



KEEPING PERSPECTIVE:


Focusing on HI (Human Intelligence) When AI Is Everywhere

Tim Hudson, Ph.D.

Chief Learning Officer, Discovery Education

Table of Contents:

Introduction	3
Real-World AI Examples	4
Applying Human Skills to Technology	7
How AI Works and “Thinks”	8
How We Can Prioritize Human Intelligence	9
Designing Lessons and Assessments for Humans	10



Introduction

As Artificial Intelligence (AI) makes an impact across every industry, educators are working to navigate the complexities of this new, powerful technology. In a survey conducted during an AI webinar in fall 2023, 78% of educators responded that they were simultaneously excited and concerned about AI in education. This curious, but cautious, attitude toward AI is a clear indication that educators hope to make use of new technologies responsibly but don't assume it will be simple. They recognize that it's critical to incorporate AI in ways that enhance learning without compromising the value of human interactions, pedagogy, and ethical behavior.

In this eBook, we'll share background about AI, positive and negative examples for AI uses, and discuss what it might mean if we say that AI "thinks." We'll also outline strategies to prioritize Human Intelligence by empowering students as critical thinkers, problem solvers, and compassionate individuals who responsibly use AI as a tool in their journey.

Real-World Examples for How AI Use Can be Helpful, Limiting, and Harmful

We have a wealth of AI-powered tools that can offer unprecedented opportunities to enhance everything from daily tasks, to work, and to education. The following examples demonstrate how AI can improve our lives, while also highlighting why human intervention and review of AI's inputs and outputs remains essential.

Examples of Good AI Uses

Using Chatbots for time-consuming tasks:

Many people have found generative AI helpful for activities like reformatting citations, creating prompts for thought starters, developing slides for a presentation, or drafting a lesson plan. When used in this manner, AI helps users save time on mundane tasks or generates an initial starting point from which to build on, thereby enables them to refocus efforts on projects in which AI may not be as useful.

Using AI for guidance on a project:

AI programs, such as Pi (personal assistance AI), can be a great resource for helping solve problems, access information quickly, or ask clarifying questions. These types of bots often adopt a conversational and encouraging tone to help the user feel motivated and comfortable exploring more.


Using AI to generate captions for images:

AI image recognition has advanced rapidly, which means a user can provide an image and ask AI to describe it in three, six, and nine words. The most cutting-edge AI can complete this task in seconds with attention to detail, thereby providing simple alt-text for better accessibility for people with visual impairments.

THREE WORDS: Seaside dining table

SIX WORDS: Dining table overlooking ocean with beer

NINE WORDS: Outdoor restaurant table with ocean view and a beer




Examples of AI Limitations

AI doesn't understand purpose and intent:

When a user gave a popular chatbot the Taco Bell menu with each item's calorie count and asked it to create a balanced, 1000-calorie meal suggestion, the chatbot suggested a 920-calorie meal that consisted mostly of beverages and hot sauce packets. While the calorie count was accurate, AI couldn't provide a reasonable meal, let alone one that's balanced and tasty.

One of the most popular early uses of generative AI has been to create images and illustrations based on simple or descriptive prompts. For example, AI can quickly generate an image of "a border collie wearing a suit on a motorcycle." Some users have leveraged this capability to extrapolate what might exist beyond the borders or existing photographs or works of art, such as the Mona Lisa and the Beatles' Abbey Road album cover. Creating a new image or the extrapolation of extra graphics using original art is impressive for a machine to accomplish. However, it's important to note that AI isn't giving us "the rest of the picture" because Da Vinci and the Beatles' art was exactly how they intended it to be.



AI can rapidly create harmful disinformation and misinformation:

The AI-generated products in these examples aren't accurate, realistic, or necessarily harmful, but the technology itself creates new challenges and risks to our society because of how rapidly and realistically fake images can be created and spread. For example, AI can generate images of two world leaders shaking hands, hugging, or fighting, even if those two leaders have never met in person. With deepfake technology and AI-generated video like Sora recently demonstrated, people with an intent to harm could create entire conversations or speeches generated that will create confusion and conflict. We must take steps to create safeguards that ensure these AI tools aren't able to put people, communities, or countries in danger because of intentional disinformation.

Problematic AI Behaviors

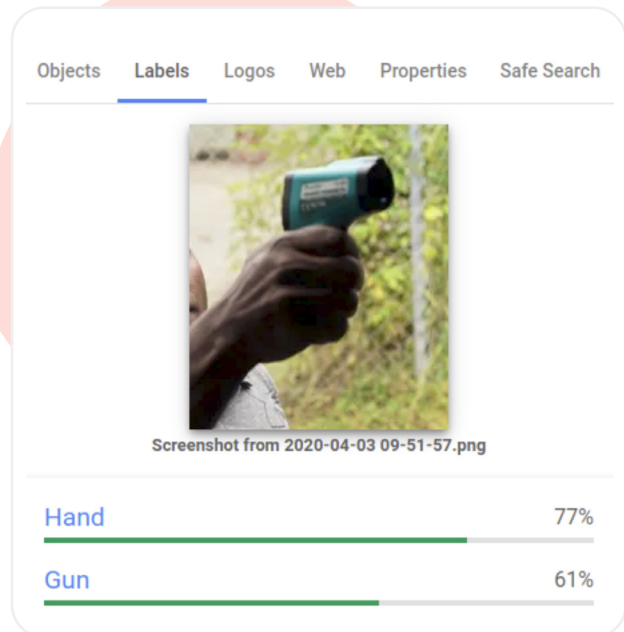
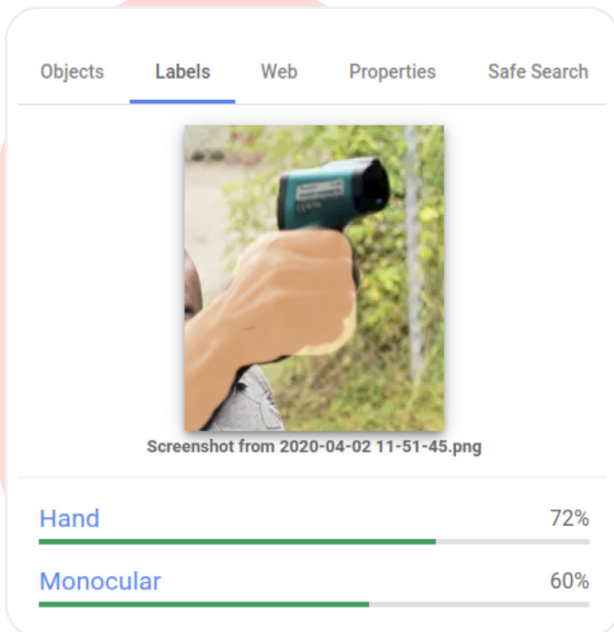
AI can generate harmful behavior learned from users:

In 2016, Microsoft launched a Twitter-based AI chatbot named Tay that used the handle @TayTweets. Tay was intended to have a playful, youthful personality. However, within hours of its launch, Tay started producing offensive and inappropriate tweets, including racist, sexist, and inflammatory content. It learned these behaviors from interacting with users who deliberately tried to provoke the bot and teach it offensive language and viewpoints.

AI adopts biases from its data inputs:

Tay is one example of how AI is not impartial and can assimilate the biases of its inputs. Another example of bias in AI was in a photo recognition program developed by Google's Vision AI. In an experiment by AlgorithmWatch, Google's Vision AI labeled an image of a dark-skinned individual holding a thermometer as a "gun", while a similar image featuring a light-skinned individual was labeled as a "Monocular." Systems like Google Vision Cloud process millions of pictures to detect patterns. However, these pictures are labeled by humans, which can introduce bias, whether they are aware of it or not.

These incidents highlighted the risks of developing AI systems that consume un-curated input. Like humans, AI learns from the inputs it's provided. Therefore, it's essential to carefully control, understand, and make transparent the inputs of AI as we also monitor its outputs. Biased outputs reflect biased inputs. We must be aware of the implications and impact of biased AI outputs and take steps to reduce harmful stereotype bias or prejudice.



Applying Human Skills to Technology

For AI, like all technology, a central purpose we should have is creating opportunities to help humans perform mundane or routine tasks faster, with more accuracy, and at scale. The time freed up can be spent on more complex tasks and relationships, for example. One great example of technology freeing up human time is the hand-held calculator. We input information into the machine, the calculator follows a set of coded logic and instructions, and then quickly produces a computational result. The calculator cuts down a manual process and frees up time for users to focus on bigger ideas.

When calculators first came into existence, many math teachers were concerned that students would rely too much on the technology, or that students would focus too much on the answer instead of on understanding key mathematical concepts. Some worried that students' mental math skills and number sense would suffer because of over-reliance on calculators. In reality, the existence of calculators creates even more of a need for students to understand concepts and have strong number sense.

To effectively use a calculator, students must know the operations and values to enter into the calculator and they must be able to recognize if the calculator produces a reasonable output. The tool can only generate correct computational results if students have applied the critical thinking and analysis they would've needed even before the calculator became available. The calculator frees up class time to focus on ideas instead of mundane computation.

Like the calculator, humans must provide the right inputs, apply the same critical thinking, and carefully analyze outputs generated by AI. One key difference between calculators and AI is that a calculator is programmed using number systems with well-defined, consistent rules, while AI uses large language models, and language is inherently far more complex and less logical than the logic of number systems. Calculators don't produce correct results because they were "trained using billions of computations." Which means you can trust the result of a calculator as being logical; calculators don't "hallucinate" like AI sometimes does.



How AI Works and “Thinks”

A Quick Definition of Large Language Models

- 1** ChatGPT and other generative AI software use large language models (LLM). These models work by processing large amounts of text data to learn patterns between words and phrases.
- 2** LLMs use word vectors, which use relational models to organize words with similar meanings while capturing the context, semantics, and associations between words in the text data on which they are trained.
- 3** The models use this knowledge to generate human-like text based on input prompts by predicting the next word in a sequence and generating coherent and contextually relevant responses.

Does AI mean that machines can think?

Machines can perform tasks like processing data, understanding language, and recognizing patterns; and do these complicated tasks far more quickly and expansively than humans, but without feeling or awareness. While it may seem like AI is “thinking,” these machines execute these tasks through programmatic models and algorithms that humans have developed to appear functionally equivalent to human speech.

As a result, the accuracy of technology like AI can hinge on the accuracy of the models, algorithms, texts, and instructions humans have provided. These tools are merely regurgitating what humans have said or written, and they are limited to knowledge of the texts they’re trained on.

This is not dissimilar to human experience; our knowledge is often shaped by the books we’ve read, the places we’ve lived, our life experiences, and the news we’ve consumed. Like AI, what we have access to and the experiences we have informs what we know and how we grow. Some key differences are that AI has access to vastly more information than an individual human ever will, and with the ability to process all of it. But humans have actual experiences, interests, feelings, and curiosity that AI doesn’t have.

How Educators Can Prioritize Human Intelligence

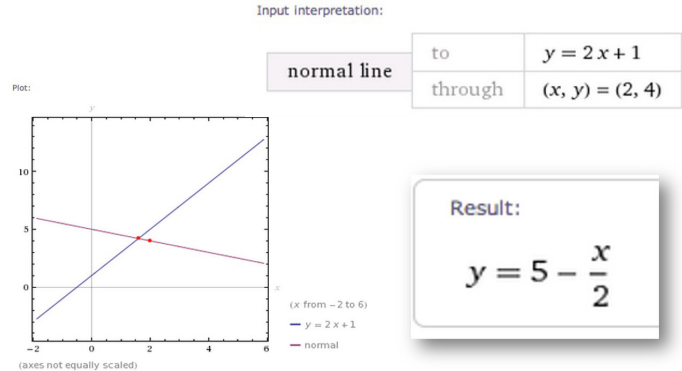
Just as we reflect and scrutinize the information and data used for training AI programs, we must apply that same thought process to our curricular decisions and how we’re educating young people. What opportunities are we giving students? What are we giving them to think with? What contexts and experiences should they have to inform their “outputs”? What are the outcomes we seek from education?

As noted earlier, students will encounter deep fakes, false information, and mal-information (which is intended purely for harm) in this new age of AI, which means it’s even more important to ensure students can think critically about information and ideas they encounter. Regardless of whether students are analyzing an AI-written article, or an article written by a human, they need to have the skills to assess the validity of that content and respond appropriately.

Making Lessons and Assessments for Humans

Before the generative AI technology advancements of 2023, math teachers had been operating with advanced math technologies for decades. For over ten years, programs like Wolfram Alpha have been able to quickly and easily solve text-based math problems such as:

What is the equation of the line perpendicular to $y=2x + 1$ through $(2,4)$ ”?



Wolfram Alpha’s answers were outstanding, and teachers agreed that if they were to grade this as a test, Wolfram Alpha deserved an A. That helped math teachers realize it was time to rethink some things about what we assess. If Wolfram Alpha can get an A on the test, then we need to make a better test. Today, we have been able to apply this same approach to ChatGPT—if ChatGPT can get an A on the test, we need to make a better test. We can do this simply by engaging students in tasks that AI can’t do or write about.

Here are three principles and examples for making AI-proof lessons and assessments that center students’ humanity:

- 1** AI isn’t capable of authentic curiosity, so invite and assess students’ questions, not just their answers.
 - Ask students to generate questions they’re wondering about a character, situation, or event, including why they’re wondering that, thereby leveraging the curiosity that makes us human.
- 2** AI is limited by its inputs, so invite students to respond to situations AI wouldn’t be aware of.
 - Ask students to describe what they notice and how they would respond to a unique situation that’s not broadly known, thus engaging and honoring their intuition, ideas, and observations.
- 3** AI doesn’t know about each student’s life experiences, so invite and require personal reflection.
 - Ask students to compare and contrast a character or situation to someone or something they’ve experienced, requiring unique details about personal connections and their community.

By better understanding generative AI technology, educators can have a more informed perspective about how it might guide learning, complete tasks, or explore ideas. The goal should always be to focus on areas that improve human development, connection, and understanding. We can use AI to better prioritize time, and help students learn things they can’t ever get from a chatbot.

For more discussion, you can watch a recorded version with these thoughts and more [here](#).

About the Author



Dr. Tim Hudson serves as Chief Learning Officer at Discovery Education, where he supports partner districts and internal teams as they develop and implement research-based, innovative, and effective learning resources for teachers and students. Prior to joining Discovery Education, Tim spent 12 years with DreamBox Learning (now a part of Discovery Education) where he oversaw the creation of digital, personalized curricula and assessments in math and literacy that engage students in critical thinking while also providing teachers and administrators with insights to support differentiation and improve student achievement. Before joining DreamBox, Tim spent over 10 years in public education, first as a high school math teacher, and then as a K–12 district mathematics coordinator and strategic planning facilitator in suburban St. Louis, MO. As an education leader, Tim writes and speaks about innovations in pedagogy, school systems, personalization, assessment, and technology.

About Discovery Education

Discovery Education is the worldwide edtech leader whose state-of-the-art digital platform supports learning wherever it takes place. Through its award-winning multimedia content, instructional supports, innovative classroom tools, and corporate partnerships, Discovery Education helps educators deliver equitable learning experiences engaging all students and supporting higher academic achievement on a global scale. Discovery Education serves approximately 4.5 million educators and 45 million students worldwide, and its resources are accessed in over 100 countries and territories. Inspired by the global media company Warner Bros. Discovery, Inc. Discovery Education partners with districts, states, and trusted organizations to empower teachers with leading edtech solutions that support the success of all learners. Explore the future of education at www.discoveryeducation.com.