READYAI

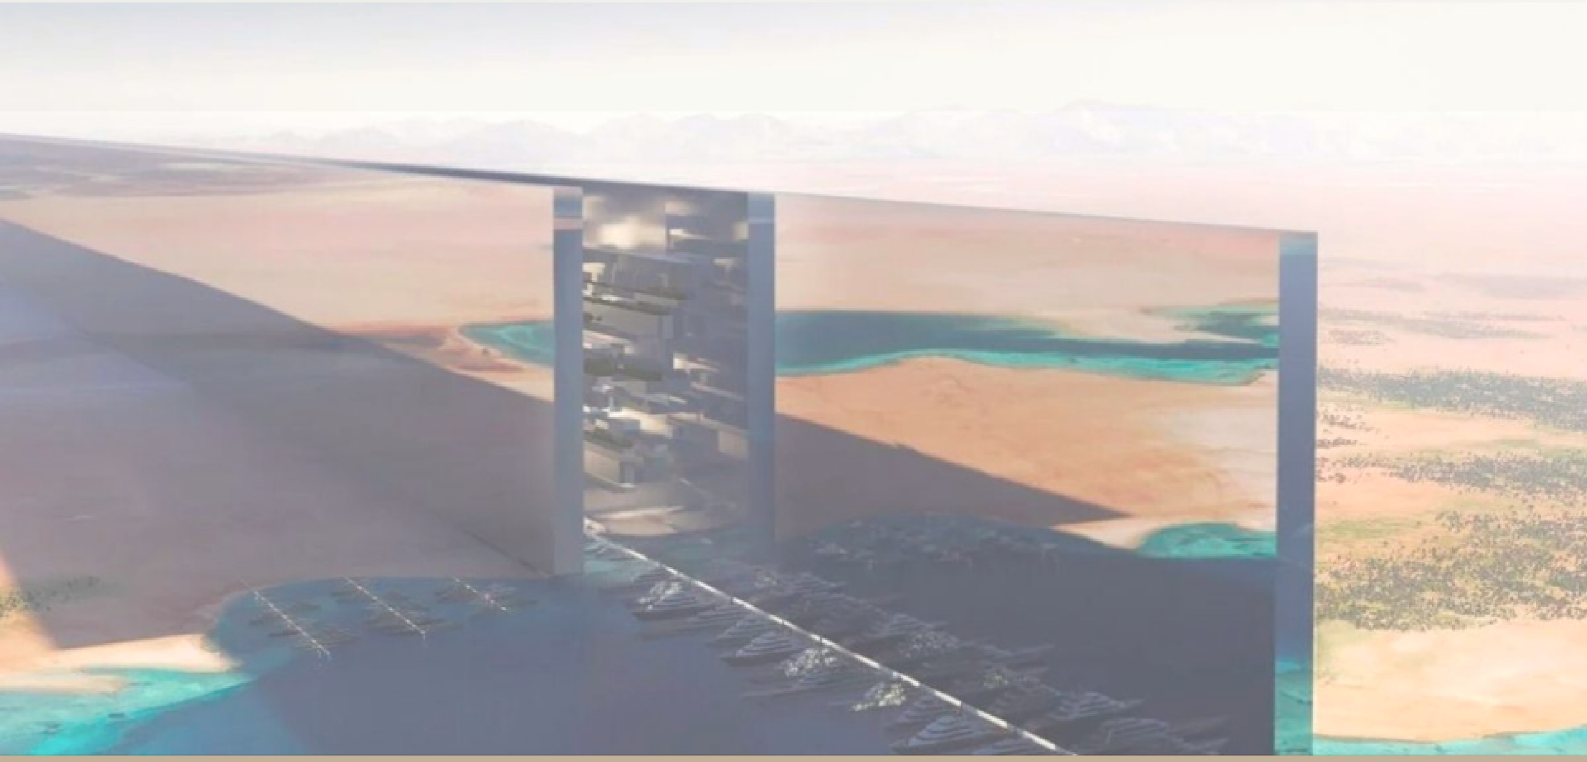

2022-2023 NEOM AI REPORT

INTRODUCTION

ReadyAI and Neom Community School's renewed partnership for the Academic Year of 2022-2023 produced a bespoke, standalone CS curriculum infused with first-of-its-kind learnings in artificial intelligence. This work was aligned to the provided units of inquiry and fully integrated interdisciplinary assessment strategies to reinforce IB practices and dispositions. The curriculum aimed to expand the hours of direct instruction while deepening engagement with artificial intelligence and computer science concepts. Though we have arrived in a place of pivoting to integrate CS and AI into core curricular strategies at NCS, ReadyAI stands firmly committed to the promises of co-design in order to create future ready learners in our ongoing collaborations. As suggested by the Crown Prince, Mohammed bin Salman, we exist in a world that has provided ample reason to …*"think and work hard to take advantage of artificial intelligence and unleash its full potential to advance our societies and economies."* Our work over the past year aimed to solidify this promise of AI in society as embodied in the teachers and students of NCS.

The report that follows offers insight into the documentation produced by ReadyAI as well as the ways in which our efforts dovetail with the goals and dreams of the Neom Project.

Rooz Aliabadi, Ph.D.
CEO ReadyAI

Jordan Mroziak, Ed.D.
CTO ReadyAI

IMPLEMENTATION AND MOVING FORWARD TOGETHER

In today's rapidly evolving technological landscape, artificial intelligence (AI) has emerged as a significant force shaping various aspects of our lives. As AI becomes increasingly prevalent, it is crucial to equip the younger generation with a foundational understanding of this transformative technology. Integrating AI education into K-12 curriculum through inquiry and project-based learning (PBL) approaches, as aligned with International Baccalaureate standards, offers incredible value by fostering critical thinking, problem-solving skills, and ethical awareness. The importance of developing an AI curriculum spanning across grade levels that utilizes inquiry and PBL has demonstrable benefits for their intellectual growth, career readiness, and responsible participation in future-facing society.

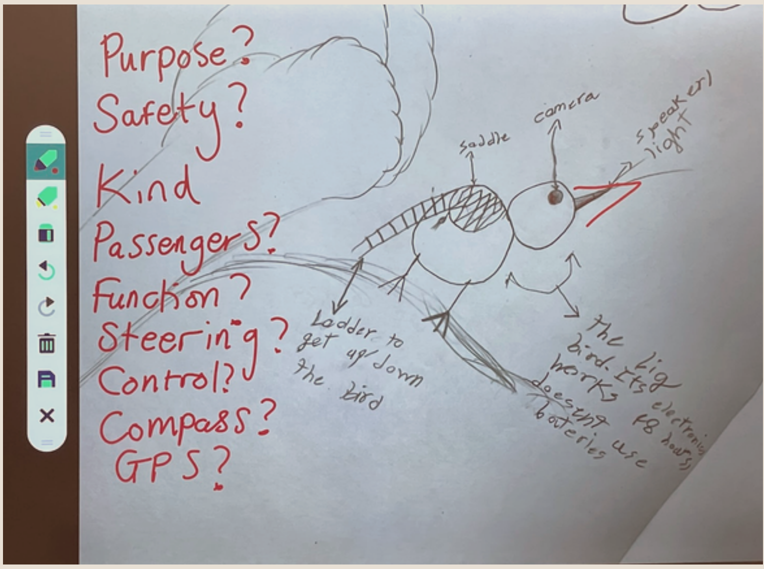

Implementing this comprehensive curriculum in artificial intelligence (AI) holds significant value for K-12 learners, as it provides structured and systematic learning opportunities in this rapidly advancing field. Our AI curriculum offers a range of benefits that go beyond isolated lessons, integrating AI education across subjects and grade levels. By introducing AI concepts, theories, and practical applications, a curriculum in AI ensures that students develop a deep and nuanced understanding of this transformative technology.

Our AI curriculum prepares NCS students for the future by equipping them with in-demand skills alongside the cognitive flexibility to adapt to the needs of the future. As AI becomes increasingly prevalent in society, the demand for individuals with expertise in AI-related fields continues to rise. By introducing AI concepts, programming languages, and problem-solving techniques, a curriculum in AI enables students to develop valuable skills that are highly sought after in the job market. These skills include critical thinking, data analysis, algorithmic reasoning, and computational thinking, which are essential for success in an AI-driven world.

Similarly, ReadyAI's curriculum promotes explorations beyond computer science (CS) learning by fostering connections between various subjects. AI draws from diverse fields such as mathematics, computer science, statistics, and psychology, making it an ideal platform for this exploration. By incorporating AI into the curriculum, NCS can create cross-curricular projects and activities that encourage students to apply AI principles in different contexts. This approach not only enhances students' knowledge and skills in AI but also strengthens their understanding of other subjects, promoting a holistic and interconnected understanding of the world. As mentioned previously, ReadyAI and NCS have decided to rededicate our ongoing partnership to the thoughtful construction of both existing and to-be-developed content into a truly interdisciplinary approach to AI and CS learning. This integrated approach will thoughtfully and critically enmesh AI dispositions and skills into core curricular explorations, further setting NCS apart as a truly unique and innovative school both in Saudi Arabia and across the world.

Integrating AI education through inquiry and PBL engages students in active learning experiences that promote critical thinking and problem-solving skills. Inquiry-based learning encourages learners to ask questions, investigate, and seek answers independently. By exploring AI concepts and applications, students develop an inquisitive mindset, honing their ability to analyze information and think critically with cutting edge technology as a means of engaging the world around them. They learn to evaluate the strengths, limitations, and ethical implications of AI technologies, enabling them to make informed decisions in the future.

Project-based learning complements inquiry-based learning by providing students with real-world problems to solve collaboratively. By engaging in AI-related projects, such as designing chatbots or developing machine learning models, students learn to apply theoretical knowledge in practical contexts for the good of their communities. This hands-on approach enhances their problem-solving skills, as they identify challenges, generate solutions, and iterate on their designs. Moreover, PBL cultivates creativity and innovation, allowing students to think outside the box and develop novel AI applications that address societal needs. This approach to teaching AI through inquiry and PBL equips K-12 learners with essential technological literacy and prepares them for the careers of the future. As AI technologies continue to shape industries, professions, and daily life, possessing an evolving understanding of AI becomes crucial for navigating the modern workforce. By engaging with AI concepts and tools, students develop digital literacy skills, allowing them to become discerning consumers and responsible creators of AI-driven solutions.

Inquiry-based learning exposes students to a wide range of AI applications, fostering awareness of how AI permeates various sectors, such as healthcare, finance, transportation, and entertainment. This exposure enables students to envision potential career paths in AI-related fields, encouraging them to explore STEM subjects and pursue advanced studies in computer science, data science, or AI research. By fostering interest and enthusiasm early on, educators can help bridge the existing skills gap and ensure a pipeline of qualified AI professionals in the future.

Teaching K-12 learners about artificial intelligence through inquiry and project-based learning has immense value in their intellectual growth, career readiness, and responsible participation in the digital era. By fostering critical thinking, problem-solving skills, and ethical awareness, AI education equips students with the tools they need to navigate an increasingly AI-driven world. Through inquiry-based learning, they develop curiosity and the ability to think critically, while project-based learning allows them to apply their knowledge in practical contexts. Moreover, AI education enhances technological literacy, preparing students for future careers, and fosters ethical awareness, empowering them to participate responsibly in shaping the future of AI. By embracing inquiry and PBL approaches, educators can inspire and empower the next generation to become informed and engaged citizens in an AI-driven society.

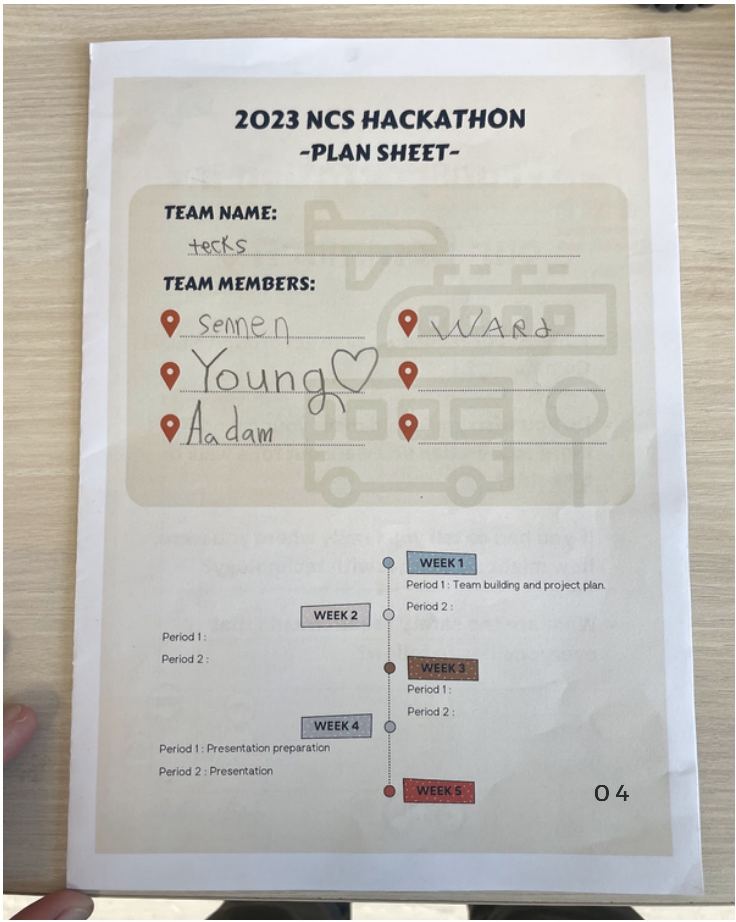

While the co-design and collaboration work of creating a thoroughly integrated AI curriculum will continue, the products created over the past year demonstrate cutting edge learning design in concepts such as facial recognition, neural nets, speech recognition, and computer vision.

In conclusion, implementing a comprehensive curriculum in artificial intelligence offers numerous advantages for K-12 learners. By promoting interdisciplinary learning, preparing students for the future, fostering innovation and creativity, and nurturing ethical awareness, our developing AI curriculum equips students with the knowledge, skills, and mindset necessary to thrive in an AI-driven world. As AI continues to reshape various aspects of society, a well-structured curriculum in AI ensures that students are not merely passive consumers of technology but active participants and contributors in shaping its future.

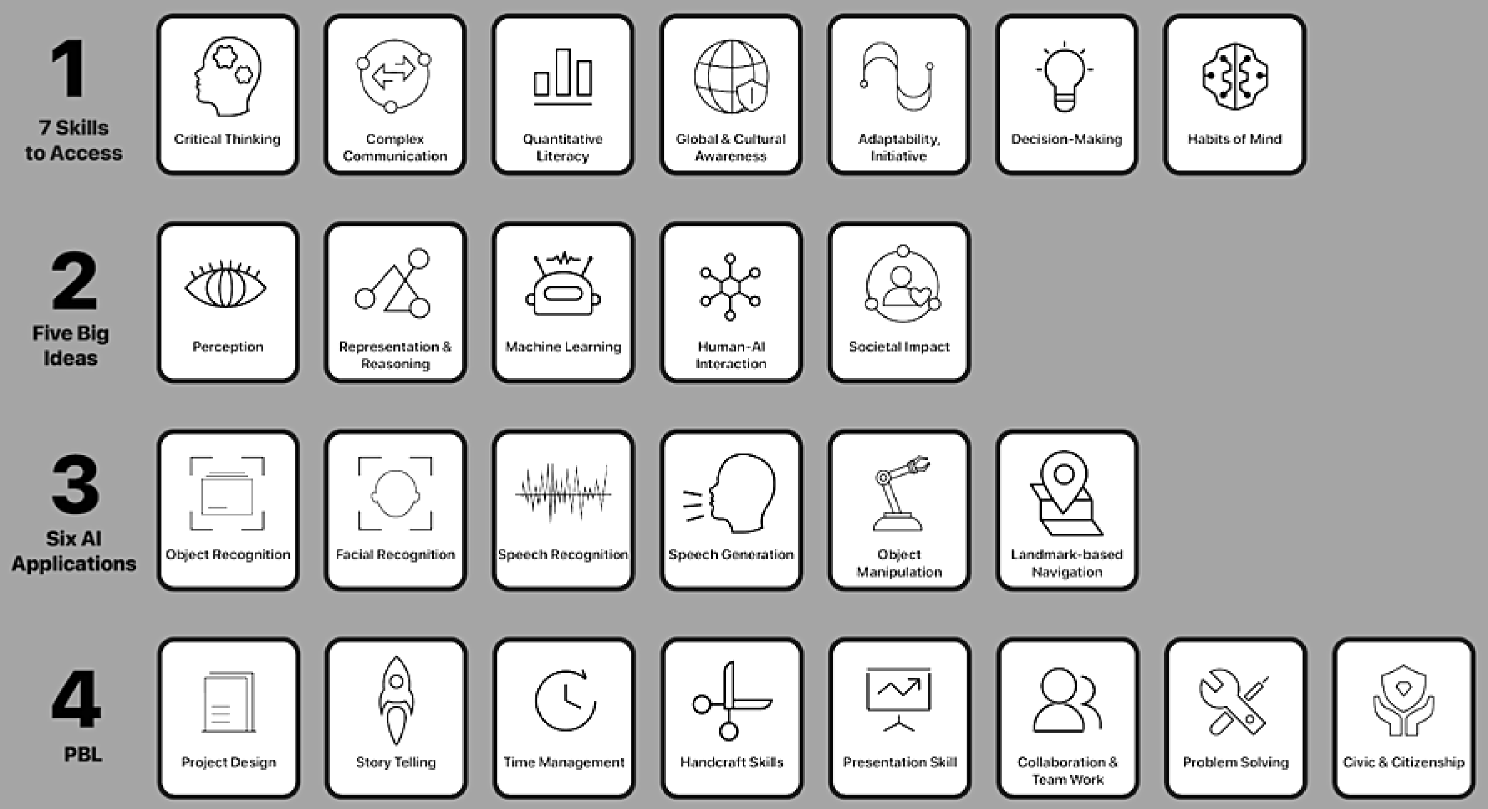

Figure 1 - 7 Skills, 5 Big Ideas, 6 AI Applications, and PBL (Project Based Learning)

CREATING THE CURRICULUM

CO-DESIGN PROCESS

As we utilize research-based processes to create and sustain effective shifts in pedagogy and instructional delivery, our work leverages co-design methodologies to support targeted teams. In this process, design is understood to be an iterative process with opportunities to (re)evaluate, make improvements, and step into previous stages to re-examine and critique. Similarly, we understand that design for learning in classroom spaces demands a flexibility to processes given the various, idiosyncratic ways in which learners grow and develop over time. To that end, we leverage an ADDIE model as informed by Kemp's Instructional Design model, transforming instructional/curriculum planning into an explorative process that is interactive and inclusive. This approach allows learners, educators, and other collaborators opportunities to share planning, teaching, and assessment ideas. We embrace this design process as it allows for innovation, experimentation, and critical reflection. Our work this year was limited in partnership, primarily relying on connection to Abbas Hawsawi as a means of collaboration and co-design. We understand that, moving forward, we will have the benefit and opportunity of partnering with the larger Instructional Technology team at NCS which will allow for deeper collaboration and active implementation.

KEMP'S MODEL

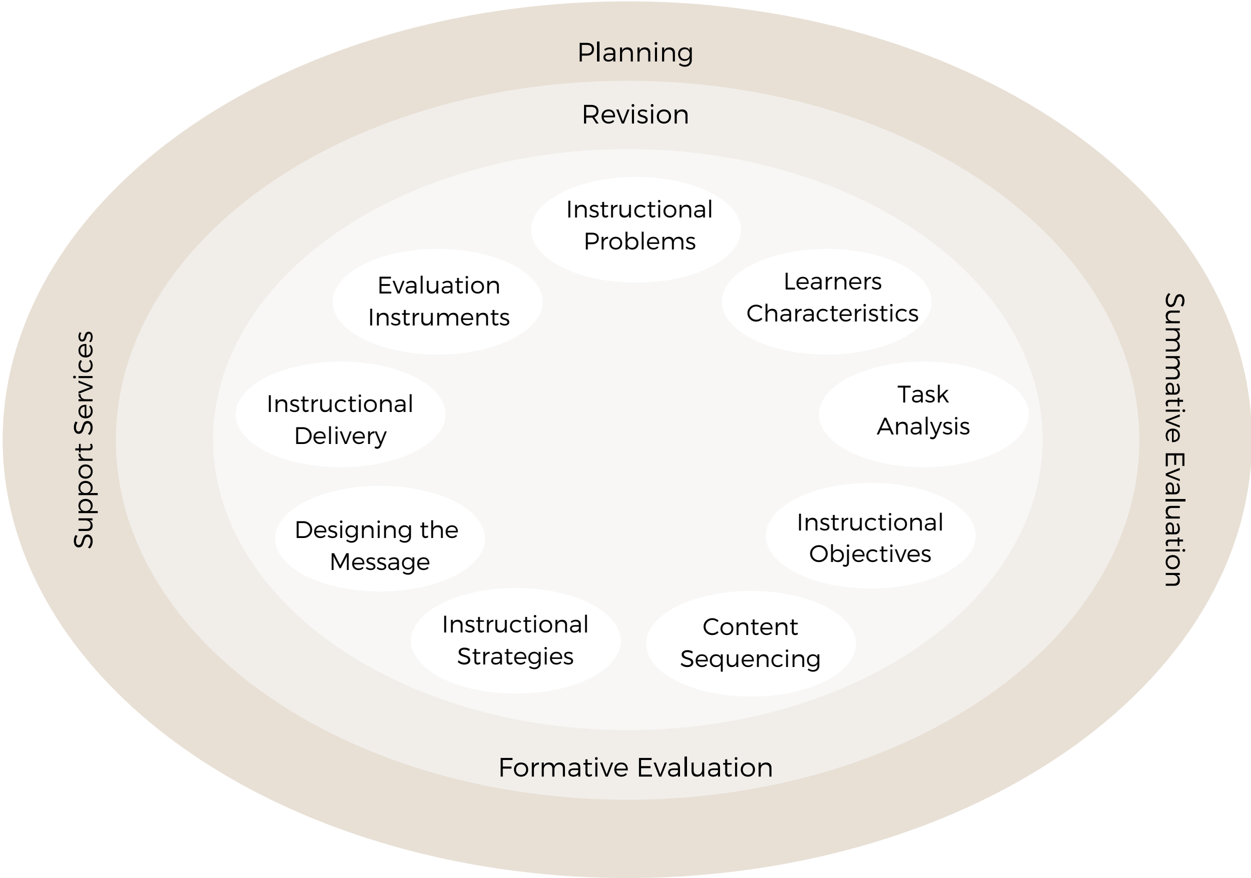

ADDIE MODEL OF INSTRUCTIONAL DESIGN

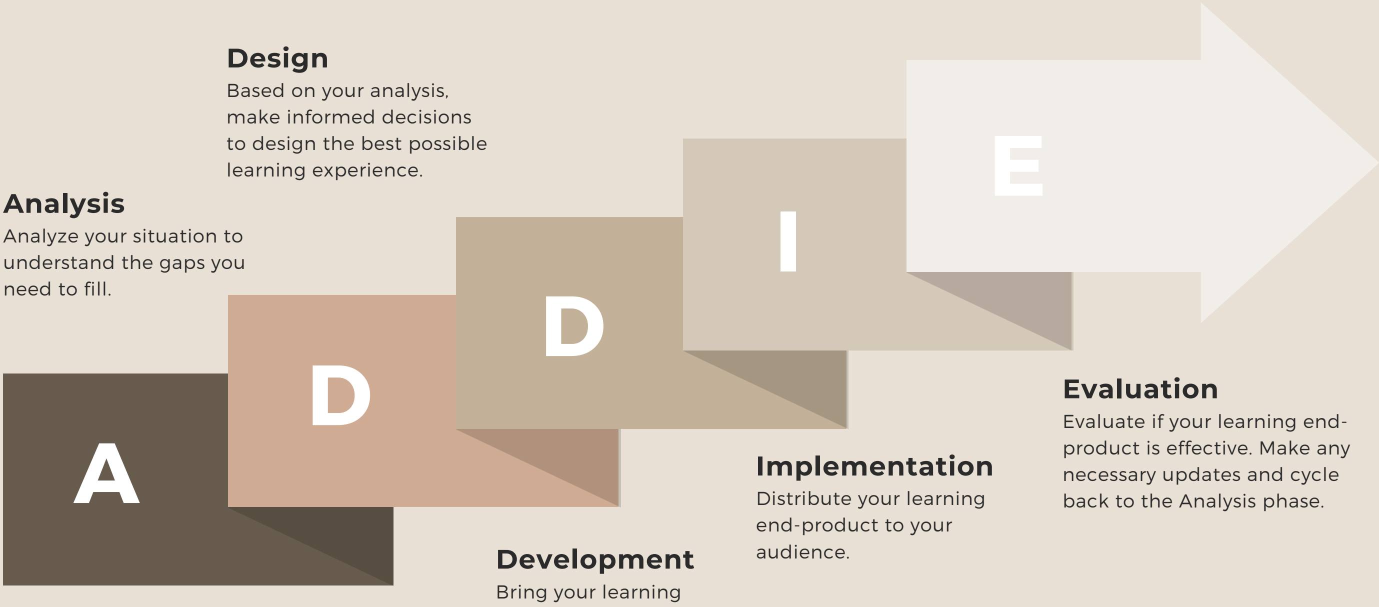

Analysis
Analyze your situation to understand the gaps you need to fill.

Design
Based on your analysis, make informed decisions to design the best possible learning experience.

Development
Bring your learning experience to life by building your end-product

Implementation
Distribute your learning end-product to your audience.

Evaluation
Evaluate if your learning end-product is effective. Make any necessary updates and cycle back to the Analysis phase.

UNIT OF INQUIRY SELECTION

As provided by our collaborator, Abbas, our anchor lessons rotate through the designated Units of Inquiry indicated in the scope and sequence provided. This aimed to ensure that integral elements of the International Baccalaureate curriculum were foundational to the emerging artificial intelligence curriculum across all grade levels. The follow list provides the anchor lessons in parentheses with corresponding Unit of Inquiry:

Who We Are (Facial Recognition)

- An inquiry into the nature of the self; beliefs and values; personal, physical, mental, social and spiritual health; human relationships including families, friends, communities, and cultures; rights and responsibilities; what it means to be human.

How We Organize Ourselves (Decision Trees and Machine Learning)

- An inquiry into the interconnectedness of human-made systems and communities; the structure and function of organizations; societal decision-making; economic activities and their impact on humankind and the environment.

Where We Are In Place And Time (Computer Vision & Its Possibilities)

- An inquiry into orientation in place and time; personal histories; homes and journeys; the discoveries, explorations and migrations of humankind; the relationships between and the interconnectedness of individuals and civilizations, from local and global perspectives

How The World Works (Speech Recognition)

- An inquiry into the natural world and its laws; the interaction between the natural world (physical and biological) and human societies; how humans use their understanding of scientific principles; the impact of scientific and technological advances on society and the environment.

How We Express Ourselves (Neural Nets)

- An inquiry into the ways in which we discover and express ideas, feelings, nature, culture, beliefs and values; the ways in which we reflect on, extend and enjoy our creativity; our appreciation for the aesthetic.

Sharing the Planet (Climate Change and AI)

- An inquiry into the rights and responsibilities in the struggle to share finite resources with other people and with other living things; communities and the relationship within and between them; access to equal opportunities; peace and conflict resolution.

ASSESSMENT OF EXISTING PLANS

Upon receiving documentation, we began categorization as to how the various types of activities presented might best be reformulated into a more traditional curricular document. While the existing documents place activities/ideas into what might best be described as a scope and sequence, the difficulty remains of how to differentiate learning into specific grade level bands to create a cohesive learning experience from KG3 through Secondary. Specifically, there is not an AI activity outlined so the assessment of specific learning design and outcomes here is not possible. We proceeded to build a comprehensive lesson plan and activity set from scratch to accommodate this need. The excerpt below reflects a piece of the existing documentation. Though links exist to external resources, Code.org most notably, we have worked extensively to supplement these learnings with lesson plans on computer science in addition to translating the document into a more formalized curriculum map. It should also be noted that spaces/gaps in programming here are locations where substantial lesson support have been given and now no longer exist due to our processes.

Line of inquiry (Key concepts)			WEEK	Period	Activity	Activity	learning outcomes	
3	4	5					A Part	B Part
Factors impact values and beliefs (Perspective) Rights and responsibilities of children (Responsibility) Connections with members of a diverse global community (Connection)	Cultures and traditions (Form) Different components of culture (Perspective) Compare and contrast what connects us culturally (Connection)	as an assessment to learn about the hard and soft. use world (Scratch)	1	p1	c-1-1 How to act in the lab c-1-2 How to Handle devices	c-1-1 How to act in the lab c-1-2 How to Handle devices		
				p2	b-3-3 Free buillding	b-3-3 Free buillding		
			2	p1	c-2-7 We the Digital Citizens	c-2-7 We the Digital Citizens		
				p2	c-2-10 Digital Trail	c-2-10 Digital Trail		
			3	p1	a-3-1 Hardware & Software	a-3-1 Hardware & Software		
				p2	~~b-2-13 Good Morning Machine~~	a-3-3 Creating Algorithms		
			4	p1	a-3-3 Creating Algorithms	a-3-4 Intro to Scratch Programming		
				p2	a-3-4 Intro to Scratch Programming		b-2-13 Good Morning Machine	intro to scratch robot programming
			5	p1		a-3-6 Programming A Sprite		
				p2		a-3-8 Debugging In Scratch		
Types of goods and services (Form) Production, distribution and consumption of goods and their environmental impact (Function) Interdependence of producers and consumers (Connection	Entrepreneurship and economic principles (Function) Economic development in history (Change) Impact of the economy on humankind and the environment (Responsibility)		6	p1	a-3-11 Conditionals (unplugged)	a-3-9 Animations Using Loops		
				p2	a-3-12 Making Decisions Using Conditionals	a-3-10 Sprite Conversations		
			7	p1	~~b-2-10 Big Bus (G3)~~	a-3-11 Conditionals (unplugged)		
		focacing on conditions using color *events		p2	Driving base with color sensor (G4,5)	a-3-12 Making Decisions Using Conditionals		
			8	p1	c-3-3 Mini PBL goods distribution system, Service evaluation system	a-3-13 Keeping Score With Variables		
				p2				
			9	p1			b-... Dri... sor (G4,5) — focacing on conditions using color *events scratch robot programming, revision of the 6th lesson using robot.	focacing on conditions using color *events scratch robot programming, revision of the 6th lesson using robot.
				p2				
			10	p1	c-3-5 Showcase (in class presentation+video)			
				p2				
Important discoveries and explorations (perspective) Current application of discoveries (form)	Reasons why people migrate (Causation) Migration of humankind (Change)	lesson about how the communication has been changed through the time and how the AI get involved in that evolution	11	p1	c-1-4 Digital literacy (Evolution of communication tech	c-1-4 Digital literacy (Evolution of communication technolo	lesson about how the	
				p2	AI related activity	AI related activity	Intruduce the different types of text input and the basics of "pages"	
			12	p1	c-1-4 Digital literacy (Textinputand "Pages")	c-1-4 Digital literacy (Textinputand "Pages")	Build a service robot	
				p2				

The image below is an example of the newly developed curriculum map with a scope and sequence for Grade 7 to demonstrate the evolution of content from what ReadyAI started with to where we have arrived.

Curriculum Map - G7

Subject: Computer Science **Time Frame:** 2022-2023 School Year

Grade: Seventh Grade

Essential Question:
- How can I use Python programming to create art?
- Can a robotic arm be created with robotics?
- What ethical considerations do programmers need to consider when programming autonomous cars?
- Can I simulate the heart using a robotics kit and what benefits might this lead to in biomedical engineering?

Objectives:
- Students will design a vinyl cutout for their laptop computers and program the design using Python programming and Turtle Art
- Students will create a mini drum kit using Makey Makey and experiment with musical sounds
- Students will explore the heart and use a robotics kit to simulate a beating heart
- Students will discover how robotic arms are used in society both for manufacturing purposes and prosthetics

Standards:
- **ISTE 1.2 Digital Citizen:** Students recognize the rights, responsibilities, and opportunities of living, learning, and working in an interconnected digital world, and they act in ways that are safe, legal, and ethical
- **ISTE 1.6 Creative Communicator:** Students communicate clearly and express themselves creatively for a variety of purposes using the platforms, tools, styles, format, and digital media appropriate to the goals
- **ISTE 1.4 Innovative Designer:** Students use a variety of technologies within a design process to identify and solve problems by creating new, useful or imaginative solutions.
- **Five Big Ideas of AI: Idea #5:** Societal Impact - AI can impact society in both positive and negative ways

Skills:	Coding Programs:	Materials:	Activities:
• Biomedical engineering with robotic arms • Machine learning in autonomous vehicles • Art and coding with vinyl designs • Connecting LEDs and motors to robotic kits	• Python with AutoAuto • LEGO Sprike Coding • iRobot Coding • Turtle Art • Makecode • AutoAuto Python	• Chromebooks/ iPads • Root Robot • Makey Makey • AutoAuto Cars • Hummingbird Kits • Hydraulic Kits • LEGO Spike Prime	• Use Turtle Art to create patterns on vinyl for laptop • Use AutoAuto cars to avoid humans & pets • Create a heart with a • Hummingbird Kit • LEGO Spike robotic arm with coding • Makey Makey drum kit creation

Resources:
- Brainbird Technologies teaching unit
- LEGO Spike Prime App
- iRobot education learning guide
- Moxie robot app and tutorials
- Common Sense Media: Digital Skills

Assessment:
- Rubric for Turtle Art coding and design
- LEGO Spike robotic arm goals
- AutoAuto online Python assessments
- Hummingbird Kit project completion

Computer Science

SCOPE & SEQUENCE

Seventh Grade | 2023

UNIT 1 | UNIT 2 | UNIT 3 | UNIT 4

INTRO TO CLASS	
WEEK 1:	**Title: Robotics and Tools Guidelines** 1. Review new equipment and projects that the 7th graders will use 2. Have students practice using the machines and equipment in class
WEEK 2:	**Title: Create a Story on Scratch to Represent Your Break** 1. Using Scratch, create a sprite and background to tell about break 2. Add movement and sound effects to your Scratch story and share
UNIT 1: Autonomous Vehicles	
WEEK 3:	**Title: AutoAuto Car Maze** 1. Program the cars to complete a maze in the virtual world 2. Create a maze using the mats and race the cars in the maze
WEEK 4:	**Title: AutoAuto Python** 1. Create a course with two obstacles, a cat and a tree 2. Discuss the decision making that programmers make for self driving cars, as well as the impact on society
WEEK 5:	**Title: Code Monkey Chatbot** 1. Complete the Code Monkey Activity to code a chatbot 2. Discuss multiple ways that chatbots are used in today's society
WEEK 6:	**Title: Hydraulic Crane** 1. Learn about how hydraulics are used by machines and robots 2. Create a crane using hydraulics, specifically syringes and tubes
WEEK 7:	**Title: Hydraulic Press** 1. Examine how a hydraulic press works and the strength of the press 2. Create a miniature hydraulic press in class
UNIT 2: Beating Heart Hummingbird Kit	
WEEK 8:	**Title: Hummingbird Makecode** 1. Use Makecode to add coding blocks that control the Hummingbird 2. Practice plugging in the white, red, and black wires into the Servo port
WEEK 9:	**Title: Hummingbird LEDs & Sensors** 1. Program the light sensor to control a single LED when it is dark or light 2. Write a program that makes an LED turn on and off with another sensor
WEEK 10:	**Title: Hummingbird Heart** 1. Using cardboard, create a heart attached to a Hummingbird kit 2. Have the heart light up like a heart beat and try out different rhythms
WEEK 11:	**Title: Common Sense Media Ad Detective** 1. Discuss how companies target their ads towards certain individuals 2. Discuss privacy settings and complete the media ad detective activity
WEEK 12:	**Title: Create an Online Appropriate and Inappropriate Post** 1. Practice writing posts with photos that are appropriate to post online 2. Discuss information that is important to keep private and not share on social media for security or privacy reasons
UNIT 3: Robotic Arms	
WEEK 13:	**Title: LEGO Spike Robotic Arm** 1. Use the LEGO Spike app to complete the robotic arm design 2. Discuss how the motors and servos work to make the arm move
WEEK 14:	**Title: LEGO Spike Code** 1. Program the robotic arm to pick up LEGO pieces 2. Add functionality so that the arm could be used in more situations
WEEK 15:	**Title: Talk to a Biomedical Engineer** 1. Write questions to ask a biomedical engineer who utilizes AI 2. Interview the engineer to learn about the engineering design process
WEEK 16:	**Title: Root Robot & Fireplace** 1. Discuss how a fireplace robot might change colors to imitate fire 2. Use the root robot to flicker different colors in a mini fireplace
WEEK 17:	**Title: Root Robot** 1. Discuss how iRobot makes Root Robot and Roomba 2. Discuss how Root Robot could be used for cutting the grass, etc.
UNIT 4: Turtle Art & Vinyl Cutouts	
WEEK 18:	**Title: Turtle Art & Python Coding** 1. Use the turtle art tutorial to learn about drawing shapes with Python 2. Write a program to draw different sized spirals
WEEK 19:	**Title: Turtle Art & Vinyl Cutting** 1. Create an new, original design to display on a laptop 2. Using Python, program the Turtle Art design and use the vinyl cutter to attach to the students' laptop
WEEK 20:	**Title: Explore Chrome Music Lab & Blob Opera** 1. Experiment with Blob Opera and explore how the app uses AI 2. Compose a simple melody using Chrome Music Lab
WEEK 21:	**Title: Makey Makey Controller** 1. Connect Makey Makey to the Pinball game on the app 2. Control the game using tinfoil cutouts attached to pool noodles
WEEK 22:	**Title: Makey Makey Drum Kit** 1. Students create their own drum kits with partners 2. Connect a Makey Makey and play the drums to create a new beat
WEEK 23:	**Title: ToxiCode** 1. Complete the ToxiCode Compute It Hour of Code 2. Discuss strategies for following the code and completing the activity
END OF CLASS	
WEEK 24:	**Title: Create a slideshow for a student in the future** 1. Each student creates one slide to prepare a future student for CS 2. Use Generative AI to add an image to the slideshow
WEEK 25:	**Title: Robot Playground** 1. Using pool noodles for barriers, control Dash Robot and Root Robot 2. See who can most successfully dash to the other side of the room

COLLABORATIVE CREATION

As we approach this project through a lens of sustainability, ideally transferring ownership of the design and development over time to internal members of the CS/AI faculty and staff at NCS, we understand that singular external content development is not enough. Our objective is to increase internal capacity and expertise while supporting technology integration so as to assure a consistent pedagogical model. Working in this way, increased levels of integration can deepen cohesion, collaboration, and creation, leveraging all members of the team's skills. The trade off in this approach is that it requires a deeper initial investment of time and effort for internal team members to adjust to new methods, navigate differences in style, and forge joint pathways when they already have very full workloads. The model below, adapted from Bamber & Stefani's research, is useful for understanding the process we have leveraged in collaborative creation.

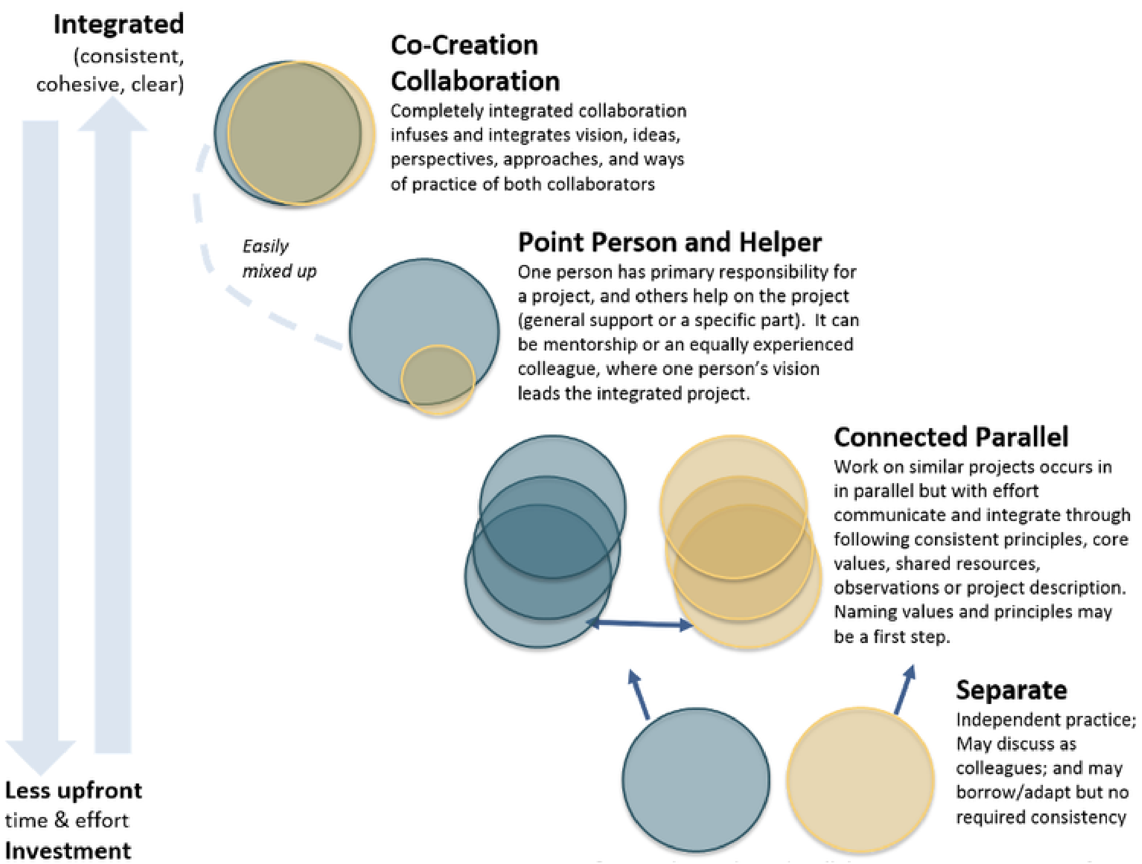

We also approach the design or processes, projects, and goals through this lens of co-design and collaboration. This best ensures that the content developed is integrated with robust fidelity so as to hit targeted learning outcomes and alignment to IB-established curricular goals. We recognize that this was a deep difficulty this past year and aim to achieve a greater level of collaboration and co-design with the newly introduced Technology Integration team. This will allow learners to develop the core IB learner profile characteristics. The collaboration model below highlights the progression of how co-designing each element facilitates scaffolding of development into more complex tasks.

COLLABORATION - MODEL

The 5 Models for Collaboration

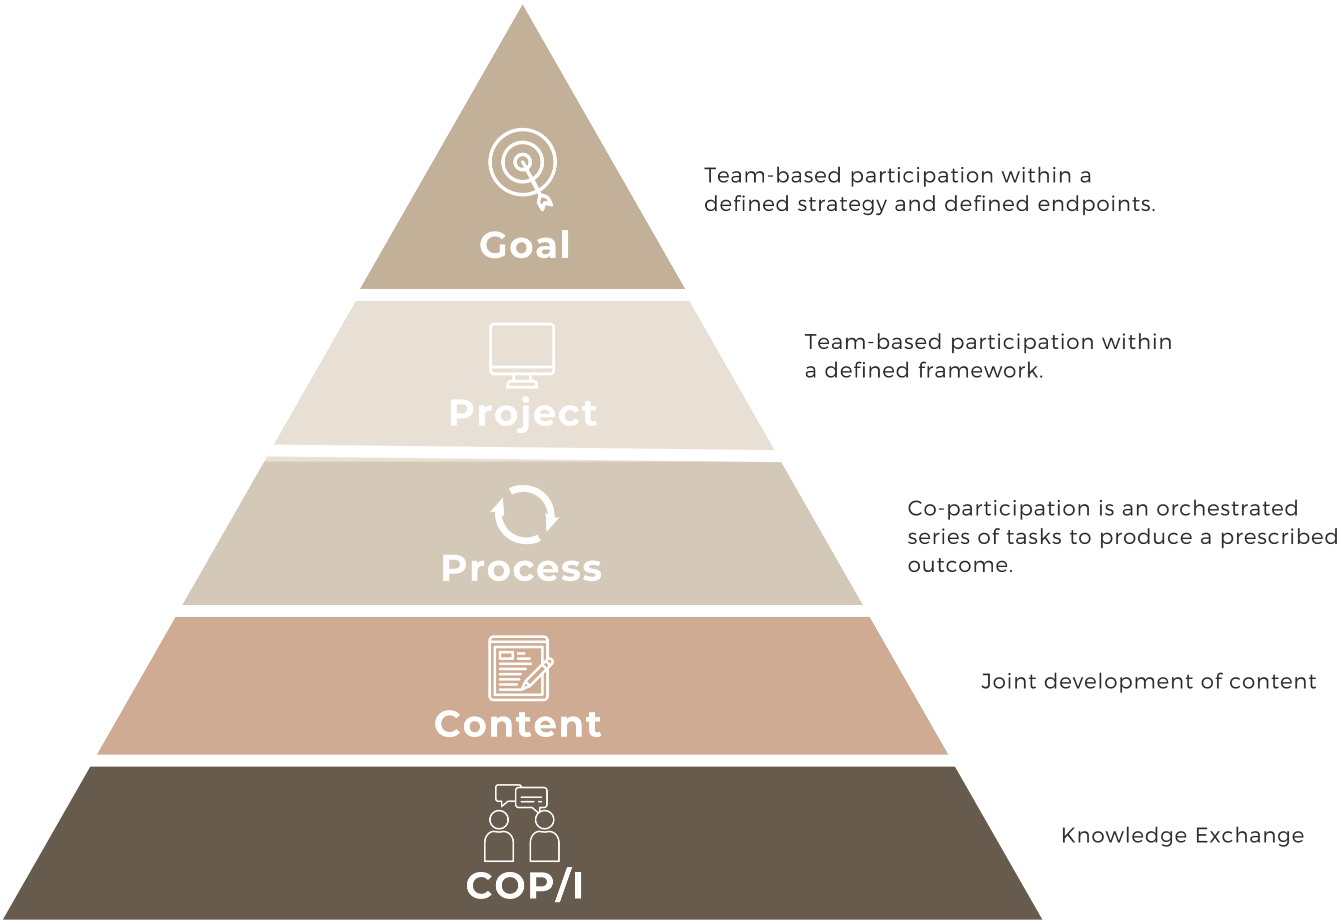

FORMATIVE/SUMMATIVE ASSESSMENT CRITERIA

As we continue to move forward in the dissemination and testing of anchor lessons, we have intentionally built in mechanisms and instruments to make use of IB-aligned practices. In formative assessment, we have been diligent in suggesting in-class prompts and dialogue to support qualitative and quantitative data gathering to ensure adaptive practices can be used by on-site educators. Far too often, the rigidity of lesson plans have been used to adhere to pacing expectations in moving the curriculum forward.

However, in our meetings, we have encouraged consistent adaptation while marking out project-based learning approaches that demonstrate successive mastery of key learning objectives. The daily observations of students as well as the completion of projects act as key evidence for on-site teachers to make critically informed decisions about how best to serve their students. The model below indicates IB-suggested practices that we have attempted to develop throughout our learning materials and curriculum.

FORMATIVE ASSESSMENT AS A DIALOGUE AND PROCESS

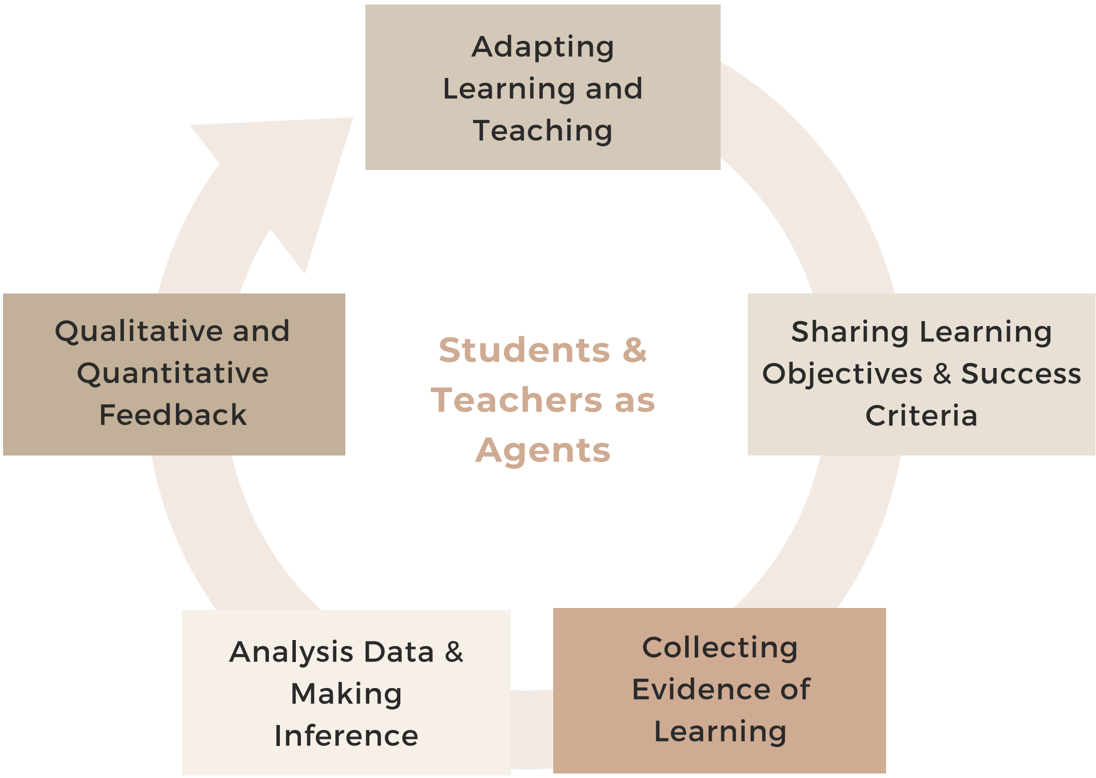

As we consider summative assessment in these efforts, we are mindful that the core of IB's mission is to support learners who contribute to the development of "a better world". As such, the projects that have been developed in the curriculum replicate situations, contexts, and problems that emerge as part of the real-world complexities of AI and CS. Projects enable students to demonstrate knowledge, growth, skills, and dispositions in idiosyncratic ways while also generating objective evidence that supports the educator in determining 'grades' and achievement.

ANCHOR LESSON EXAMPLES

The following are two examples of the designed anchor lessons. The two presented, Neural Nets and Speech Recognition, demonstrate the depth of learning created by our team at ReadyAI while scaffolding the experience across grade levels.

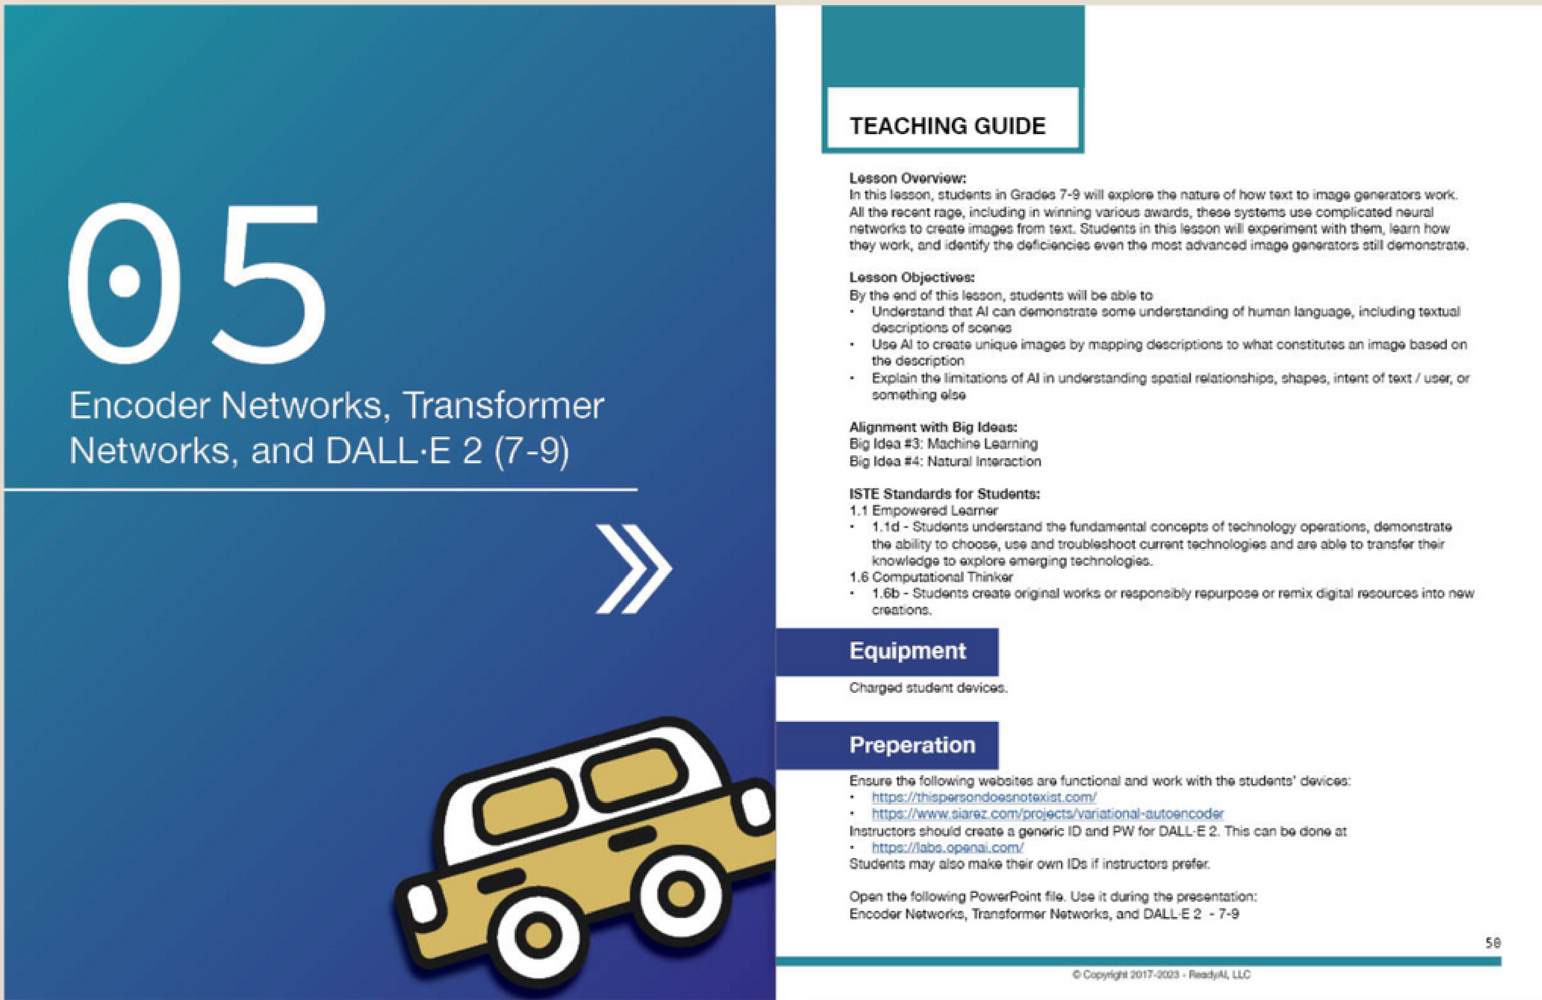

05

Encoder Networks, Transformer Networks, and DALL·E 2 (7-9)

TEACHING GUIDE

Lesson Overview:
In this lesson, students in Grades 7-9 will explore the nature of how text to image generators work. All the recent rage, including in winning various awards, these systems use complicated neural networks to create images from text. Students in this lesson will experiment with them, learn how they work, and identify the deficiencies even the most advanced image generators still demonstrate.

Lesson Objectives:
By the end of this lesson, students will be able to
- Understand that AI can demonstrate some understanding of human language, including textual descriptions of scenes
- Use AI to create unique images by mapping descriptions to what constitutes an image based on the description
- Explain the limitations of AI in understanding spatial relationships, shapes, intent of text / user, or something else

Alignment with Big Ideas:
Big Idea #3: Machine Learning
Big Idea #4: Natural Interaction

ISTE Standards for Students:
1.1 Empowered Learner
- 1.1d - Students understand the fundamental concepts of technology operations, demonstrate the ability to choose, use and troubleshoot current technologies and are able to transfer their knowledge to explore emerging technologies.

1.6 Computational Thinker
- 1.6b - Students create original works or responsibly repurpose or remix digital resources into new creations.

Equipment

Charged student devices.

Preperation

Ensure the following websites are functional and work with the students' devices:
- https://thispersondoesnotexist.com/
- https://www.siarez.com/projects/variational-autoencoder

Instructors should create a generic ID and PW for DALL·E 2. This can be done at
- https://labs.openai.com/

Students may also make their own IDs if instructors prefer.

Open the following PowerPoint file. Use it during the presentation:
Encoder Networks, Transformer Networks, and DALL·E 2 - 7-9

58

© Copyright 2017-2023 - ReadyAI, LLC

Lesson Orientation

Warm-Up Exercise Review slides 1-7 in the PowerPoint. Highlight how the open software DALLE 2 has revolutionized art by using AI to generate high quality images. Instructors may also share the following video that demonstrates how AI art has won competitions over real human artists: AI generated art piece takes first at the state fair[1] **(slide 2)**	5 minutes
Presentation: Auto-encoder networks & Denoising auto-encoders Explain that the earlier slides help demonstrate what Transformer Networks and Image Generators are capable of. Explain to students that we will explore these and how they work in this lesson **(slide 8).** **Slides 9-11** Computers are capable of generating images, but this takes a lot of training. People begin with a real image, break it into code, and recreate the image from the code. How does the computer learn to do this? It uses a type of neural network called an auto-encoder, like this:	10 minutes

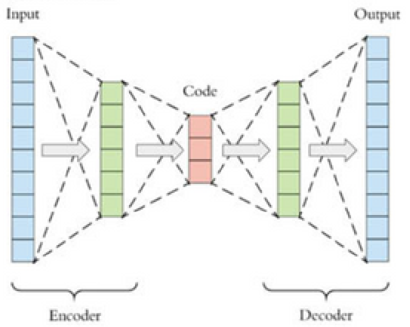

1 https://youtu.be/yYTjyfYb0qY

We start by collecting a bunch of clean images (no noise). Let's say they're all images of people's faces or numbers.

Slide 11: The encoder reduces the image to its essential features. It turns the image into a code.

Slide 12: The decoder can reconstruct the image from the code. It often looks slightly different than the original.

Slide 13: But now computers can do a lot more than just show us what we gave it in the first place. It can turn a "noisy" image into a more clean image.

This is useful because noisy images might come from many places in real life: a shaky hand with a camera, photographs taken on a hazy day, a scan of an image in a newspaper. Explain that computers can learn to remove noise from images.

See Slides 14-17.

Now we train the auto-encoder by selecting an image and making a noisy version. The noise could be speckle-type noise, or blurring, or something else.

We feed the noisy image to the encoder network (blue box on the left) and show the original clean image to the decoder (blue box on the right) and we tell it: this clean image is what you should be producing.

Initially the decoder produces garbage, but the neural net learning algorithm adjusts the weights in the network so that eventually the decoder's output looks pretty much like the clean image it's supposed to produce.

Once we've trained up the network, we can show it a new noisy face it hasn't seen before, and it will do a pretty good job of removing the noise.

This only works if we have enough training examples; we'll need thousands of faces and a fast computer.

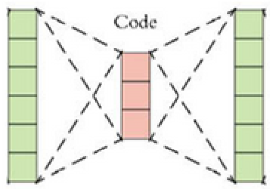

This auto-encoder is also doing something else interesting. Notice that the pink layer in the middle, labeled "Code", has fewer units than the green layer, which in turn has fewer units than the blue layer:

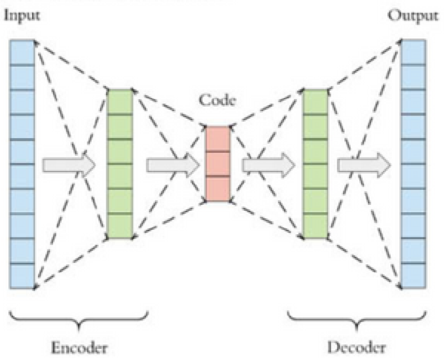

What this means is the auto-encoder has come up with a way of extracting the "essence" of a face image that requires fewer numbers than the original image. It's able to do this because faces have similar structure: hair on top, then two eyes below that, a nose in the middle, a mouth and chin at the bottom. And eyes have a certain general shape, as do noses. So it doesn't need to keep track of every pixel; it can describe each face more succinctly. This will be important later.

If you feed in a new face to the encoder part of the network, the resulting Code (pink layer) serves as a kind of shorthand for the face, using fewer numbers. Whenever you like, you can feed that code into the corresponding decoder network and get back something that is pretty close to the original face. But what if you made up a new code at random and fed that into the decoder? Then you'd get back a different face!

Noisy images might come from many places in real life: a shaky hand with a camera, photographs taken on a hazy day, a scan of an image in a newspaper. Explain that computers can learn to remove noise from images. Show them pictures of noisy images and the de-noised versions:

Slide 18: On the left, we see a noisy image. Computers can "denoise" the image and produce something close to the original.

Slide 19: How does this happen? Basically, there is a noisy image, either taken that way or made to be noisy to train the computer. The computer removes the noise and renders the image as it thinks it should look:

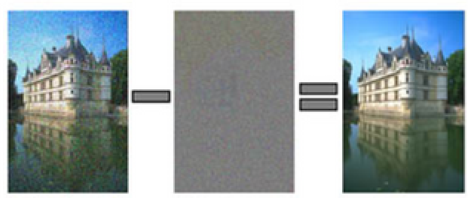

Practice: Variational Auto-Encoders & "This Face Does Not Exist" **Slides 19-22** Show students the Variational Auto-Encoder website at https://www.siarez.com/projects/variational-autoencoder	10 minutes

This site uses a neural network where the code is just 10 numbers, and we can control those numbers using the sliders. (The fact that there are 10 sliders and 10 distinct digits is just a coincidence; there is no association between the nth slider and the nth digit.) If you click on one of the 10 digit images it will run that image through the network and show you both the decoded version and the slider positions that produce that decoded image. If you move the sliders around you are changing the Code and you can see how the output changes. You can also draw a digit using the mouse and run that through the network. Playing around with this is pretty cool. Allow students a chance to play with the code using the sliders.

This Person Does Not Exist
Slides 23-31
Visit www.ThisPersonDoesNotExist.com and examine some generated faces. It's doing what we described above: making up random codes and decoding them to create new faces. But it's not perfect.

Show them how to look for imperfections such as mismatched earrings, incomplete glasses, funny eyes, etc. The neural net isn't perfect, but it's pretty good!

Ask students to begin to identify inconsistencies in images they see. Allow some students to share their images by connecting their devices to the projector.

Production: Transformer Networks & DALL·E 2

20 minutes

Slides 32-33: Just as the auto-encoder turns images into codes and back to images, transformer networks turn text into codes and back to text.

For instance, Google uses them to help understand search queries and also to generate some of the answers you get back from a search. ChatGPT uses them to generate text. Much of today's natural language processing is done using transformers.

Slide 34: Introduce DALL·E 2. From the user's point of view, it converts text to images. But what it's really doing is converting text codes to image codes. We use the transformer to turn the text into text codes, and we train a neural network to generate image codes from those text codes. Then we use the decoder (from the variational autoencoder) to turn those novel image codes into images.

Share the slides that introduce DALL·E 2. Text to image generators use the encoder half of a transformer network. Then, a neural network translates the code into image code. The computer then shows an image of what it thinks we wanted to see, such as a "A dog sitting under a palm tree."

Slide 35: However, explain that the network produces pictures that may be inconsistent with what they want. Sometimes the images are inaccurate based on the numbers given to it, such as in "Three cats and three dogs sitting on the grass."

Have them play with DALL·E 2 that maps text to images. Run through some examples below. Show some cases where the images are malformed, and some cases where it misunderstood the language. Explain the process that is happening in the background, i.e. the text is turned into code, the code is run through a neural network to produce different image code, and the computer then renders the image code as the image we see.

Let's see how the system does with numbers first. Try some of these experiments:
- One dog on the street.
- Two dogs on the street.
- Three dogs on the street.
- Four dogs on the street.
- Five dogs on the street.
- One cat and one dog sitting on the grass.
- One cat and two dogs sitting on the grass.
- One cat and three dogs sitting on the grass.
- Two cats and one dog sitting on the grass.
- Two cats and two dogs sitting on the grass.
- Two cats and three dogs sitting on the grass.

- Three cats and one dog sitting on the grass.
- Three cats and two dogs sitting on the grass.
- Three cats and three dogs sitting on the grass.

Slide 36: We can continue to test this system by introducing a new form of language to it, i.e. descriptions. Let's do the first one together:
- A vehicle composed of two wheels held in a frame one behind the other, propelled by pedals and steered with handlebars attached to the front wheel.

Does this meet the requirements set out by the description? Yes–using natural language, the AI system is understanding the language input. It generates the image based on what is being offered. It understands colors and spatial relationships. It's not as good with counting things or ordering who is doing what, as you learned earlier.

But is it really what that description is looking for? No–this helps us see that AI often lacks the 'common sense' reasoning that we associate with true intelligence.

Slide 37: Try this now! Let's borrow a description of an animal from a dictionary.
- A small domesticated carnivorous mammal with soft fur, a short snout, and retractable claws. It is widely kept as a pet or for catching mice, and many breeds have been developed.

What animal is being defined? A cat: But is this what DALL·E 2 will produce? Or will it produce some cats and some mice, as it did in this test? Try it out.

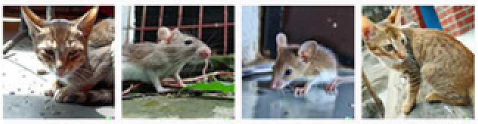

Why don't you try some of these? Which of these text prompts does DALL·E 2 do a good job on, and which does it fail?
- A large motor vehicle carrying passengers by road, typically one serving the public on a fixed route and for a fare.
- A small vessel propelled on water by oars, sails, or an engine.
- A connection point by which firefighters can tap into a water supply.
- A machine next to a parking space in a street, into which the driver puts money so as to be authorized to park the vehicle for a particular length of time.
- A device consisting of a circular canopy of cloth on a folding metal frame supported by a central rod, used as protection against rain or sometimes sun.
- A separate seat for one person, typically with a back and four legs.
- An appliance or compartment which is artificially kept cool and used to store food and drink.
- A mechanical or electrical device for measuring time.
- An instrument used for cutting cloth, paper, and other thin material, consisting of two blades laid one on top of the other and fastened in the middle so as to allow them to be opened and closed by a thumb and finger inserted through rings on the end of their handles.
- A large plant-eating domesticated mammal with solid hoofs and a flowing mane and tail, used for riding, racing, and to carry and pull loads.
- A long curved fruit which grows in clusters and has soft pulpy flesh and yellow skin when ripe.
- A domesticated carnivorous mammal that typically has a long snout, an acute sense of smell, nonretractable claws, and a barking, howling, or whining voice.
- An organ of soft nervous tissue contained in the skull of vertebrates, functioning as the coordinating center of sensation and intellectual and nervous activity.
- An American multinational technology company that focuses on artificial intelligence, search engine, online advertising, cloud computing, computer software, quantum computing, e-commerce, and consumer electronics.

- A large keyboard musical instrument with a wooden case enclosing a soundboard and metal strings, which are struck by hammers when the keys are depressed. The strings' vibration is stopped by dampers when the keys are released and can be regulated for length and volume by two or three pedals.
- A type of digital currency in which a record of transactions is maintained and new units of currency are generated by the computational solution of mathematical problems, and which operates independently of a central bank.
- A large thick-skinned semiaquatic African mammal, with massive jaws and large tusks.
- A machine resembling a human being and able to replicate certain human movements and functions automatically.

Conclude by reiterating how DALL·E 2 works: Just as we have encoder networks for images, we also have encoder networks for language, called transformer networks.

So we can take a description like "a dog sitting underneath a palm tree" and use the transformer to turn that into a collection of numbers.

Now, to train DALL·E 2, we need a large collection of images with captions that describe the image. For each image/caption pair, we use the transformer to turn the caption into a caption code, and we use the autoencoder to turn the image into an image code. Then we train a third neural network to map the caption code to the image code.

Training takes a long time and many image/caption pairs. After the training is complete, we can make up a new caption, and this gets turned into a new caption code by the transformer, which is then used to generate a new image code.

This new image code is run through the decoder network and we get a new image, never seen before. And that's how DALL·E 2 is able to turn your descriptions into images!

Extension

Give students an assignment that requires them to use DALL·E 2 to produce work for the class. Ask them to identify images that work or make sense and images that do not work or make sense based on the assignment prompt.

04
Interacting with
Intelligent Agents (7-9)

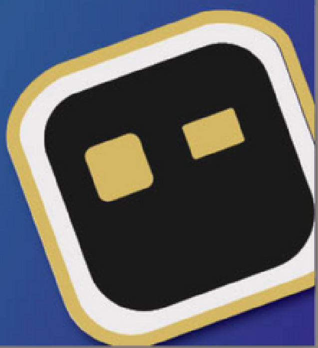

TEACHING GUIDE

Lesson Overview:
Interacting with intelligent agents has become a part of everyday life. In this lesson, students will explore first how sound works and is represented on a computer. Then, students will use the SpeechDemo tool to determine how computers render the sound and make informed guesses as to what users have said. This also presents limitations that students should be aware of. Finally, using Scratch, students will create a tool that allows them to engage the programming of intelligent agents.

Lesson Objectives:
By the end of this lesson, students will be able to
- Demonstrate how different sounds in English result in different spectrogram patterns.
- Explaining how ambiguity in the speech signal leads to multiple possible transcriptions
- Create a Scratch program that employs speech recognition.

Alignment with Big Ideas:
Big Idea #1: Perception
Big Idea #4: Natural Interaction

ISTE Standards for Students:
1.2 Digital Citizen
- 1.2b - Students engage in positive, safe, legal and ethical behavior when using technology, including social interactions online or when using networked devices.

1.3 Knowledge Constructor
- 1.3d - Students build knowledge by actively exploring real-world issues and problems, developing ideas and theories and pursuing answers and solutions.

1.7 Global Collaborator
- 1.7a - Students use digital tools to connect with learners from a variety of backgrounds and cultures, engaging with them in ways that broaden mutual understanding and learning.

Equipment

Charged student devices. If desired, instructors may also provide the premade demo found here: Intelligent Assistant with Keywords (Cognimates) Activity[1]

Preperation

Ensure the following websites are functional and work with the students' devices:
- https://spectrogram.sciencemusic.org/
- http://www.cs.cmu.edu/~dst/SpeechDemo
- https://codelab.cognimates.me/

1 https://docs.google.com/document/d/1epxNTvVwrpt67sGz74uHkmXxQkwGg0vwzF0-JhcXMbo/edit?usp=share_link

Attributions

Much of this lesson was derived from the following resources:

- Touretzky, D. (2022). Waveforms and Spectrograms: Activity Guide. *AI4K12.org*. Retrieved from https://ai4k12.org/wp-content/uploads/2022/11/Waveforms-and-Spectrograms-Activity-Guide.pdf
- Touretzky, D. (2022). SpeechDemo: Activity guide. *AI4K12.org*. Retrieved from https://ai4k12.org/wp-content/uploads/2022/11/SpeechDemo-Activity-Guide-4.pdf
- Gardner-McCune, C. & Touretzky, D. (2022). Intelligent Assistant with Keywords (Cognimates). AI4K12.org. Retrieved from https://ai4k12.org/wp-content/uploads/2022/12/Intelligent-Assistant-with-Keywords-Cognimates-Activity-Guide.pdf

Lesson Orientation

Warm-Up Exercise Ask students how they make sounds. Have students try to identify as many sound organs as they can. These include mouth, teeth, tongue, larynx, diaphragm, lungs, and others. Watch this short video: #NIDCD: How Does the Human Body Produce Voice and Speech?	5 minutes
Presentation: Spectrogram Activity Sound is a rapid variation in air or water pressure. There is no sound in outer space because there is no air, but there is sound underwater. A microphone converts pressure variations into voltage variations that a computer can detect. The computer measures the voltage thousands of times each second. If we plot the voltage values over time, we can see the "waveform": the pattern of pressure changes that we perceive as sound. An oscilloscope is a device for visualizing waveforms. Rapid changes back and forth (i.e., in pressure or voltage) are called "oscillations", so an oscilloscope is a device for seeing oscillations. "Frequency" refers to how quickly the air pressure is oscillating. Frequency determines the pitch of the sound. Fast oscillations (high frequencies) result in higher pitch sounds. Here is what a sound might look like on an oscilloscope: 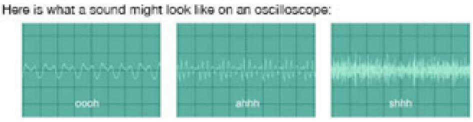 A spectrogram is another way to look at the audio signal. As with the oscilloscope display, the x-axis denotes time. But in a spectrogram, the y-axis indicates pitch (frequency) rather than amplitude. Since the speech signal can have contributions from multiple frequencies, instead of plotting a single line like a waveform, we use color to show the amount of energy in each frequency band. High energy, shown as yellow or red in the display, means that this frequency contributes significantly to the mixture. Here is the waveform for high C (one octave above middle C) along with the corresponding spectrogram:	10 minutes

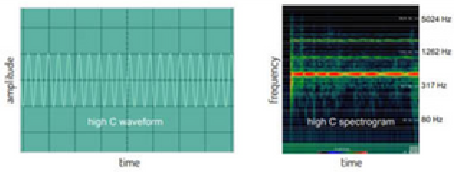

Let's practice now with a spectrogram:
1. Open https://spectrogram.sciencemusic.org/.
2. Click on the microphone icon twice to begin recording.
3. Speak into your laptop's microphone and observe the spectrogram.
4. Press the space key to freeze the display so you can examine the spectrogram more closely. Press the space key again to unfreeze it when you're ready to collect new audio.
5. The spectrogram below shows a person singing a scale middle C-E-G-high C, or do-mi-so-do. Try singing some scales and see what that looks like.

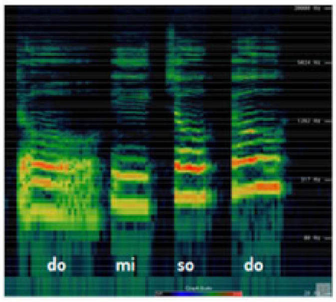

6. **Formants.** Formants are the major frequency components (high energy bands) in the speech signal. They are labeled, F1, F2, F3, etc., with F1 being the lowest frequency. Formants move up and down with intonation (voice pitch). Early speech recognition systems paid a lot of attention to formants, using algorithms called "formant trackers". Modern systems based on neural networks take a more holistic approach and look at the entire spectrogram. The spectrogram below shows the formants when a person is saying "the goat wrote," with "goat" having a rising intonation and "wrote" a falling

intonation. The white boxes enclosing hiss-like sounds (no formants visible) denote the puff of air released after the "t" in "goat" and "wrote."

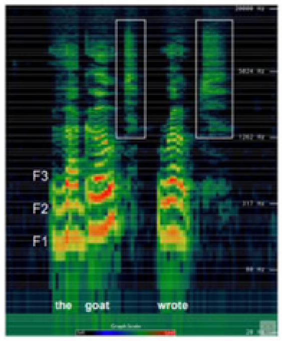

7. Try making your own spectrograms of different vowel sounds and observing how they differ. Vowels are easier to distinguish than consonants because they can be held like a sung note, whereas consonants are brief events whose distinguishing features can be quite subtle. For example, the spectrograms for "pat" and "cat" or "do" and "goo" look nearly identical. This is one of the things that makes speech recognition difficult.
8. Analyzing speech. Taking everything we've learned about spectrograms we can look at a segment of human speech and tease apart some of the structure. Here is a spectrogram of someone slowly saying the phrase "artificial intelligence."

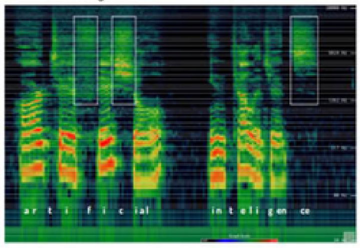

9. Try having different people around you say "artificial intelligence" slowly. What do you notice about the renderings in the spectrogram? Note how that overall, they are very similar regardless of who is speaking. Given how similar one person's use of language may be to another, computers can then begin to use artificial intelligence to come up with informed determinations of what a person says based on what that word or string of words sounds like when said by most humans.

Practice: Speech Demo — 15 minutes
1. Ask students to navigate to http://www.cs.cmu.edu/~dst/SpeechDemo
2. Speak into the microphone and your voice will be transcribed to text
3. If there are several plausible transcriptions, all leading candidates will be displayed.
4. Click on the "pause" button to temporarily disable the microphone.
5. Click on "resume" to resume recording.
6. Click on the "read back" button to turn on computer readback of the recognized text.
7. Click on the pulldown menu to switch to a different language model (e.g., Spanish, Chinese, etc.)

Students should now practice with the Speech Demo. Explain that computers use data taken from the microphone and a digital rendering of the person's speech to transcribe what the user is saying. At times, this is very accurate, especially if the person speaks clearly, speaks slowly, and does not have a strong accent.

Ask students what ethical concerns might arise from speech recognition experiencing challenges in rendering speech from those with accents. (Answer: Speakers with native accents or accents adjacent to those who programmed the artificial intelligence are prioritized, while second language learners may be left out of the development process and the subsequent benefits AI can bring.)

Offer students some insights into how AI speech recognition can work better for them:

1. Importance of context: Longer sentences work better. Accuracy for single words may not be that high, but if the words are strung together in a sentence, contextual information can be used to make better guesses about each word.
   Try the following short words or phrases:
   a. • "brat"
   b. • "scat"
   c. • "sat scat spat"
   d. • "fat flat flap flag flab"
   Now try these full sentences:
   a. • "My little brother is a brat."
   b. • "I like scat singing."
   c. • "Lie flat and let the flag flap in the breeze."

2. Careful enunciation helps. Compare the following:
   a. • "heating the in a up sandwich"
   b. • "heating. the. in. a. up. sandwich." (pause between each word)
3. Homophones can often be disambiguated by context. Try these examples:
   a. • "They're claiming their dog is there."
   b. • "Which witch is which?"
   c. • "If I teach from 12 to 2 can I eat lunch too?"
   d. • "He led the study of lead poisoning."

However, there are limitations to speech recognition. Speech recognition was not all trained on human speech. So, common sayings or phrases may take precedent over what a person is actually saying.

1. Try these examples:
   a. Here is famous quotation, from The Hobbit, that uses "ere" Does Google get it?
      • "We must away ere break of day."
      What if we change a few words?
      • "You must depart ere noon today."
   b. Sing the chorus to "Three blind mice." What if we change some words?
      • "Three blind mice. Three blind mice."
      • "Three blind mites. Three blind mites."
      • "See blind mice. See blind mice."
   c. This famous quote is from John Donne ("Meditation XVII"). What if we change the last word?
      • "No man is an island."
      • "No man is an eyelid."
      • "Now blink your eyelid."
      Google has no trouble recognizing "eyelid", except when it occurs in a context where it strongly expects to hear "island".

What do these experiments teach us?
- A language model determines not only the vocabulary that can be recognized, but also the set of phonemes (sounds) the recognizer will consider. Changing language models can have a dramatic effect on how English is "heard."
- AI is limited in its abilities right now. Given not every intelligent agent system has been trained on every model possible, errors will inevitably occur.
- AI models must be expansive and include a variety of spoken dialects of a language, including those spoken that are not considered "standard." In this way, AI can represent and benefit more people.
- Models need to be continually updated in order to reflect the most recent changes to language. Can you think of sayings that are used today that may not be recognized by AI?

Production: Cognimates

20 minutes

1. Here is a direct link to the Scratch3 page at Cognimates.me. Note that you must use this version of Scratch, not the one at scratch.mit.edu: https://codelab.cognimates.me/. You can also get to this page by going to the Cognimates.me home page and clicking on the "Code & Play" button.
2. At the bottom of the palette, you will find entries for the Speech to Text and Text to Speech extensions, which are pre-loaded. These give you the additional blocks shown below.

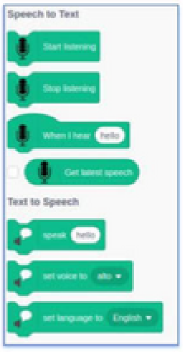

The intelligent assistant works by accepting input from the user, scanning it for keywords, and generating a canned response when a keyword matches.

The assistant understands several types of requests; we will write a separate block to handle each type. All block definitions appear at the end of this section; refer to those pages when you need to create a block in the steps below.

Go to "Variables" and create a variable called user_response. This is what you should see:

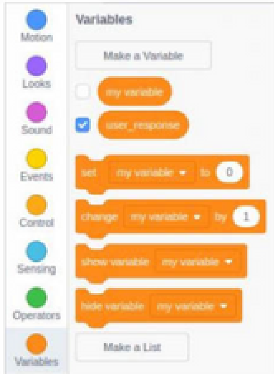

2. Go to "My Blocks" and click on the "Make Block" button. The feature we'll add allows the assistant to ask us riddles. In order to use the "riddle" block, you must create two list variables. Create a list variable called riddle_questions and set it to this list:
   - What has hands but cannot clap?
   - What has eyes but cannot see?
   - What has legs but doesn't walk?
   - What can you catch but not throw?

   The result should look like this:

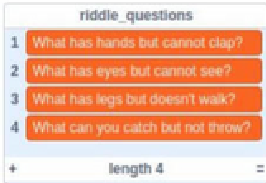

3. Create another list variable called riddle_answers and set it to this list:
   - clock
   - potato
   - table
   - cold

   The result should look like this:

4. Uncheck the boxes for riddle_questions and riddle_answers so they don't appear on the stage.
5. Create two new variables called "riddle_question" and "riddle_answer" and uncheck their boxes. Then create a new block named "riddle". Enter the definition for the riddle block using the code provided at the end of this section. Notice that riddle makes use of two-step.
6. Update the main program by inserting a "riddle" block after the "two-step" block. Now your program should exactly match the code shown in this activity guide.
7. Test your program by clicking on the green flag and saying "Assistant, ask me a riddle." Try giving the right answer. Then ask for another riddle and try giving the wrong answer.

Ask students to share what they discover. If time allows, ask students to create a new set of riddles and modify the existing code.

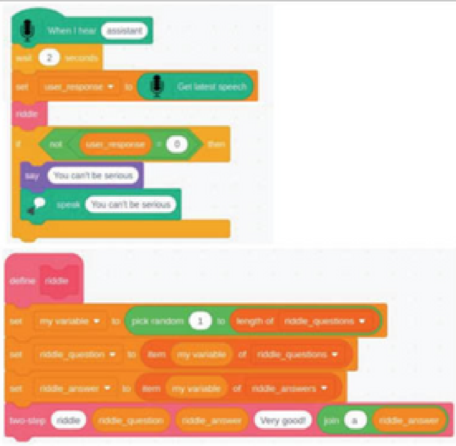

Extension

Give students the link for the Cognimates lesson link. Students can continue the lesson to create riddles and two-step processes:
https://ai4k12.org/wp-content/uploads/2022/12/Intelligent-Assistant-with-Keywords-Cognimates-Activity-Guide.pdf

Additionally, students can complete the extension activity found in the lesson plan: Have students modify the code to create their own keywords and responses. For more advanced students, have them implement a "quiz" option modeled after the "riddle" block; it should pick a quiz question at random from a list the student creates. For example, it could be a trivia quiz about their school, or a quiz about their state (e.g., state motto, state bird, state flower, state capital). Students could also modify the code to make the sprite change its appearance while it's responding to a request, or while it's waiting for them to supply an answer to a joke or riddle. (Gardner & Touretzsky, 2022, p. 8)

Students could also modify the code to make the sprite change its appearance while it's responding to a request, or while it's waiting for them to supply an answer to a joke or riddle. (Gardner & Touretzsky, 2022, p. 8)

IN CLOSING...

Though work remains in the design and implementation of our collaboratively designed and fully integrated curriculum, ReadyAI is incredibly proud of the work that was produced in partnership with NCS throughout the 2022-2023 academic year. From cutting edge lesson plans on emerging AI technologies to helping execute the first ever Hackathon at NCS, our joint efforts have created a truly unique and innovative map for AI learnings for young students. As this technology continues to rapidly evolve, it is ReadyAI's commitment to assist in the design, development, implementation, and evaluation of a robust learning ecosystem at NCS, not only for students but for the talented faculty and staff as well. In seeking together to adapt and build upon the lessons, experiences, and opportunities of this past year, we look forward to doing our part to bring about the successes of both NCS and NEOM project as a whole.

Thank you for your time, invaluable expertise, and continued support.

Roozbeh Aliabadi
Rooz Aliabadi, Ph.D.
CEO ReadyAI

Jordan Mroziak
Jordan Mroziak, Ed.D.
CTO ReadyAI