Beyond Computational Thinking: AI Thinking in K-12

David S. Touretzky
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA

MIT RAISE Seminar Series
November 30, 2021

Supported by NSF DRL-1846073

Importance of K-12 Computing Education

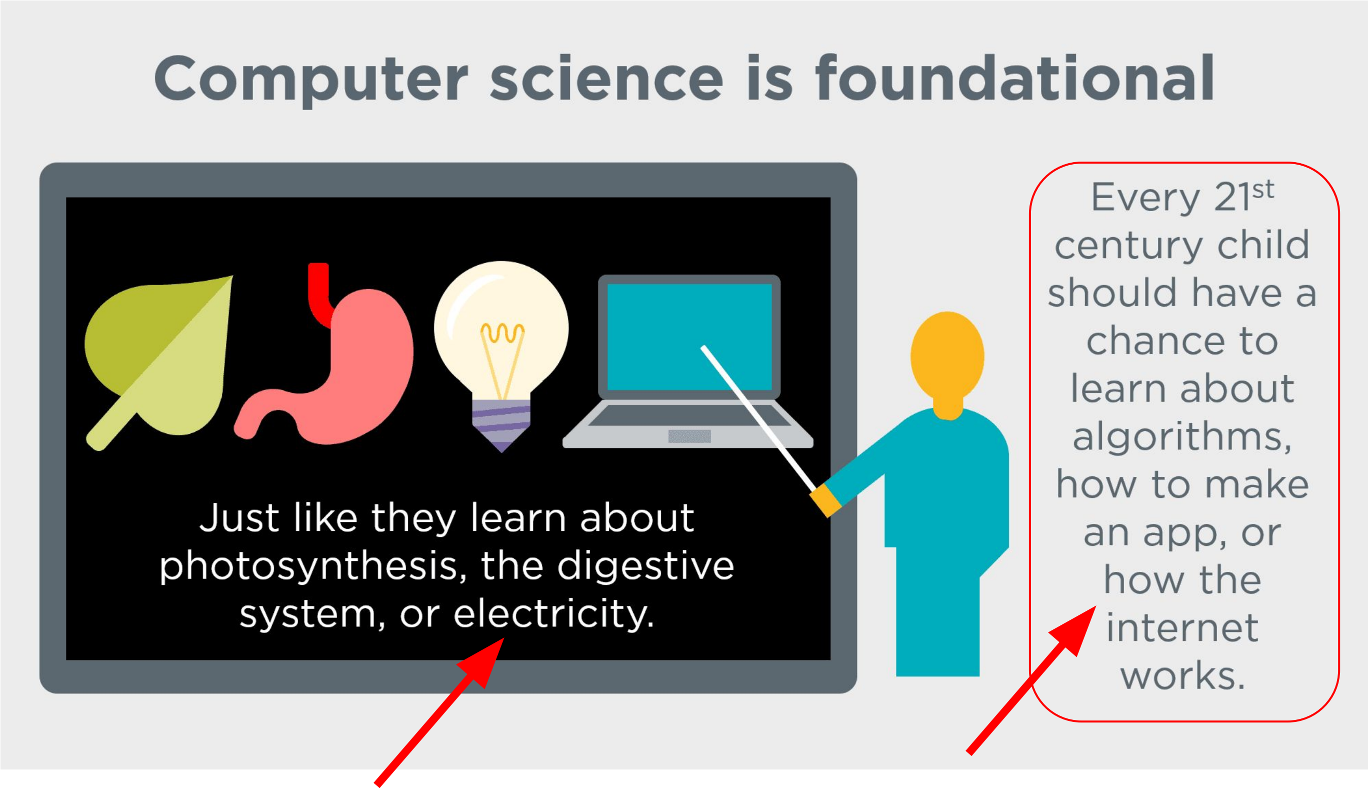

Computer science is foundational

Just like they learn about photosynthesis, the digestive system, or electricity.

Every 21st century child should have a chance to learn about algorithms, how to make an app, or how the internet works.

Code.org

Teaching Computational Thinking

“Computational thinking refers to the thought processes involved in expressing solutions as computational steps or algorithms that can be carried out by a computer.” – from k12cs.org

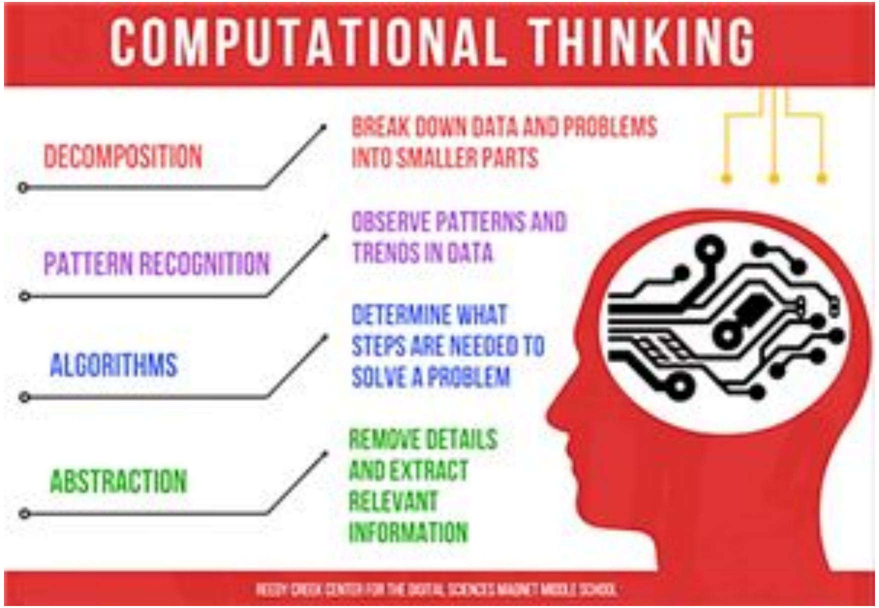

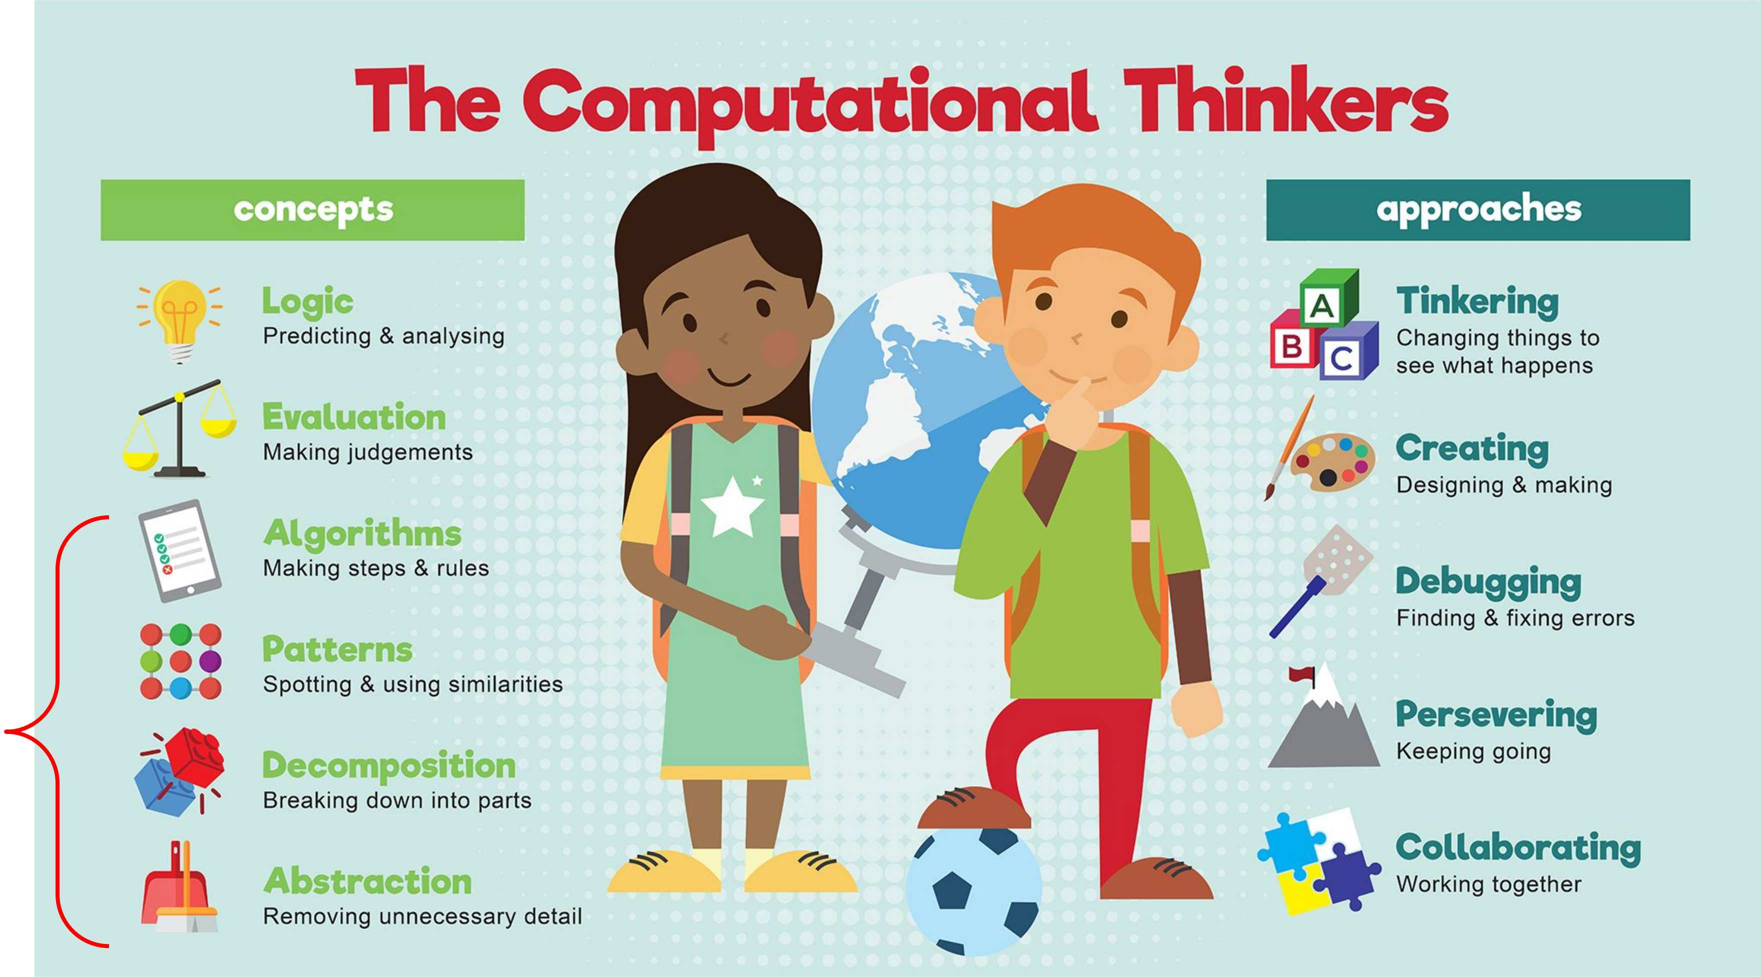

The Computational Thinkers

concepts	
Logic	Predicting & analysing
Evaluation	Making judgements
Algorithms	Making steps & rules
Patterns	Spotting & using similarities
Decomposition	Breaking down into parts
Abstraction	Removing unnecessary detail

approaches	
Tinkering	Changing things to see what happens
Creating	Designing & making
Debugging	Finding & fixing errors
Persevering	Keeping going
Collaborating	Working together

Image from Barefoot Computing, barefootcomputing.org

Importance of K-12 Computing Education

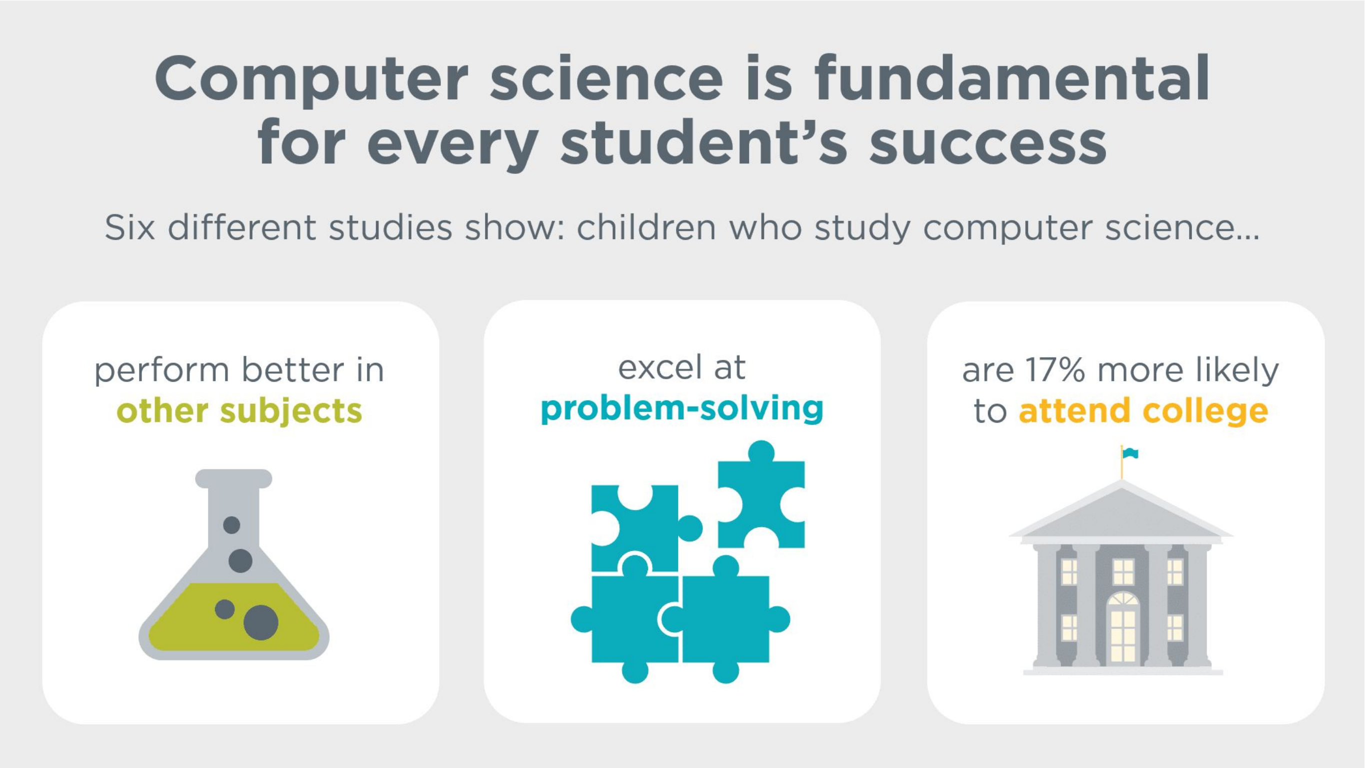

Computer science is fundamental for every student's success

Six different studies show: children who study computer science...

perform better in **other subjects**

excel at **problem-solving**

are 17% more likely to **attend college**

Code.org

Promoting Computing Education in K-12

- CSTA Computer Science Standards released in 2011, revised in 2017
    - Officially adopted by some states; recognized by others
    - Only two sentences about AI
- NSF funding computing education research through programs such as CS10K, CSforAll, STEM+C, ITEST, etc.
- Code.org providing curriculum resources, “Hour of Code” events, teacher training.
- CSforALL.org: “make high-quality computer science an integral part of the educational experience for all K-12 students and teachers…”

But K-12 Computing Education Is Not Yet Universal

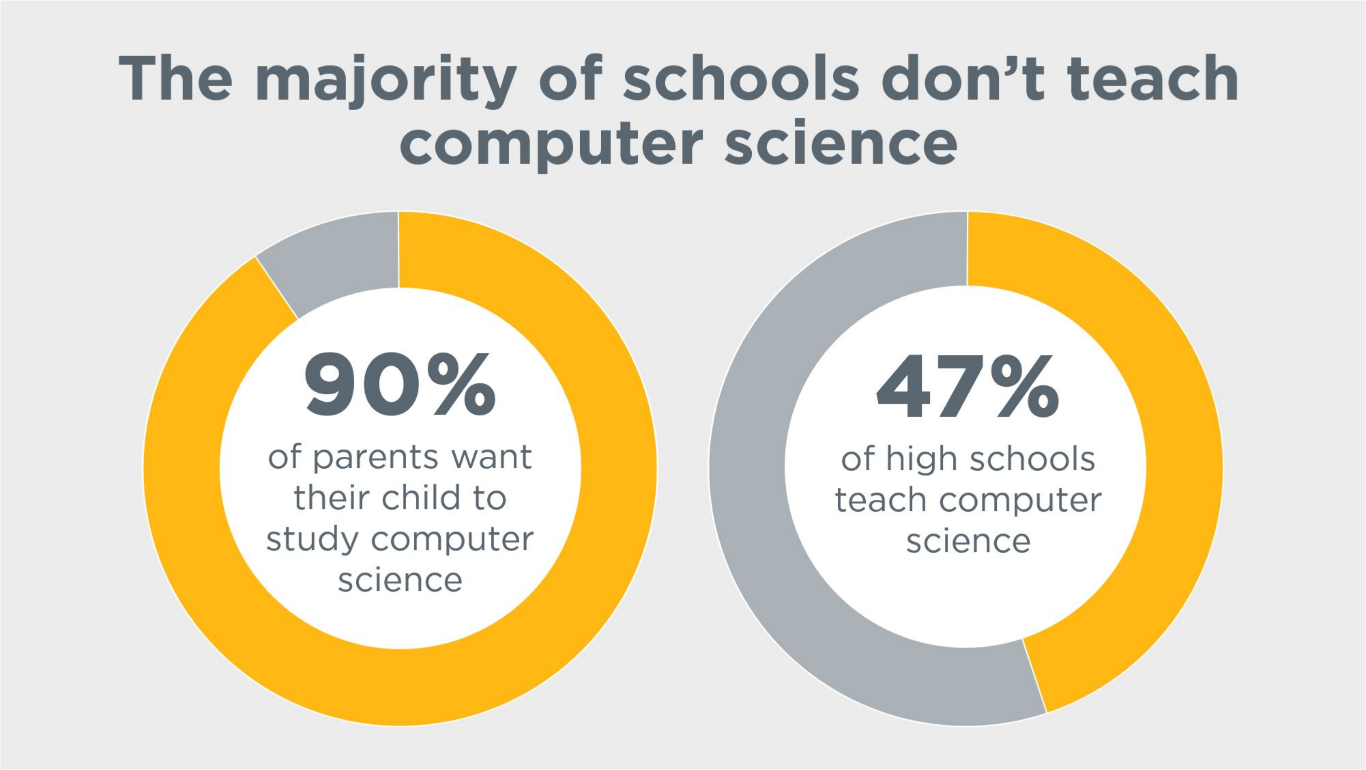

The majority of schools don't teach computer science

90% of parents want their child to study computer science

47% of high schools teach computer science

Code.org

Code.org's 9 policy recommendations to make computer science fundamental to K-12 education

1. **Create a state plan** for K–12 computer science
2. **Define computer science** and establish rigorous K–12 computer science standards
3. **Allocate funding** for rigorous computer science teacher professional learning and course support
4. **Implement clear certification pathways** for computer science teachers
5. **Create programs** at institutions of higher education to offer computer science to preservice teachers
6. **Establish dedicated computer science positions** in state and local education agencies
7. **Require that all secondary schools offer computer science** with appropriate implementation timelines
8. **Allow computer science to satisfy** a core graduation requirement
9. **Allow computer science to satisfy** an admission requirement at institutions of higher education

Code.org

No State Has Universal CS Education

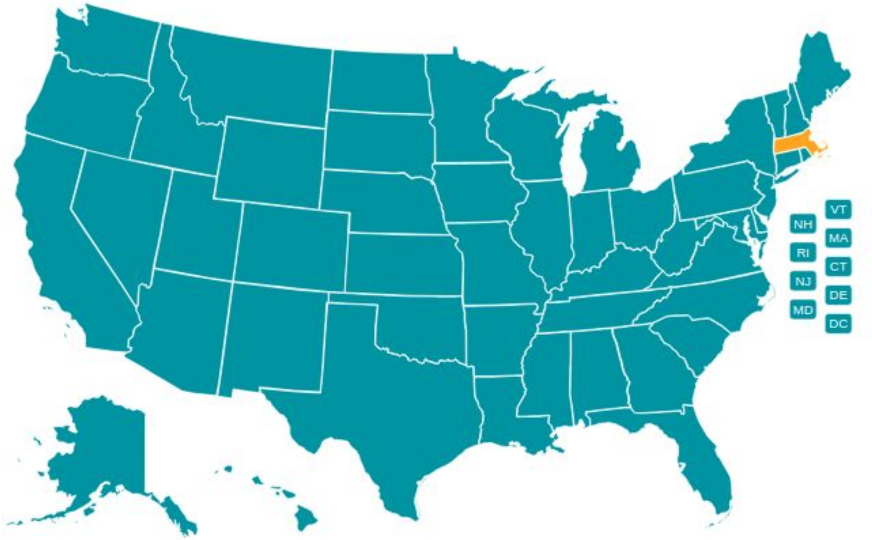

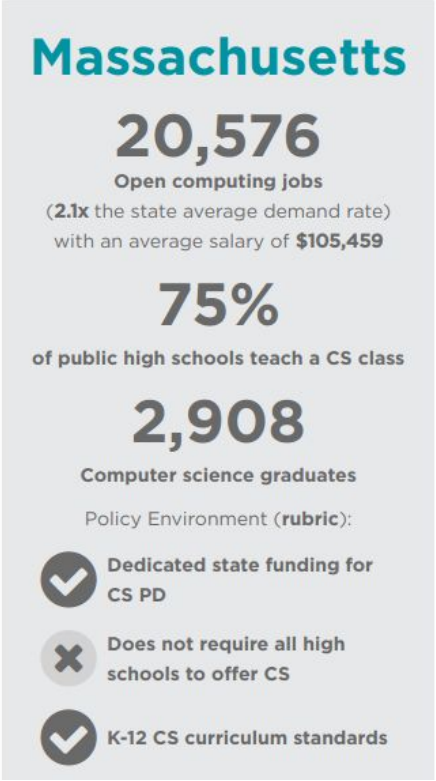

Massachusetts

20,576

Open computing jobs

(**2.1x** the state average demand rate) with an average salary of **$105,459**

75%

of public high schools teach a CS class

2,908

Computer science graduates

Policy Environment (**rubric**):

- ✓ **Dedicated state funding for CS PD**
- ✗ **Does not require all high schools to offer CS**
- ✓ **K-12 CS curriculum standards**

Code.org

No State Has Universal CS Education

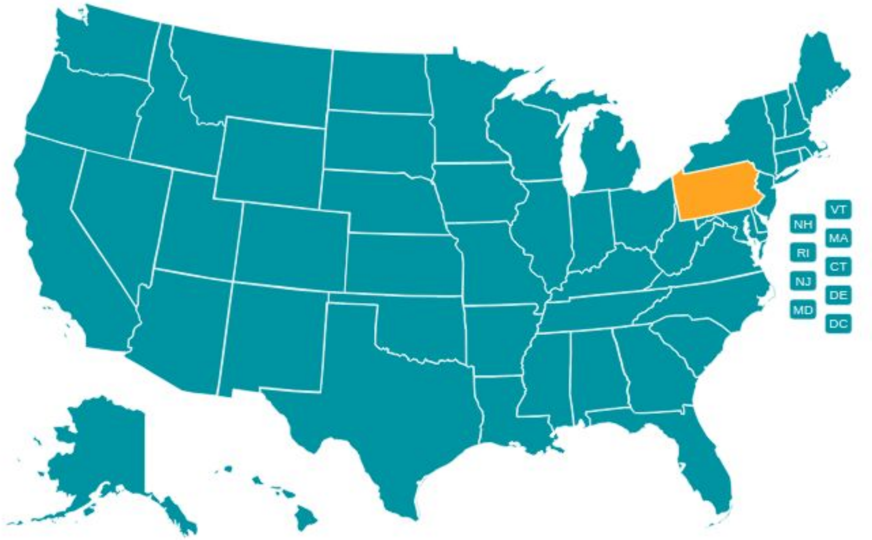

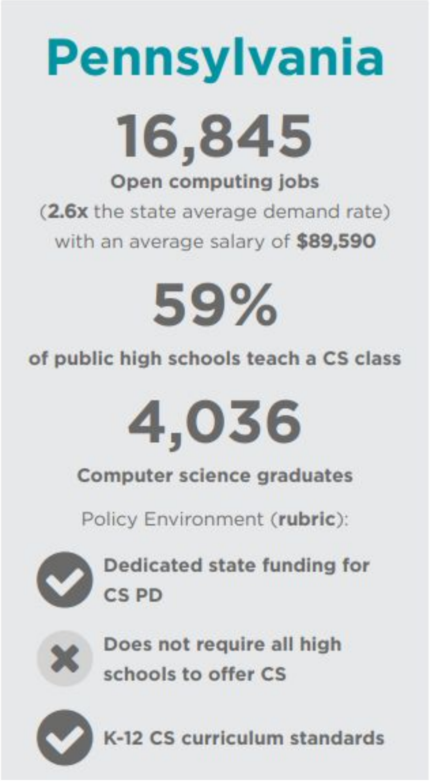

Pennsylvania

16,845

Open computing jobs

(**2.6x** the state average demand rate)
with an average salary of **$89,590**

59%

of public high schools teach a CS class

4,036

Computer science graduates

Policy Environment (**rubric**):

- ✓ **Dedicated state funding for CS PD**
- ✗ **Does not require all high schools to offer CS**
- ✓ **K-12 CS curriculum standards**

Code.org

No State Has Universal CS Education

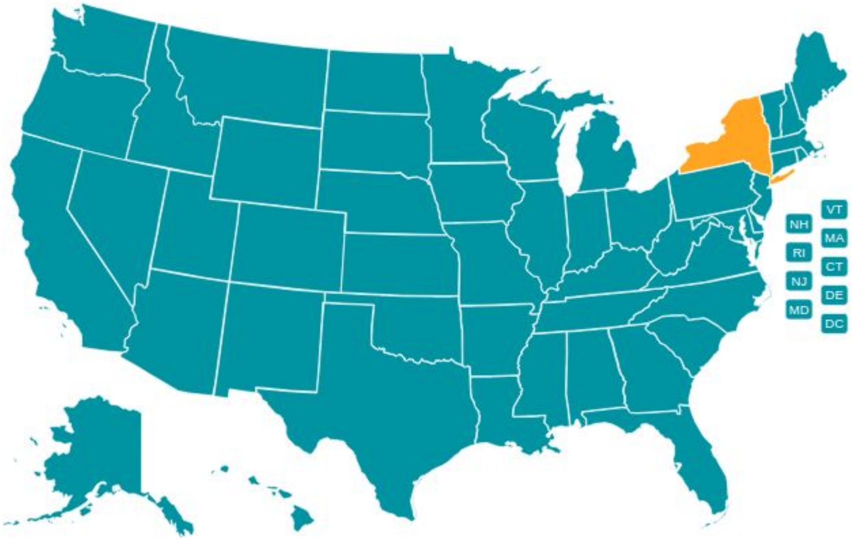

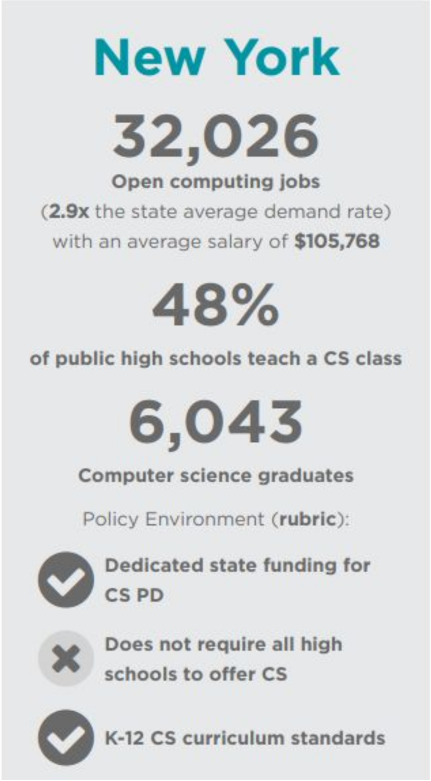

New York

32,026

Open computing jobs

(**2.9x** the state average demand rate)
with an average salary of **$105,768**

48%

of public high schools teach a CS class

6,043

Computer science graduates

Policy Environment (**rubric**):

- ✓ **Dedicated state funding for CS PD**
- ✗ **Does not require all high schools to offer CS**
- ✓ **K-12 CS curriculum standards**

Code.org

No State Has Universal CS Education

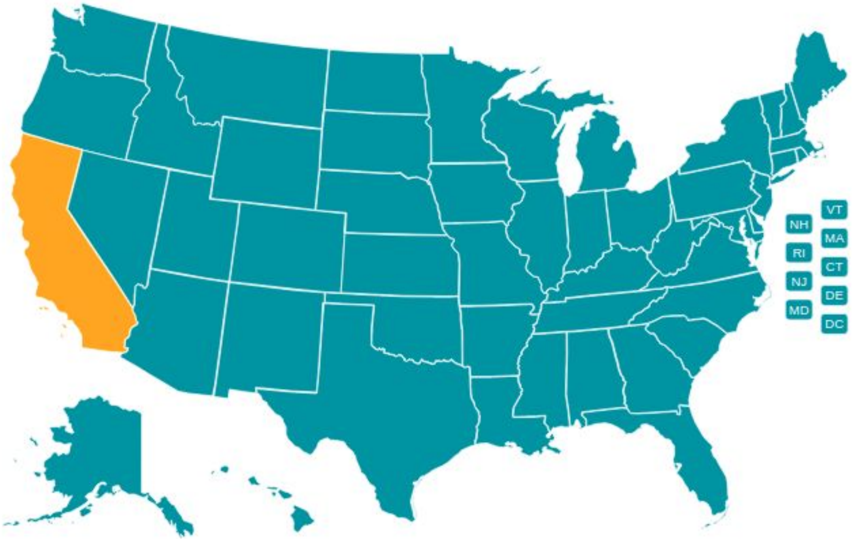

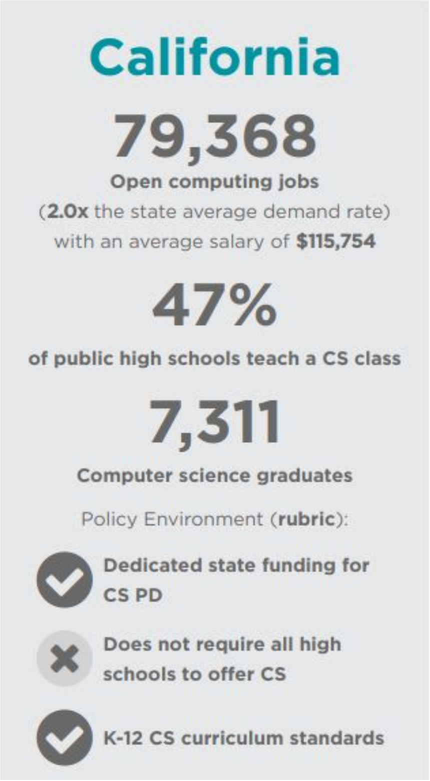

California

79,368

Open computing jobs

(**2.0x** the state average demand rate)
with an average salary of **$115,754**

47%

of public high schools teach a CS class

7,311

Computer science graduates

Policy Environment (**rubric**):

- ✓ **Dedicated state funding for CS PD**
- ✗ **Does not require all high schools to offer CS**
- ✓ **K-12 CS curriculum standards**

Code.org

No State Has Universal CS Education

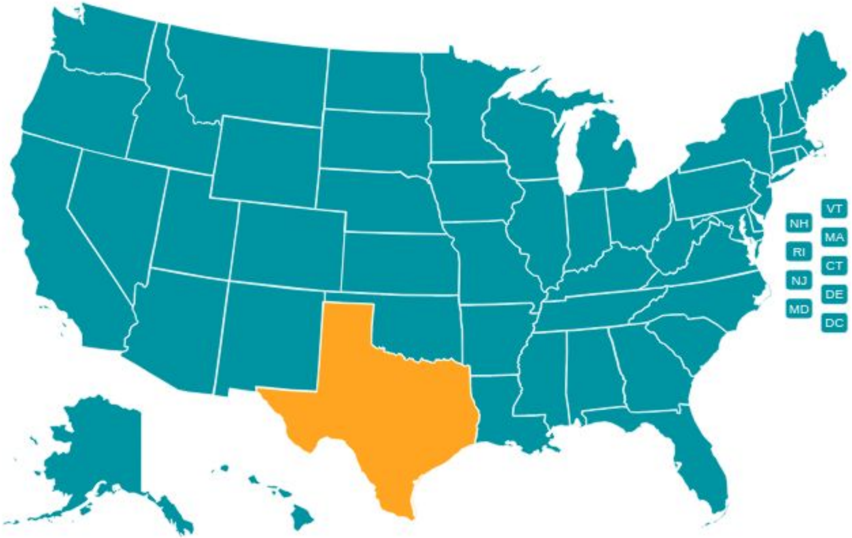

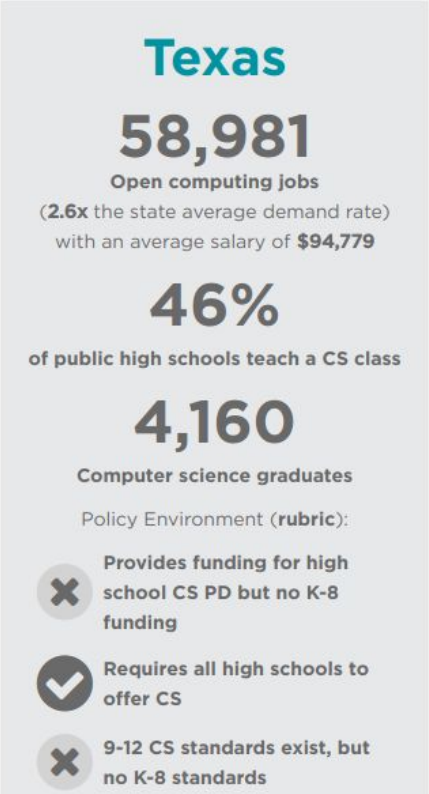

Texas

58,981

Open computing jobs

(**2.6x** the state average demand rate) with an average salary of **$94,779**

46%

of public high schools teach a CS class

4,160

Computer science graduates

Policy Environment (**rubric**):

- ✗ **Provides funding for high school CS PD but no K-8 funding**
- ✓ **Requires all high schools to offer CS**
- ✗ **9-12 CS standards exist, but no K-8 standards**

Code.org

No State Has Universal CS Education

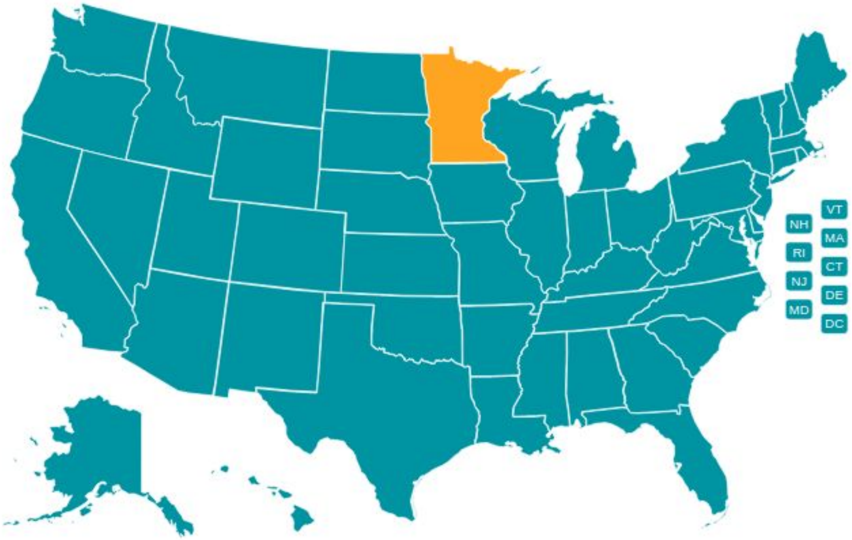

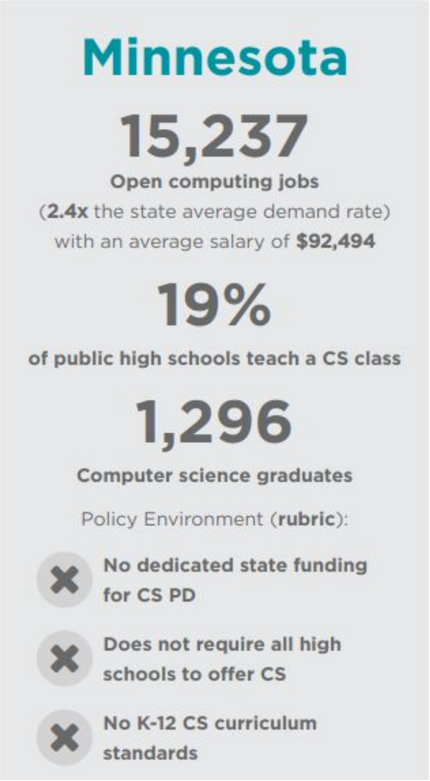

Minnesota

15,237

Open computing jobs

(**2.4x** the state average demand rate)
with an average salary of **$92,494**

19%

of public high schools teach a CS class

1,296

Computer science graduates

Policy Environment (**rubric**):

- ✗ **No dedicated state funding for CS PD**
- ✗ **Does not require all high schools to offer CS**
- ✗ **No K-12 CS curriculum standards**

Code.org

Industrial Revolutions (Grossly Oversimplified)

1. Mechanical power
    - Automated manufacturing
    - Self-powered vehicles (trains, steamboats)

2. Electrical power
    - Electric lighting; telegraph, telephone, radio; electromechanical devices

3. **Computer power**
    - Digital information processing; computer networking; Internet and World Wide Web

4. AI power
    - Computer perception; autonomous robots; automated decision making
    - Machine learning on massive datasets

CS Is Hard Enough. Why Should We Teach AI in K-12?

- AI is the new electricity.

- **Our children are growing up with AI.** By time many children arrive in kindergarten, they’ve spent two years conversing with Alexa.

- We must prepare for the next round of revolutionary disruption:
    - Autonomous robots everywhere.
    - Changing nature of work.
    - Demand for an AI-literate workforce.
    - AI policy issues regarding fairness, privacy/surveillance, disparate impacts of technology, etc.

The AI4K12 Initiative, a joint project of:

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

AI4K12 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 educators, researchers, and resource developers

Five Big Ideas in AI

1. **Perception:** Computers perceive the world using sensors.
2. **Representation and reasoning:** Agents maintain representations of the world and use them for reasoning.
3. **Learning:** Computers can learn from data.
4. **Natural interaction:** Intelligent agents require many kinds of information to interact naturally with humans.
5. **Societal impact:** AI can impact society in both positive and negative ways.

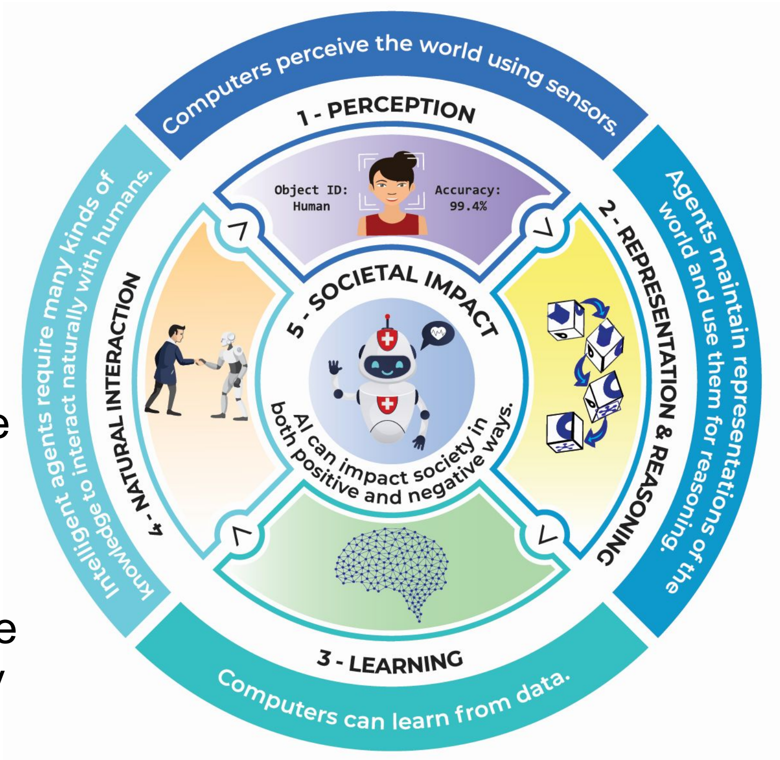

AI4K12.org

Five Big Ideas in Artificial Intelligence

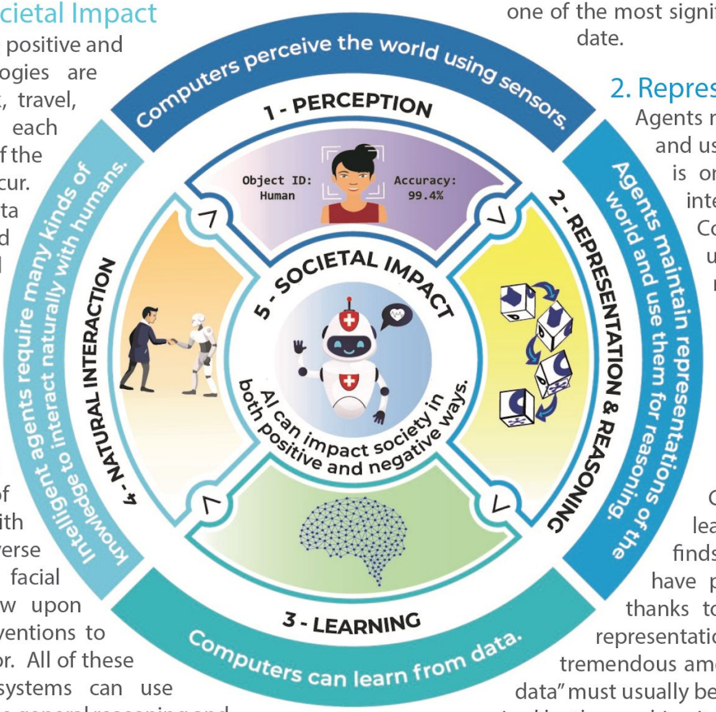

1. Perception

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 Five Big Ideas

- Now being referenced by multiple curriculum developers in the US and elsewhere.
- Big ideas poster is available in 16 languages.

Turkish

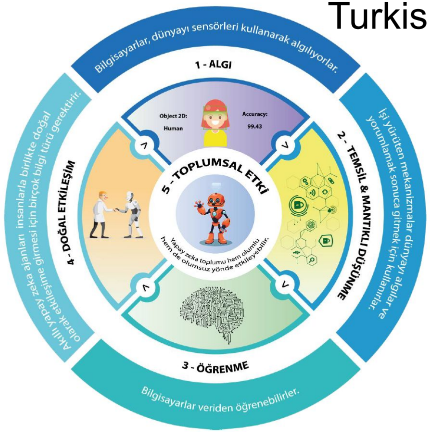

Chinese

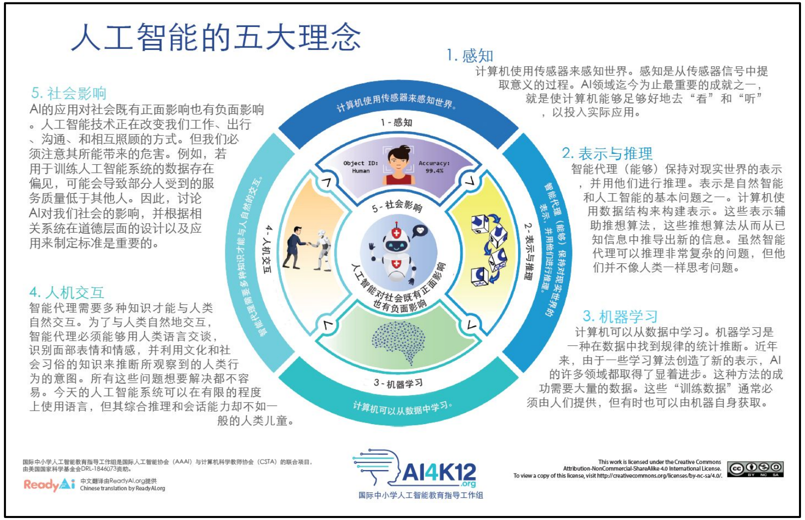

人工智能的五大理念

5. 社会影响

AI的应用对社会既有正面影响也有负面影响。人工智能技术正在改变我们工作、出行、沟通、和相互照顾的方式。但我们必须注意其所能带来的危害。例如，若用于训练人工智能系统的数据存在偏见，可能会导致部分人受到的服务质量低于其他人。因此，讨论AI对我们社会的影响，并根据相关系统在道德层面的设计以及应用来制定标准是重要的。

4. 人机交互

智能代理需要多种知识才能与人类自然交互。为了与人类自然地交互，智能代理必须能够用人类语言交谈，识别面部表情和情感，并利用文化和社会习俗的知识来推断所观察到的人类行为的意图。所有这些问题想要解决都不容易。今天的人工智能系统可以在有限的程度上使用语言，但其综合推理和会话能力却不如一般的人类儿童。

1. 感知

计算机使用传感器来感知世界。感知是从传感器信号中提取意义的过程。AI领域迄今为止最重要的成就之一，就是使计算机能够足够好地去“看”和“听”，以投入实际应用。

2. 表示与推理

智能代理（能够）保持对现实世界的表示，并用他们进行推理。表示是自然智能和人工智能的基本问题之一。计算机使用数据结构来构建表示。这些表示辅助推想算法，这些推想算法从而从已知信息中推导出新的信息。虽然智能代理可以推理非常复杂的问题，但他们并不像人类一样思考问题。

3. 机器学习

计算机可以从数据中学习。机器学习是一种在数据中找到规律的统计推断。近年来，由于一些学习算法创造了新的表示，AI的许多领域都取得了显着进步。这种方法的成功需要大量的数据。这些“训练数据”通常必须由人们提供，但有时也可以由机器自身获取。

国际中小学人工智能教育指导工作组是国际人工智能协会（AAAI）与计算机科学教师协会（CSTA）的联合项目，由美国国家科学基金会DRL-1846073资助。

ReadyAi 中文翻译由ReadyAI.org提供
Chinese translation by ReadyAI.org

AI4K12.org
国际中小学人工智能教育指导工作组

Korean

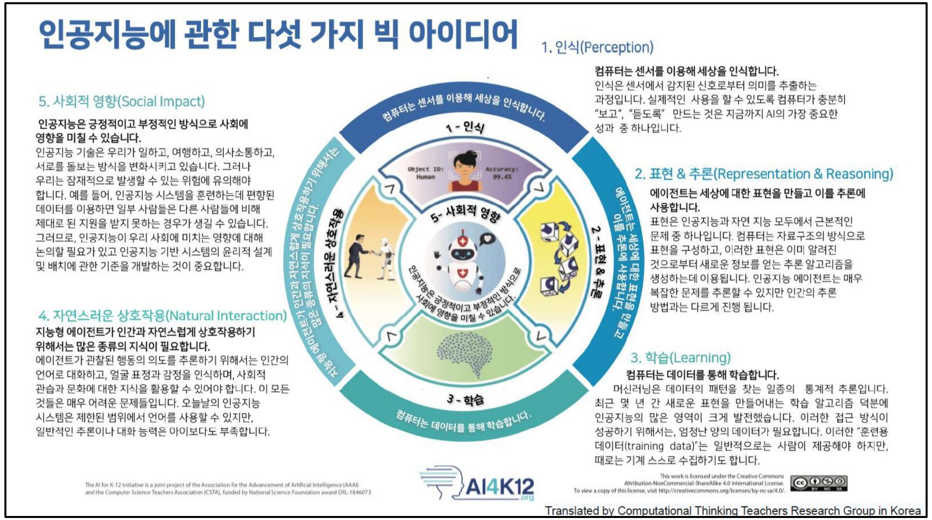

인공지능에 관한 다섯 가지 빅 아이디어

5. 사회적 영향(Social Impact)

인공지능은 긍정적이고 부정적인 방식으로 사회에 영향을 미칠 수 있습니다.
인공지능 기술은 우리가 일하고, 여행하고, 의사소통하고, 서로를 돌보는 방식을 변화시키고 있습니다. 그러나 우리는 잠재적으로 발생할 수 있는 위험에 유의해야 합니다. 예를 들어, 인공지능 시스템을 훈련하는데 편향된 데이터를 이용하면 일부 사람들은 다른 사람들에 비해 제대로 된 지원을 받지 못하는 경우가 생길 수 있습니다. 그러므로, 인공지능이 우리 사회에 미치는 영향에 대해 논의할 필요가 있고 인공지능 기반 시스템의 윤리적 설계 및 배치에 관한 기준을 개발하는 것이 중요합니다.

4. 자연스러운 상호작용(Natural Interaction)

지능형 에이전트가 인간과 자연스럽게 상호작용하기 위해서는 많은 종류의 지식이 필요합니다.
에이전트가 관찰된 행동의 의도를 추론하기 위해서는 인간의 언어로 대화하고, 얼굴 표정과 감정을 인식하며, 사회적 관습과 문화에 대한 지식을 활용할 수 있어야 합니다. 이 모든 것들은 매우 어려운 문제들입니다. 오늘날의 인공지능 시스템은 제한된 범위에서 언어를 사용할 수 있지만, 일반적인 추론이나 대화 능력은 아이보다도 부족합니다.

1. 인식(Perception)

컴퓨터는 센서를 이용해 세상을 인식합니다.
인식은 센서에서 감지된 신호로부터 의미를 추출하는 과정입니다. 실제적인 사용을 할 수 있도록 컴퓨터가 충분히 “보고“, “듣도록” 만드는 것은 지금까지 AI의 가장 중요한 성과 중 하나입니다.

2. 표현 & 추론(Representation & Reasoning)

에이전트는 세상에 대한 표현을 만들고 이를 추론에 사용합니다.
표현은 인공지능과 자연 지능 모두에서 근본적인 문제 중 하나입니다. 컴퓨터는 자료구조의 방식으로 표현을 구성하고, 이러한 표현은 이미 알려진 것으로부터 새로운 정보를 얻는 추론 알고리즘을 생성하는데 이용됩니다. 인공지능 에이전트는 매우 복잡한 문제를 추론할 수 있지만 인간의 추론 방법과는 다르게 진행 됩니다.

3. 학습(Learning)

컴퓨터는 데이터를 통해 학습합니다.
머신러닝은 데이터의 패턴을 찾는 일종의 통계적 추론입니다. 최근 몇 년 간 새로운 표현을 만들어내는 학습 알고리즘 덕분에 인공지능의 많은 영역이 크게 발전했습니다. 이러한 접근 방식이 성공하기 위해서는, 엄청난 양의 데이터가 필요합니다. 이러한 “훈련용 데이터(training data)”는 일반적으로는 사람이 제공해야 하지만, 때로는 기계 스스로 수집하기도 합니다.

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

Translated by Computational Thinking Teachers Research Group in Korea

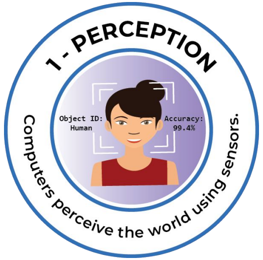

Page 1 (of 4) of the draft guidelines for Big Idea #1: Perception

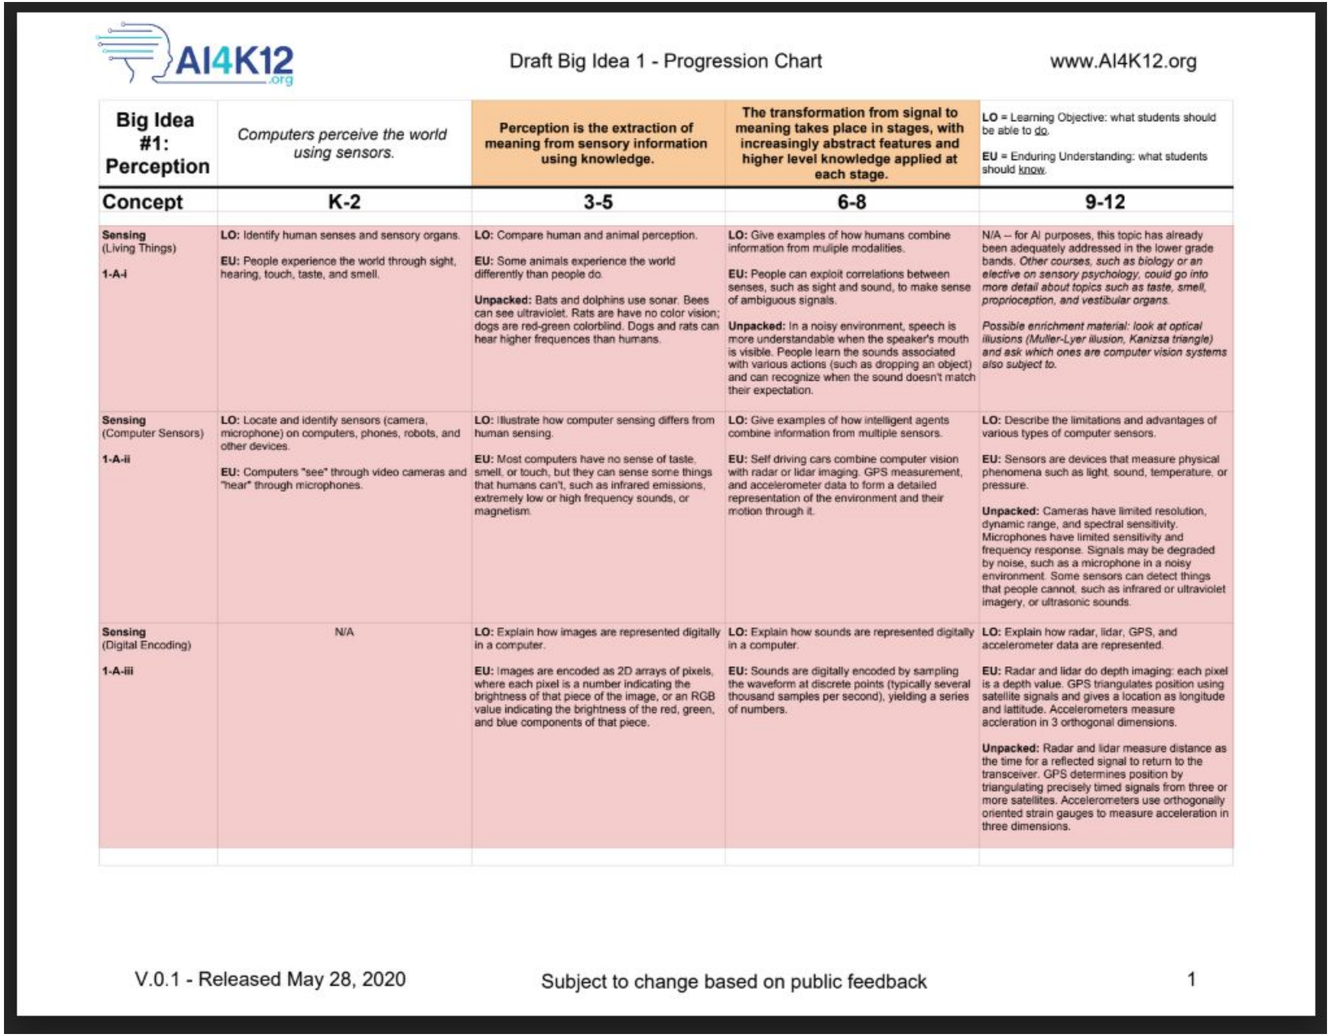

AI4K12.org

Draft Big Idea 1 - Progression Chart

www.AI4K12.org

Big Idea #1: Perception	*Computers perceive the world using sensors.*	**Perception is the extraction of meaning from sensory information using knowledge.**	**The transformation from signal to meaning takes place in stages, with increasingly abstract features and higher level knowledge applied at each stage.**	**LO** = Learning Objective: what students should be able to do. **EU** = Enduring Understanding: what students should know.
Concept	**K-2**	**3-5**	**6-8**	**9-12**
Sensing (Living Things) **1-A-i**	**LO:** Identify human senses and sensory organs. **EU:** People experience the world through sight, hearing, touch, taste, and smell.	**LO:** Compare human and animal perception. **EU:** Some animals experience the world differently than people do. **Unpacked:** Bats and dolphins use sonar. Bees can see ultraviolet. Rats are have no color vision; dogs are red-green colorblind. Dogs and rats can hear higher frequences than humans.	**LO:** Give examples of how humans combine information from muliple modalities. **EU:** People can exploit correlations between senses, such as sight and sound, to make sense of ambiguous signals. **Unpacked:** In a noisy environment, speech is more understandable when the speaker's mouth is visible. People learn the sounds associated with various actions (such as dropping an object) and can recognize when the sound doesn't match their expectation.	N/A – for AI purposes, this topic has already been adequately addressed in the lower grade bands. *Other courses, such as biology or an elective on sensory psychology, could go into more detail about topics such as taste, smell, proprioception, and vestibular organs.* *Possible enrichment material: look at optical illusions (Muller-Lyer illusion, Kanizsa triangle) and ask which ones are computer vision systems also subject to.*
Sensing (Computer Sensors) **1-A-ii**	**LO:** Locate and identify sensors (camera, microphone) on computers, phones, robots, and other devices. **EU:** Computers "see" through video cameras and "hear" through microphones.	**LO:** Illustrate how computer sensing differs from human sensing. **EU:** Most computers have no sense of taste, smell, or touch, but they can sense some things that humans can't, such as infrared emissions, extremely low or high frequency sounds, or magnetism.	**LO:** Give examples of how intelligent agents combine information from multiple sensors. **EU:** Self driving cars combine computer vision with radar or lidar imaging. GPS measurement, and accelerometer data to form a detailed representation of the environment and their motion through it.	**LO:** Describe the limitations and advantages of various types of computer sensors. **EU:** Sensors are devices that measure physical phenomena such as light, sound, temperature, or pressure. **Unpacked:** Cameras have limited resolution, dynamic range, and spectral sensitivity. Microphones have limited sensitivity and frequency response. Signals may be degraded by noise, such as a microphone in a noisy environment. Some sensors can detect things that people cannot, such as infrared or ultraviolet imagery, or ultrasonic sounds.
Sensing (Digital Encoding) **1-A-iii**	N/A	**LO:** Explain how images are represented digitally in a computer. **EU:** Images are encoded as 2D arrays of pixels, where each pixel is a number indicating the brightness of that piece of the image, or an RGB value indicating the brightness of the red, green, and blue components of that piece.	**LO:** Explain how sounds are represented digitally in a computer. **EU:** Sounds are digitally encoded by sampling the waveform at discrete points (typically several thousand samples per second), yielding a series of numbers.	**LO:** Explain how radar, lidar, GPS, and accelerometer data are represented. **EU:** Radar and lidar do depth imaging: each pixel is a depth value. GPS triangulates position using satellite signals and gives a location as longitude and lattitude. Accelerometers measure accleration in 3 orthogonal dimensions. **Unpacked:** Radar and lidar measure distance as the time for a reflected signal to return to the transceiver. GPS determines position by triangulating precisely timed signals from three or more satellites. Accelerometers use orthogonally oriented strain gauges to measure acceleration in three dimensions.

V.0.1 - Released May 28, 2020

Subject to change based on public feedback

1

Big Idea #1: Perception

Computers perceive the world using sensors.

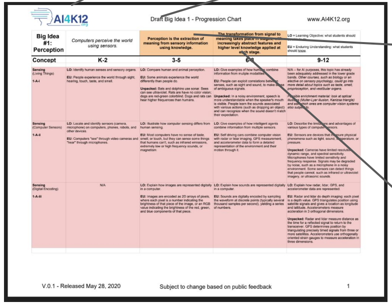

AI4K12 Draft Big Idea 1 - Progression Chart www.AI4K12.org

Big Idea #1: Perception	*Computers perceive the world using sensors.*	**Perception is the extraction of meaning from sensory information using knowledge.**	**The transformation from signal to meaning takes place in stages, with increasingly abstract features and higher level knowledge applied at each stage.**	**LO** = Learning Objective: what students should be able to do. **EU** = Enduring Understanding: what students should know.
Concept	**K-2**	**3-5**	**6-8**	**9-12**
Sensing (Living Things) **1-A-i**	**LO:** Identify human senses and sensory organs. **EU:** People experience the world through sight, hearing, touch, taste, and smell.	**LO:** Compare human and animal perception. **EU:** Some animals experience the world differently than people do. **Unpacked:** Bats and dolphins use sonar. Bees can see ultraviolet. Rats are have no color vision; dogs are red-green colorblind. Dogs and rats can hear higher frequences than humans.	**LO:** Give examples of how humans combine information from muliple modalities. **EU:** People can exploit correlations between senses, such as sight and sound, to make sense of ambiguous signals. **Unpacked:** In a noisy environment, speech is more understandable when the speaker's mouth is visible. People learn the sounds associated with various actions (such as dropping an object) and can recognize when the sound doesn't match their expectation.	N/A -- for AI purposes, this topic has already been adequately addressed in the lower grade bands. *Other courses, such as biology or an elective on sensory psychology, could go into more detail about topics such as taste, smell, proprioception, and vestibular organs.* *Possible enrichment material: look at optical illusions (Muller-Lyer illusion, Kanizsa triangle) and ask which ones are computer vision systems also subject to.*
Sensing (Computer Sensors) **1-A-ii**	**LO:** Locate and identify sensors (camera, microphone) on computers, phones, robots, and other devices. **EU:** Computers "see" through video cameras and "hear" through microphones.	**LO:** Illustrate how computer sensing differs from human sensing. **EU:** Most computers have no sense of taste, smell, or touch, but they can sense some things that humans can't, such as infrared emissions, extremely low or high frequency sounds, or magnetism.	**LO:** Give examples of how intelligent agents combine information from multiple sensors. **EU:** Self driving cars combine computer vision with radar or lidar imaging, GPS measurement, and accelerometer data to form a detailed representation of the environment and their motion through it.	**LO:** Describe the limitations and advantages of various types of computer sensors. **EU:** Sensors are devices that measure physical phenomena such as light, sound, temperature, or pressure. **Unpacked:** Cameras have limited resolution, dynamic range, and spectral sensitivity. Microphones have limited sensitivity and frequency response. Signals may be degraded by noise, such as a microphone in a noisy environment. Some sensors can detect things that people cannot, such as infrared or ultraviolet imagery, or ultrasonic sounds.
Sensing (Digital Encoding) **1-A-iii**	N/A	**LO:** Explain how images are represented digitally in a computer. **EU:** Images are encoded as 2D arrays of pixels, where each pixel is a number indicating the brightness of that piece of the image, or an RGB value indicating the brightness of the red, green, and blue components of that piece.	**LO:** Explain how sounds are represented digitally in a computer. **EU:** Sounds are digitally encoded by sampling the waveform at discrete points (typically several thousand samples per second), yielding a series of numbers.	**LO:** Explain how radar, lidar, GPS, and accelerometer data are represented. **EU:** Radar and lidar do depth imaging: each pixel is a depth value. GPS triangulates position using satellite signals and gives a location as longitude and latitude. Accelerometers measure accleration in 3 orthogonal dimensions. **Unpacked:** Radar and lidar measure distance as the time for a reflected signal to return to the transceiver. GPS determines position by triangulating precisely timed signals from three or more satellites. Accelerometers use orthogonally oriented strain gauges to measure acceleration in three dimensions.

V.0.1 - Released May 28, 2020 Subject to change based on public feedback 1

Perception is the extraction of meaning from sensory information using knowledge.

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

What Does AI Thinking Look like in K-12?

Computational Thinking

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

AI Thinking

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

Visual Perception

Computers can see:

- Faces
- Household objects
- Road scenes

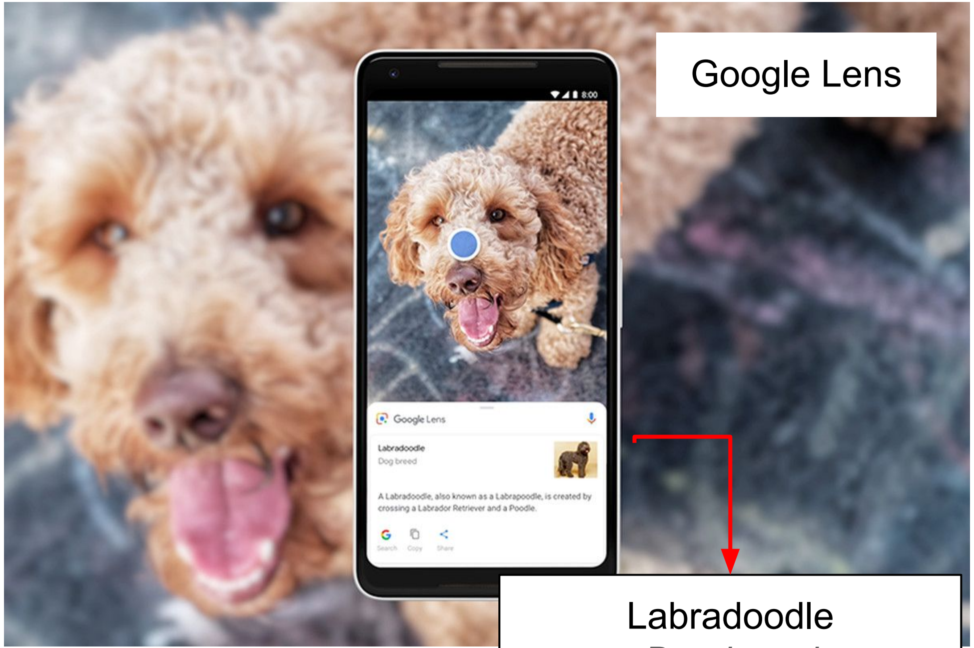

Google Lens

Labradoodle
Dog breed

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

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

Pixels

Textures

Edges

Shadows

Contours

Boundaries

Surfaces

Parts

Objects

Scenes

Neural Net Edge and Face Detection Demo

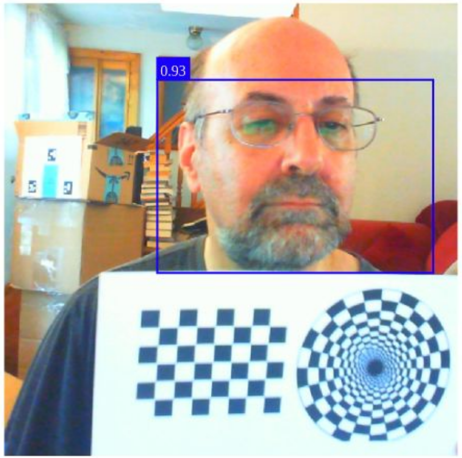

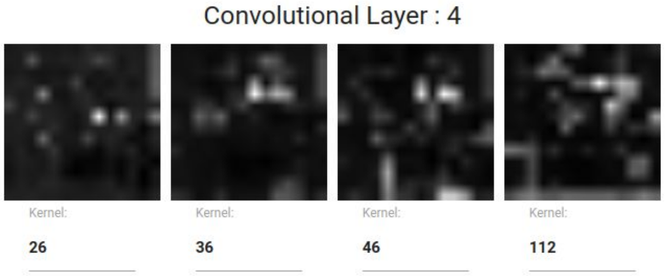

Convolutional Layer : 4

Kernel: 26

Kernel: 36

Kernel: 46

Kernel: 112

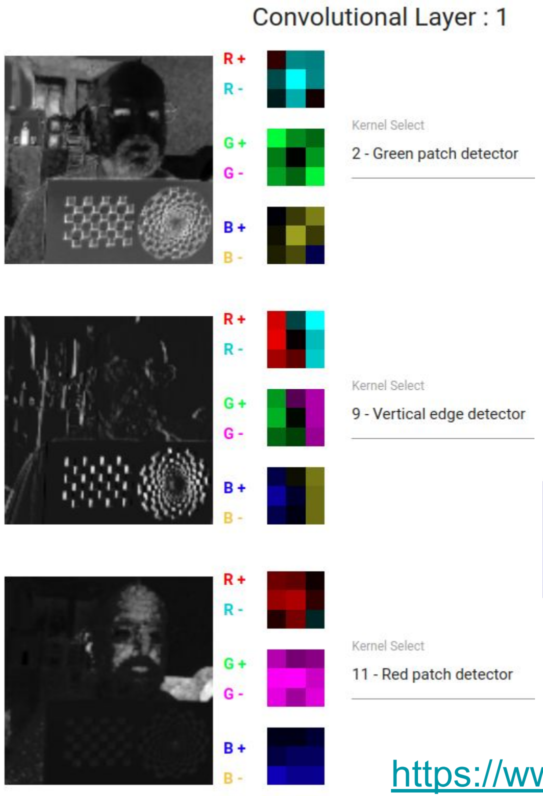

Convolutional Layer : 1

Kernel Select: 2 - Green patch detector

Kernel Select: 9 - Vertical edge detector

Kernel Select: 11 - Red patch detector

Real-time face detection by a deep neural network (TinyYoloV2)

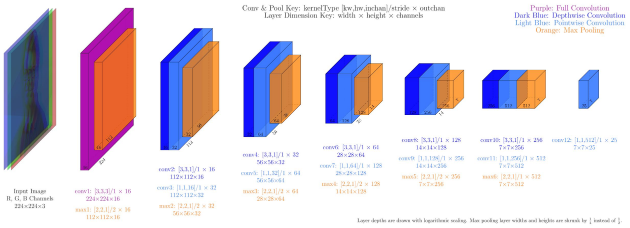

https://www.cs.cmu.edu/~dst/FaceDemo

Speech Perception

Computers can understand spoken language.

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

- "<u>They're</u> building <u>their</u> new house over <u>there</u>."

I can make artifacts that understand voice commands.

Levels of representation and linguistic knowledge

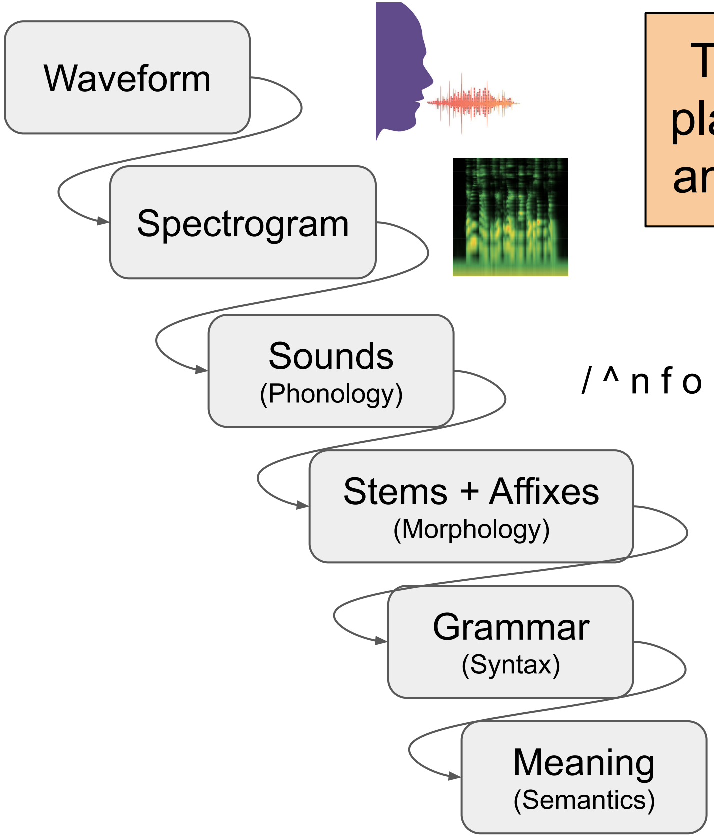

The transformation from signal to meaning takes place in stages, with increasingly abstract features and higher level knowledge applied at each stage.

/^nforgɜtabʊl#saIt/

un + forget + able   site/sight

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

“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.

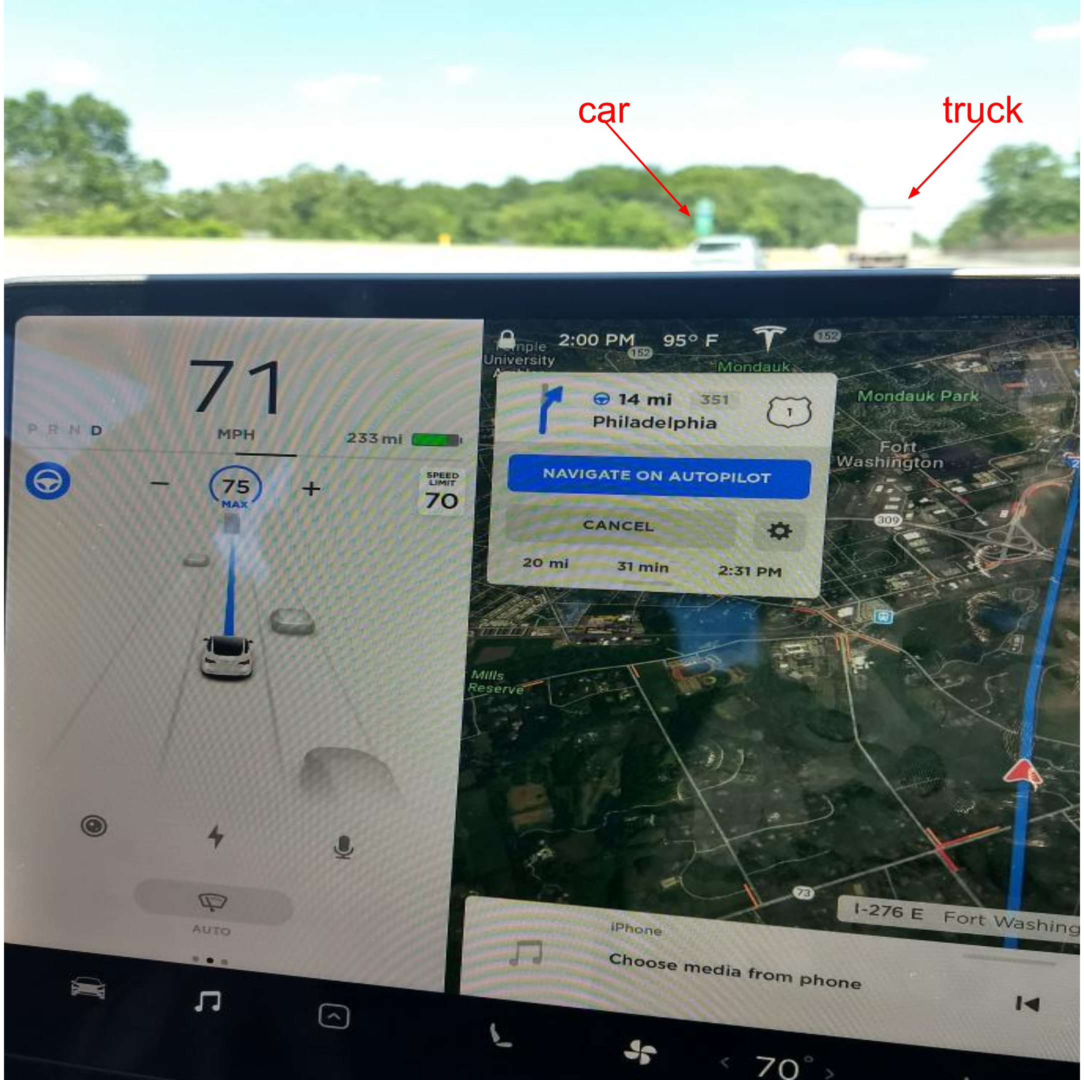

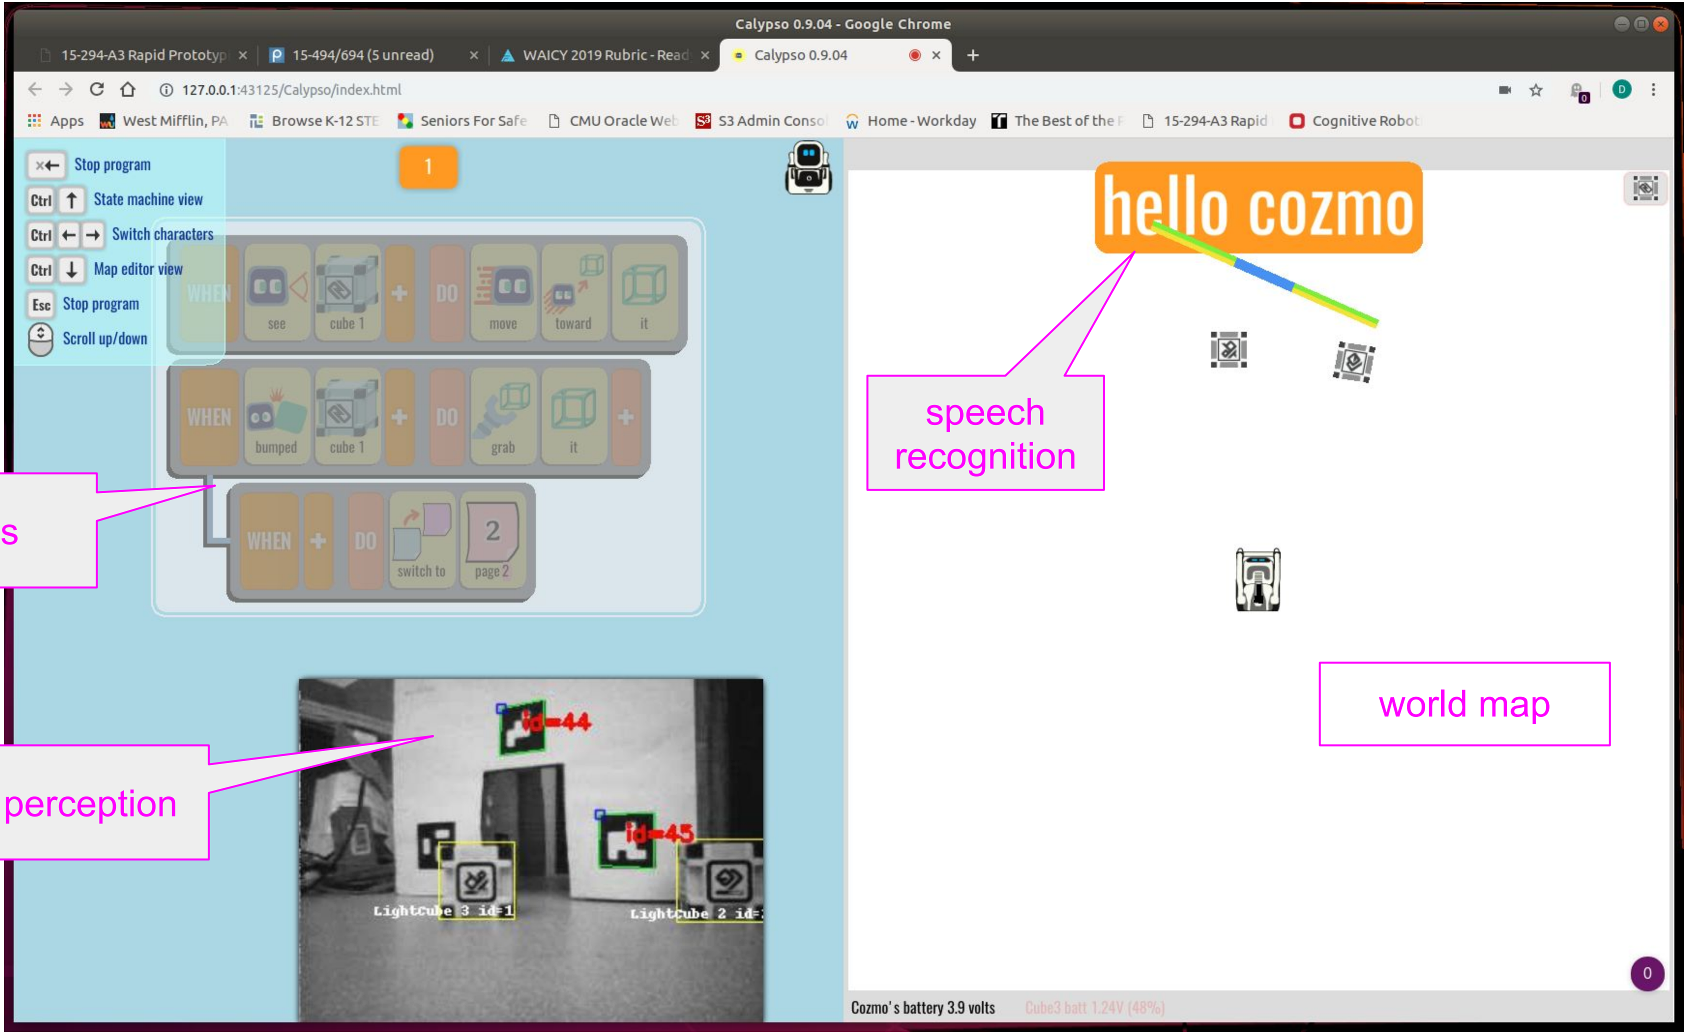

Cozmo's battery 3.9 volts

Cube3 batt 1.24V (48%)

rules

speech recognition

perception

world map

Calypso for Cozmo

- 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)

Reasoning

Types of reasoning problems:

- Classification: cat or dog?
- Search: find a path to a goal state.
- Many other types, including regression, optimization, sequential decision making, logical deduction, Bayesian inference, etc.

I can build a classifier.

I can build a reasoner.

Reasoning Algorithms

There are algorithms for each type of reasoning problem.

- Classifiers
  - Decision trees
  - Neural networks
  - Nearest neighbor

- Search algorithms
  - Breadth-first, depth-first, best-first, heuristic search, etc.

Learning: Computers Can Learn From Data

Computers don’t learn the way people do.

Machine learning constructs 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.

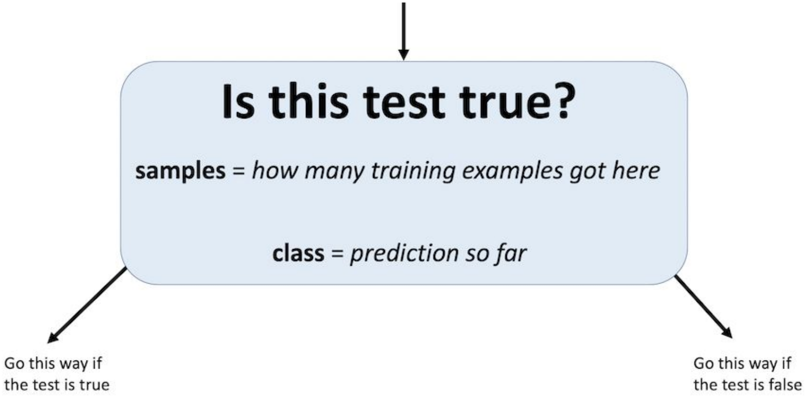

Is this test true?

samples = *how many training examples got here*

class = *prediction so far*

Go this way if the test is true

Go this way if the test is false

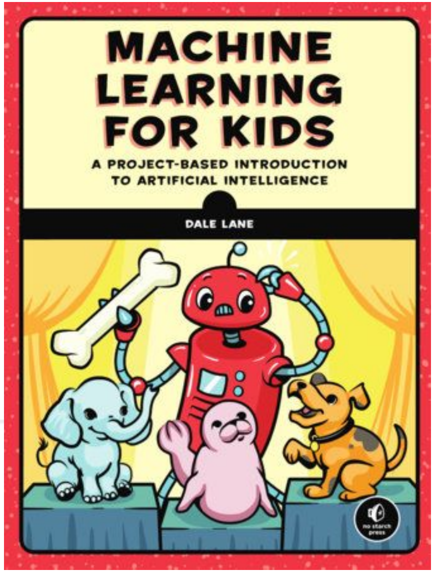

MACHINE LEARNING FOR KIDS

A PROJECT-BASED INTRODUCTION TO ARTIFICIAL INTELLIGENCE

DALE LANE

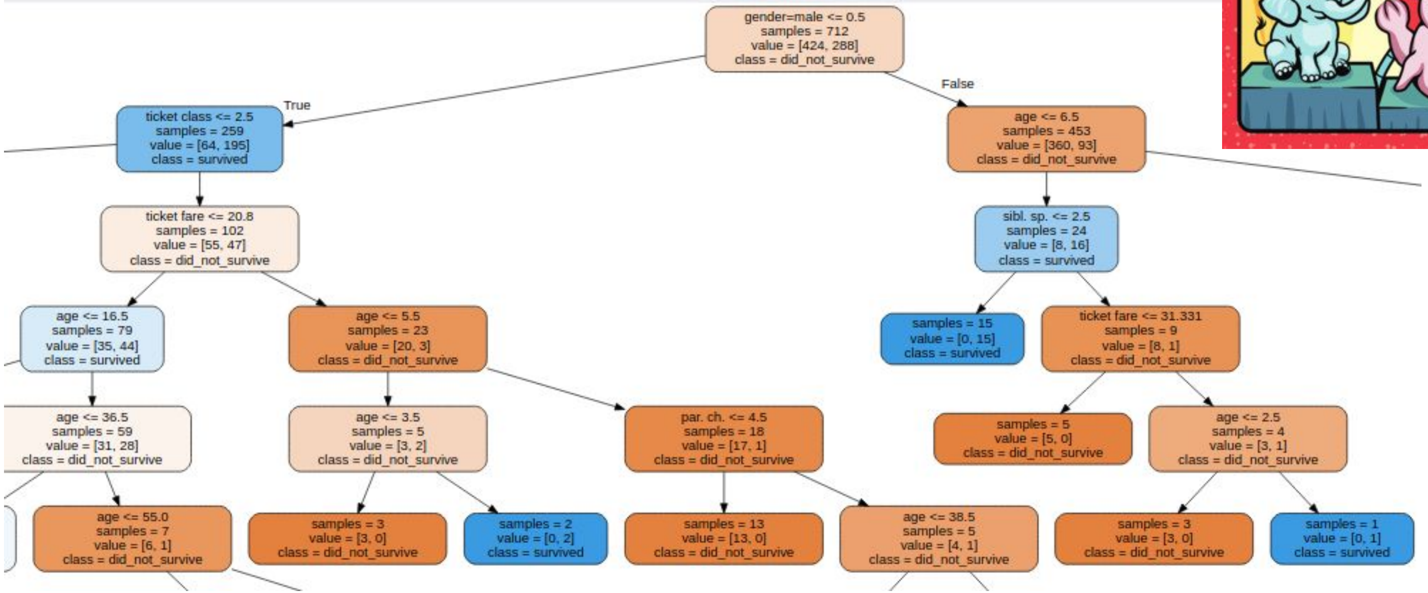

The Importance of Training Data

- Goal: generalize correctly to new instances
- The dataset needs to be representative
- Effects of biased training data

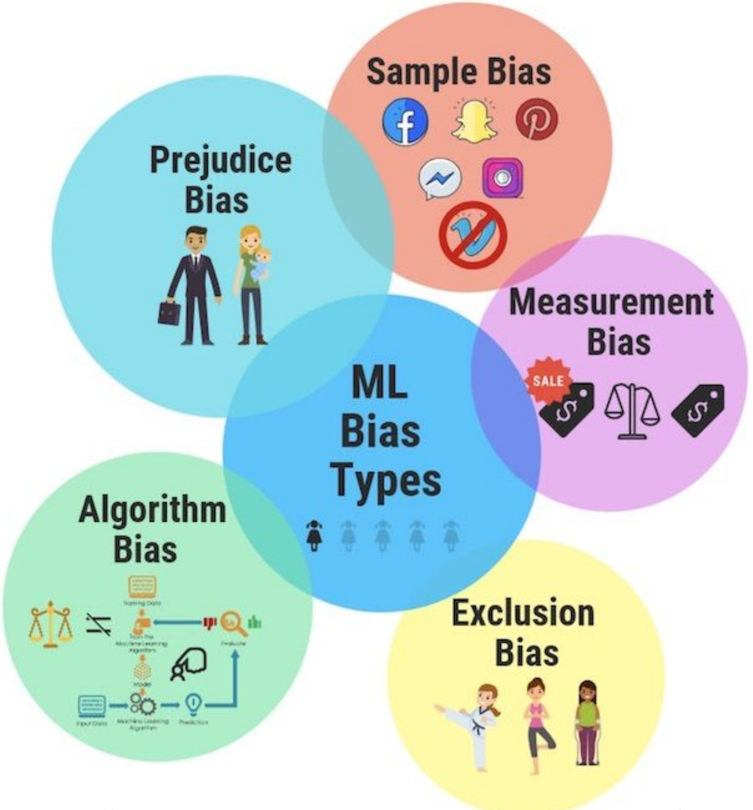

https://www.datasciencecentral.com/profiles/blogs/three-steps-to-addressing-bias-in-machine-learning

Language Understanding

- Question answering: “How much does an alligator weigh?”
- Machine translation
- Chatbots and intelligent agents
    - Intent recognition
    - Slot filling

I can build a simple chatbot.

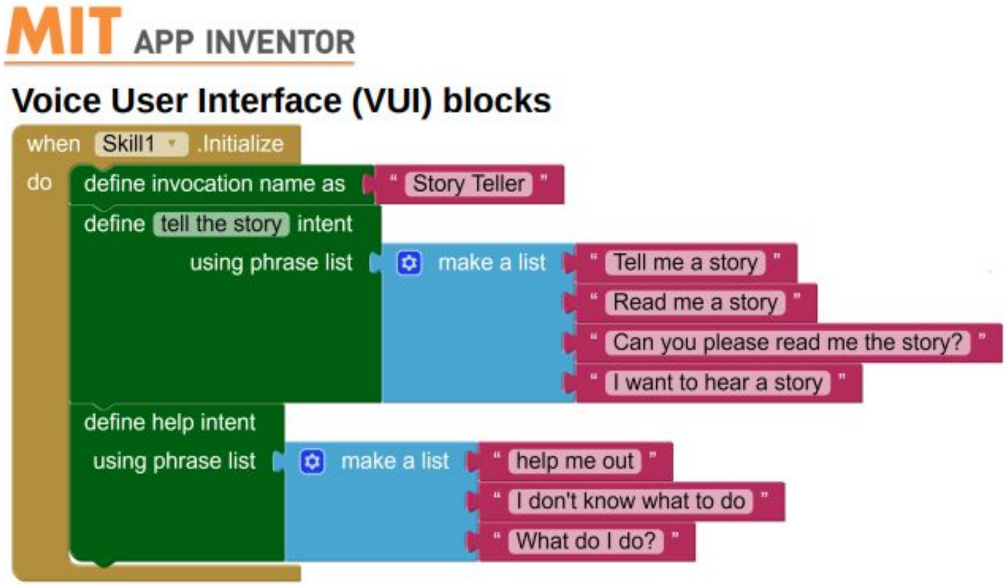

Voice User Interface (VUI) blocks

```
when Skill1 .Initialize
do  define invocation name as " Story Teller "
    define tell the story intent
           using phrase list  make a list " Tell me a story "
                                          " Read me a story "
                                          " Can you please read me the story? "
                                          " I want to hear a story "
    define help intent
     using phrase list  make a list " help me out "
                                    " I don't know what to do "
                                    " What do I do? "
```

Work by Jessica Van Brummelen and Hal Abelson

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.

Moving Forward

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

David Touretzky, CMU & Christina Gardner-McCune, UF

Funded by National Science Foundation award DRL-1846073.

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 State & Territories Completed Plans

NM, VA
Puerto Rico, Virgin Islands

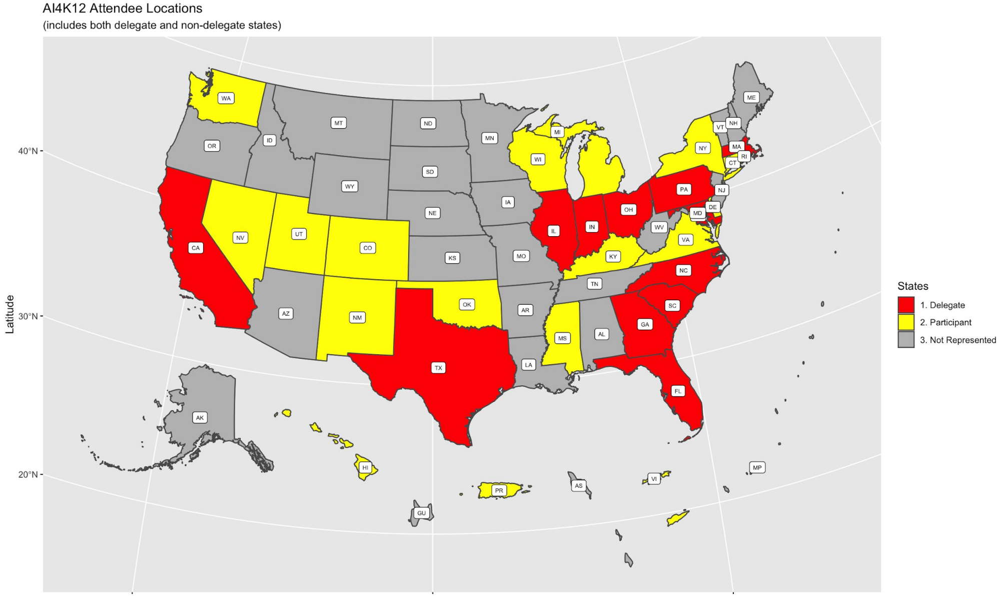

AI4K12 Attendee Locations
(includes both delegate and non-delegate states)

- 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 Us in Developing the Guidelines, or Help Grow the Community of AI Resource Developers

Visit us:

https://AI4K12.org

Join the mailing list:

https://aaai.org/Organization/mailing-lists.php

Questions?