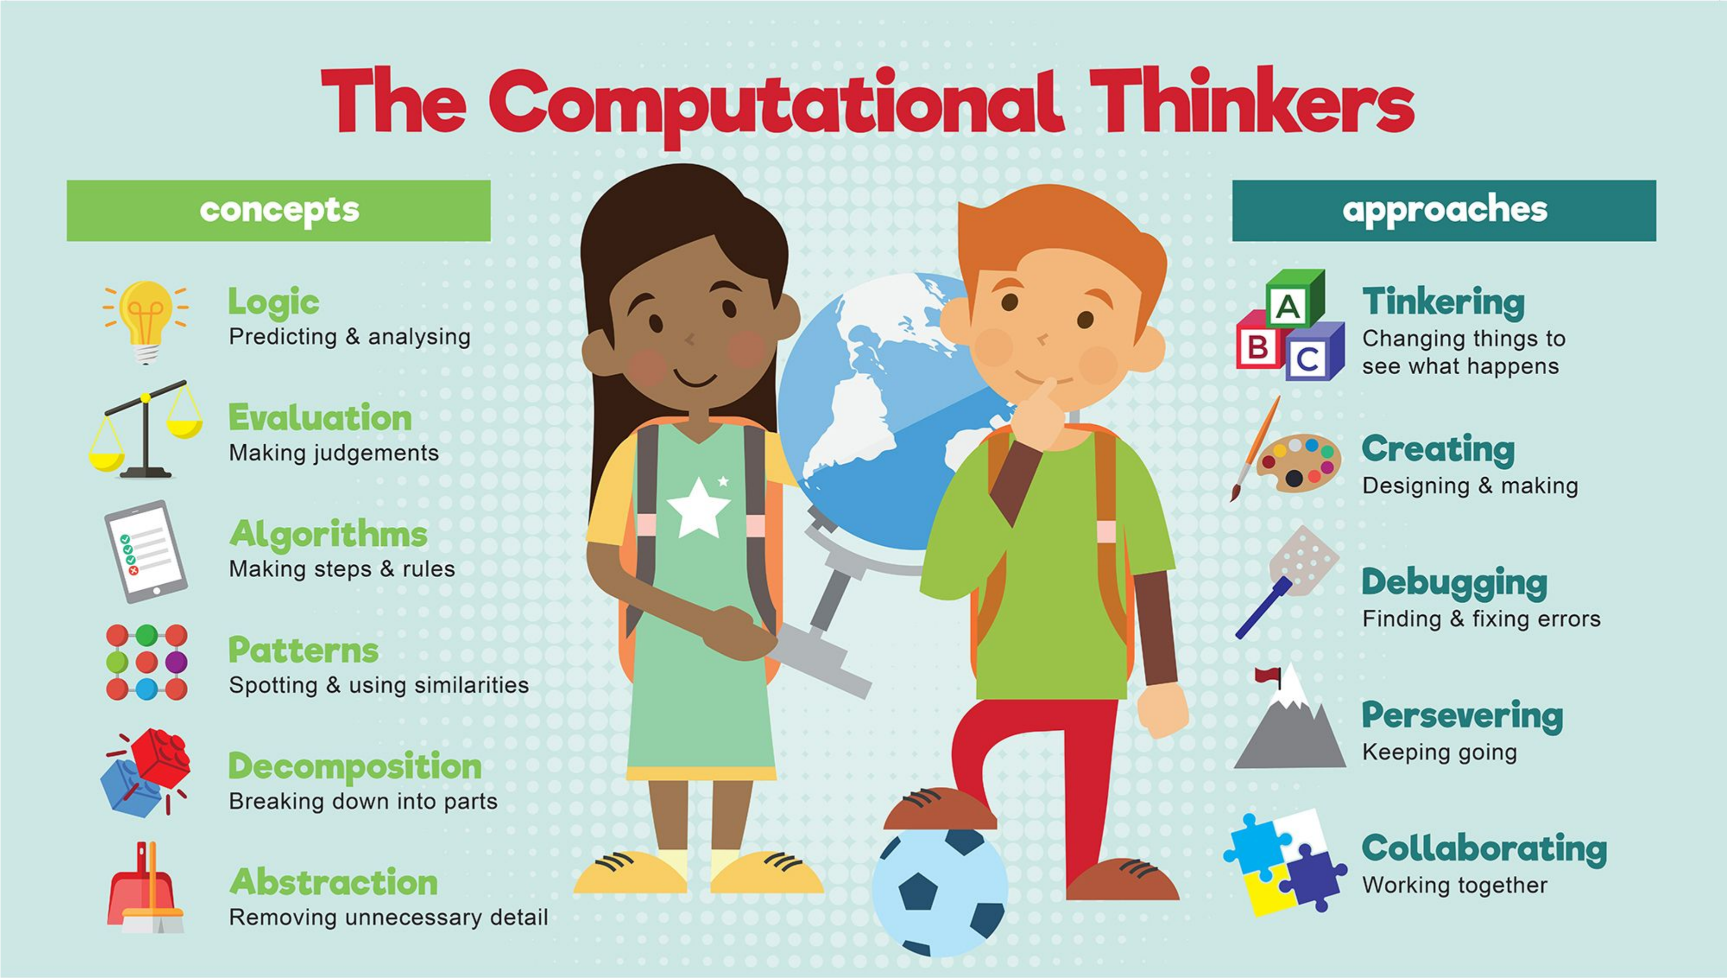

Image from Barefoot Computing, barefootcomputing.org

Computational Thinking

- Logic
- Evaluation
- Problem Decomposition
- Pattern Recognition
- Abstraction
- Algorithms

AI Thinking

- Perception (not just sensing!)
- Reasoning
- Representation
- Machine Learning
- Language Understanding
- Autonomous Robots

How AI Thinking Extends Computational Thinking

AI is built on representation and reasoning.

- Representations are data structures (**abstractions**)
- Reasoners are **algorithms**

So AI draws on the concepts and dispositions of computational thinking.

But AI asks students to consider that **computation can actually be thinking**.

Computational thinking is exactly what humans need when they try to understand how machines can think.

Visual Perception

Computers can see:

- Faces
- Household objects
- Road scenes

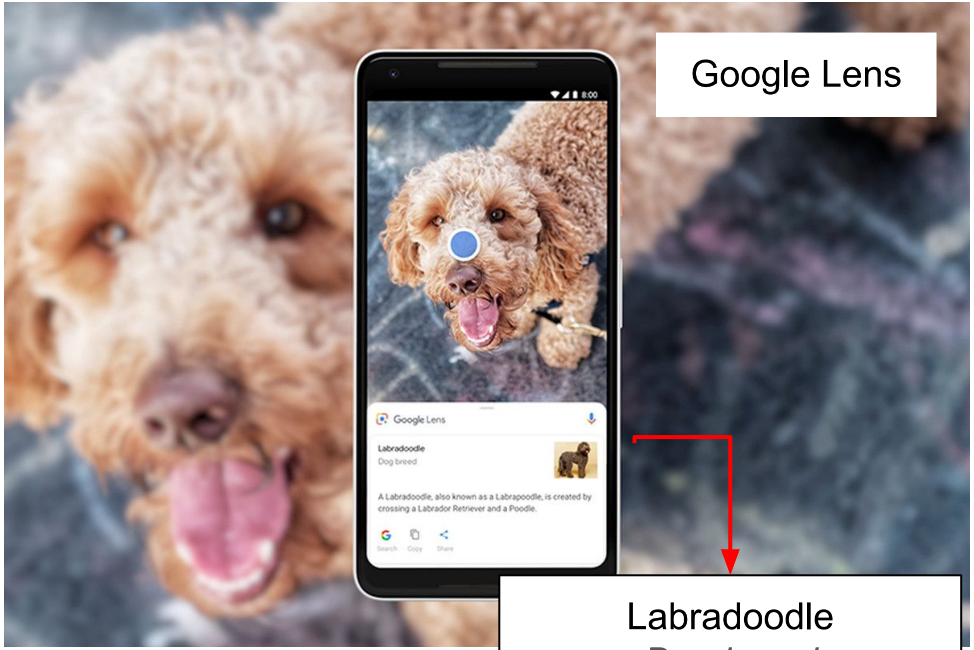

Labradoodle
Dog breed

A Labradoodle, also know as a Labrapoodle, is created by crossing a Labrador Retriever and a Poodle.

I can teach a computer to recognize what I want it to see.

I can make artifacts (programs, devices) that use computer vision.

Levels of visual structure

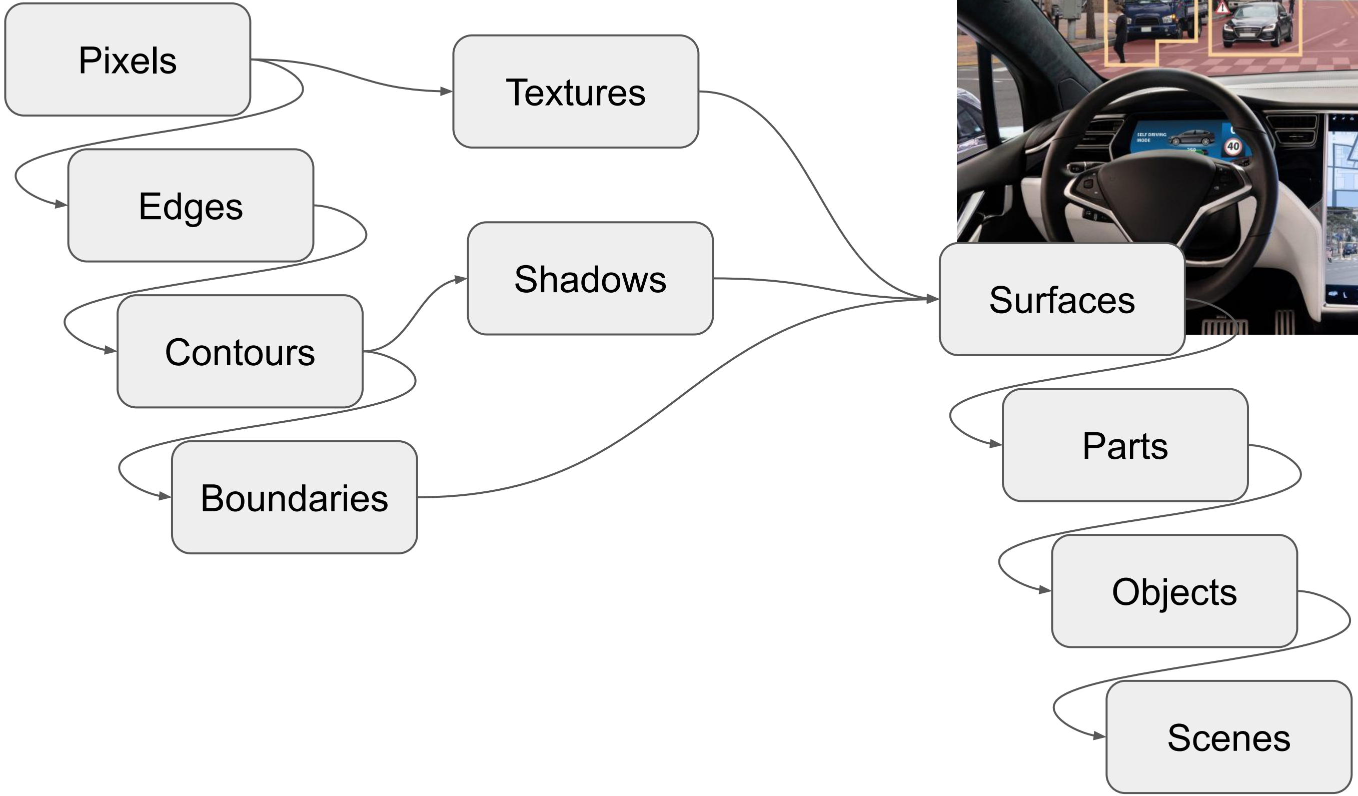

Speech Perception

Computers can understand spoken language.

Lots of knowledge is required to accurately decode the speech signal:

- "They're building their new house over there."

I can make artifacts that understand voice commands.

Levels of representation and linguistic knowledge

Waveform

Spectrogram

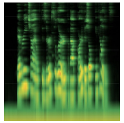

The transformation from signal to meaning takes place in stages, with increasingly abstract features and higher level knowledge applied at each stage.

Sounds
(Phonology)

/ ^ n f o r g ɜ t a b ʊ l # s al t /

Stems + Affixes
(Morphology)

un + forget + able   site/sight

Grammar
(Syntax)

(NP (ADJ “unforgettable”) (N “site”/”sight”))

Meaning
(Semantics)

“unforgettable sight”

Representation

Maps are representations of the world

Robots maintain maps of their environment

Computers build representations to aid their reasoning

Representations are data structures

- Trees
- Graphs
- Feature vectors

I can make representations and manipulate them.

Tesla's World Map

At right is an image from a real self-driving car, a Tesla, showing the road and other nearby vehicles on its world map.

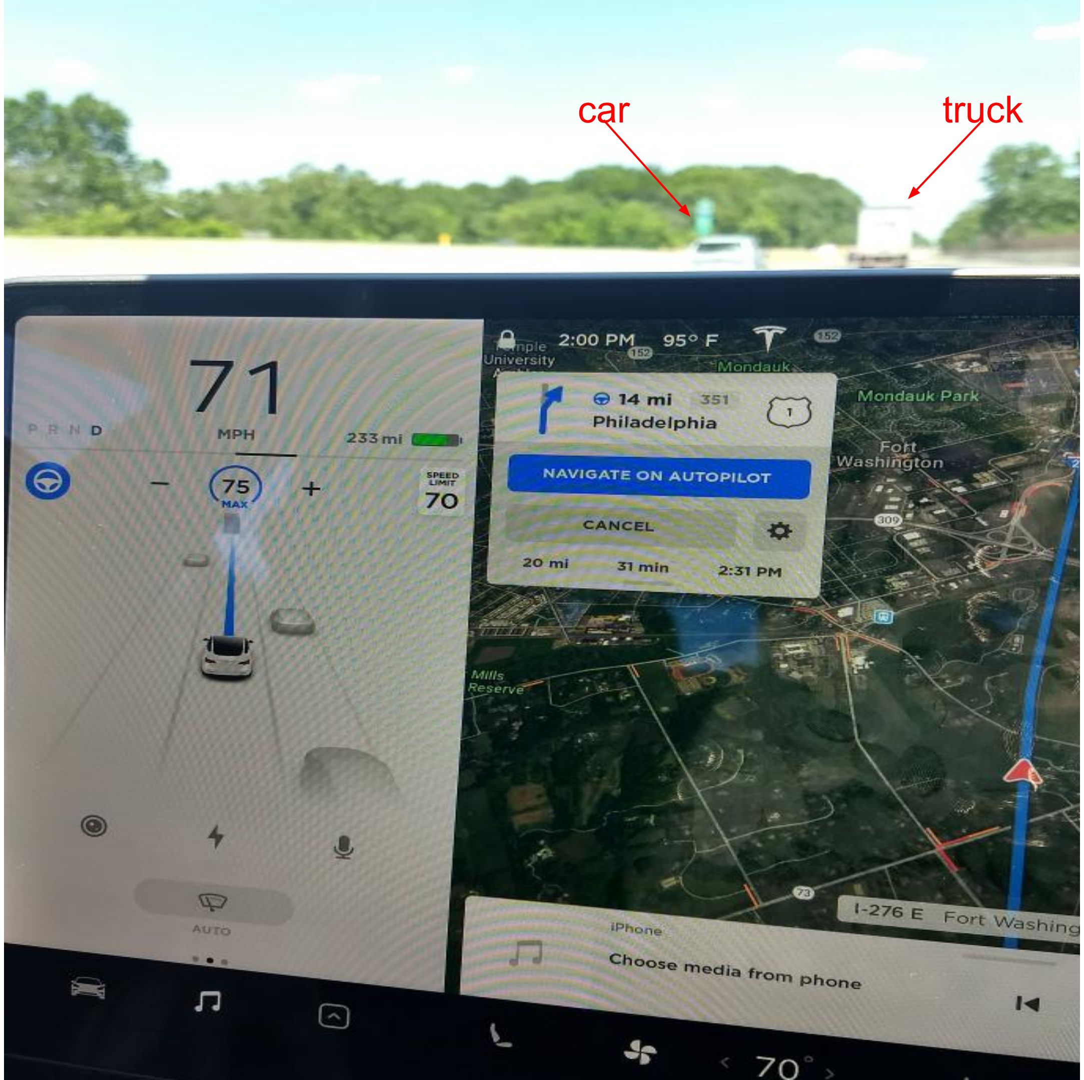

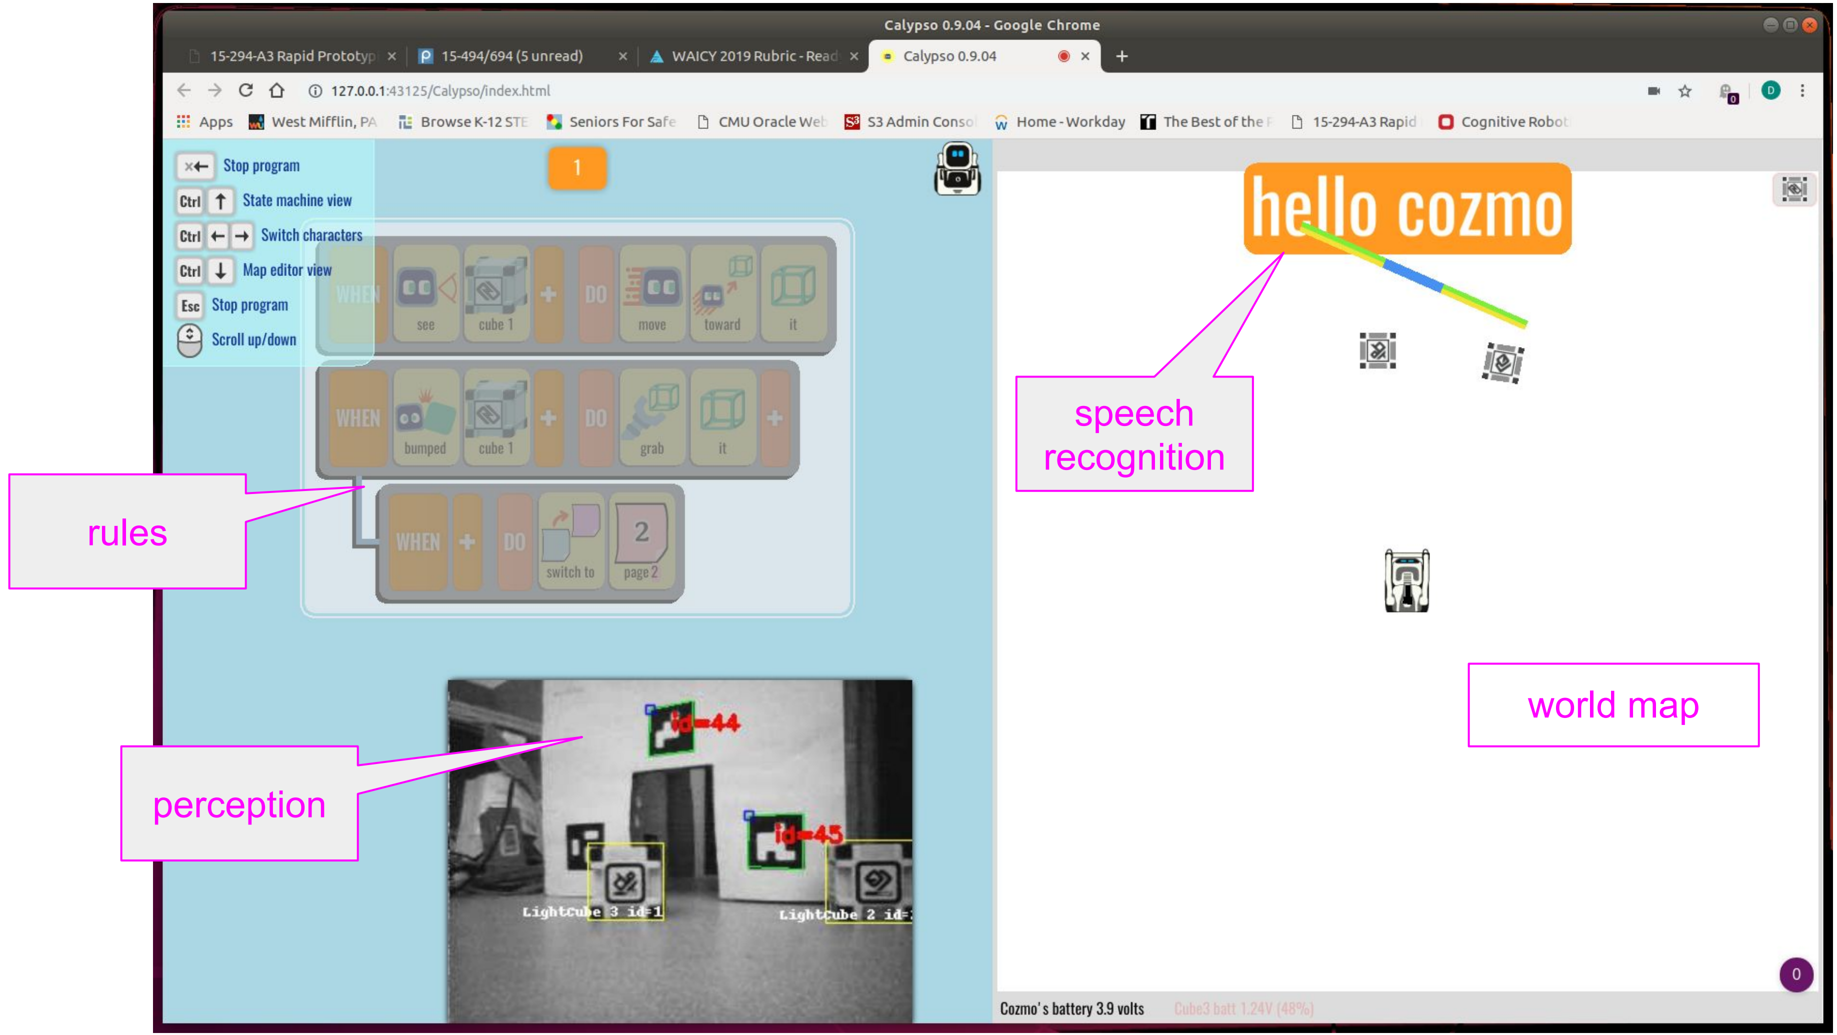
Calypso 0.9.04 - Google Chrome
127.0.0.1:43125/Calypso/index.html
Stop program
Ctrl State machine view
Ctrl Switch characters
Ctrl Map editor view
Esc Stop program
Scroll up/down
1
WHEN see cube 1 DO move toward it
WHEN bumped cube 1 DO grab it
WHEN DO switch to page 2
hello cozmo
Cozmo's battery 3.9 volts
Cube3 batt 1.24V (48%)
rules
perception
speech recognition
world map

Calypso for Cozmo

Calypso

- A robot intelligence framework that combines multiple types of AI:
  - Computer vision
  - Speech recognition
  - Landmark-based navigation
  - Path planning
  - Object manipulation

- Rule-based language inspired by Microsoft's Kodu Game Lab
- Teaches AI thinking
- Web sites:
  - https://Calypso.software (Cozmo robot version)
  - https://calypso-robotics.com (free simulator version runs in the browser)

Learning: Computers Can Learn From Data

Computers don’t learn the way people do.

Machine learning <u>constructs</u> a reasoner.

The learning algorithm uses training data to **adjust the reasoner’s internal representations** so that it produces the right answers.

What are the internal representations?

- For a decision tree, the representations are the nodes of the tree.
- For a neural network, the representations are the weights.

I can use machine learning to train a reasoner.

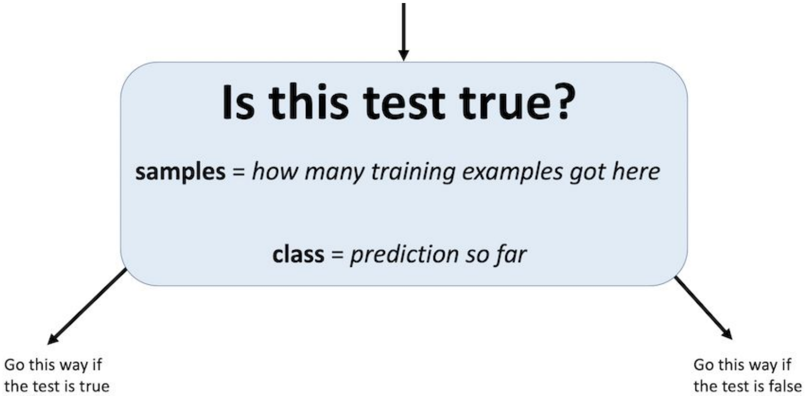

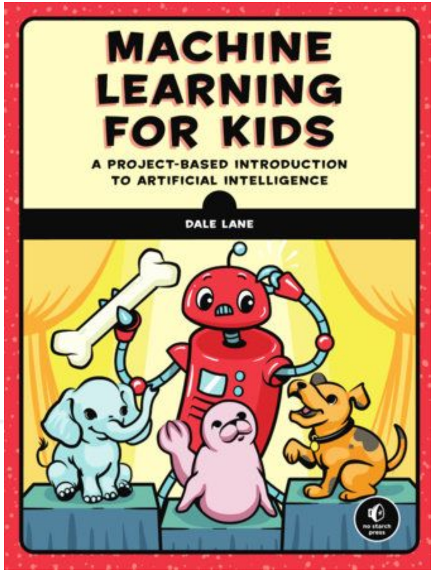

https://machinelearningforkids.co.uk

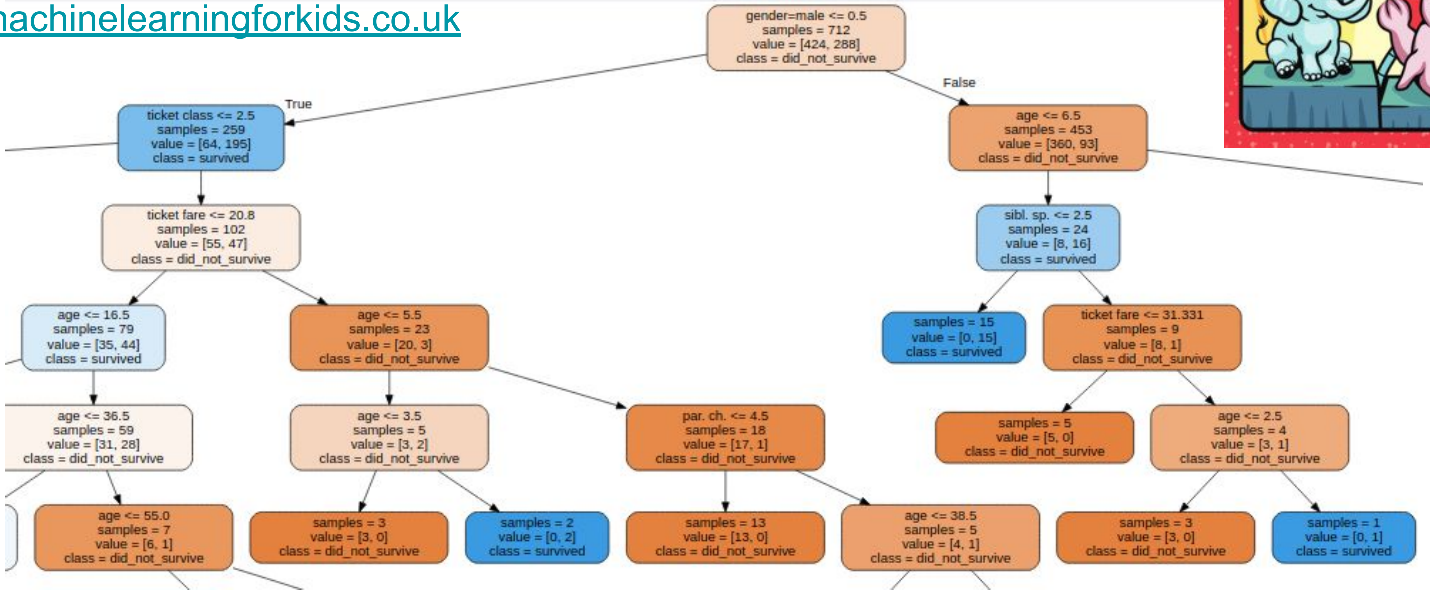

The State of K-12 AI Education in Your State: A Planning Workshop

David Touretzky, CMU & Christina Gardner-McCune, UF

https://ai4k12.org/news/the-state-of-k-12-ai-education-in-your-state-a-planning-workshop/

January 28-29, 2021
141 Participants

- 27 States
- 3 Territories

15 State Completed Plans (Jan)

CA, CT, FL, GA, HI, IL, IN, MD, MA, MS, NC, OH, PA, SC, TX,

2 New States & 2 Territories Completed Plans

NM, VA
Puerto Rico, US Virgin Islands

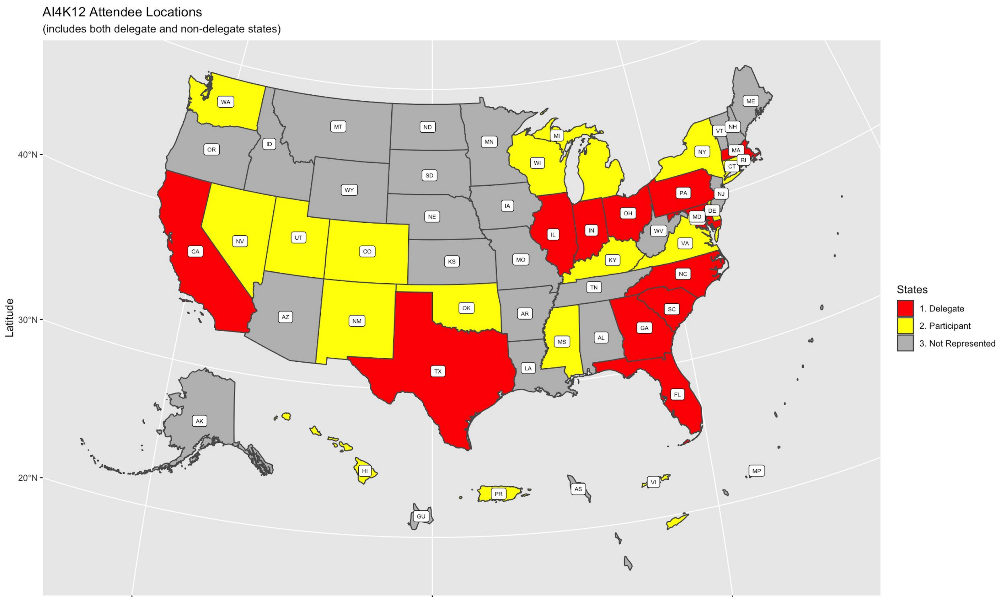

- 16 States are currently advancing their K-12 AI Implementation Plan
- 5 States developed CTE AI Course frameworks

K-12 AI Education Efforts World Wide

- United States: AI4K12.org, MIT RAISE, AI4ALL, ISTE, Code.org, many NSF projects (including our own AI4GA)
- China: government mandate that all students will learn about AI. No national standards yet. Many experiments with curriculum; multiple textbooks.
- South Korea: 2022 revised national curriculum includes AI in all grades K-12.
- United Kingdom: ComputingAtSchool advocating for AI education; teacher PD.
- European Union
  - Erasmus+ funding development of an AI curriculum adapted to European high schools
  - Many small experiments taking place in Germany, Italy, Portugal, Spain, etc.

Join the AI4K12 Community and Help Bring AI Education to K-12!

Visit us:

AI4K12.org

Join the mailing list:

Visit aaai.org/Organization/mailing-lists.php

Acknowledgments

AAAI (Association for the Advancement of Artificial Intelligence)

Association for the Advancement of Artificial Intelligence

CSTA (Computer Science Teachers Association)

CSTA

With funding from National Science Foundation ITEST Program (DRL-1846073)

Carnegie Mellon University
School of Computer Science