A STEAM Resource for Educators in a Digital World | FALL 2021 • ISSUE 4

ROOT & STEM

AI: Coding Consciousness

DIVING IN DATA

How AI can save the whales

ARTIFICIAL CREATIVITY

AI and the future of cultural jobs

METAL AND MINDS

Responsible ethics in AI

+ Plus comics, games and over 10 pages of classroom-ready teaching resources

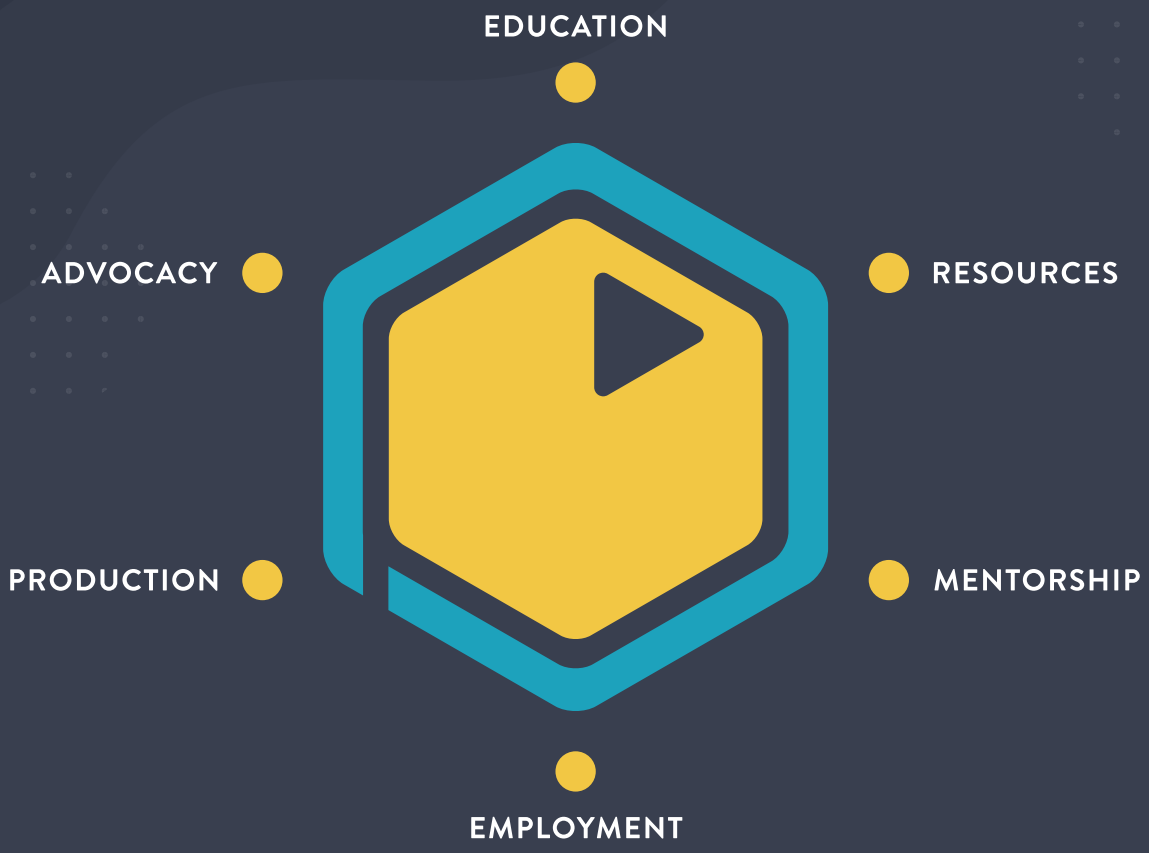

PINNGUAQ
LIFE CYCLE

Pinnguaq follows a life cycle model to support the core phases of a person's learning journey in STEAM education. We strive to provide educators and students with opportunities and resources each step of the way.

To learn more about what we do, visit our website at

pinnguaq.com

Contents

ROOT & STEM

ABOUT PINNGUAQ

The Pinnguaq Association, a not-for-profit organization, incorporates STEAM into unique learning applications that promote storytelling, health, wellness and growth in rural and remote communities. At its core, Pinnguaq embraces diversity and creates opportunities in order to empower all people.

DIGITAL TAXONOMY

Computer Science Education is more than just coding. A comprehensive approach to it includes learning skills and competencies from each of the areas listed below. Look for the following icons at the end of each article for suggested curriculum connections. *Reference: Learning for the Digital World: A Pan-Canadian K-12 Computer Science Education Framework. 2020.* ***k12csframework.ca***


CODING AND PROGRAMMING


COMPUTING AND NETWORKS


DATA


TECHNOLOGY AND SOCIETY


DESIGN

The Pinnguaq Association acknowledges the support of the Government of Canada in developing this educator resource.

Root & STEM is printed on certified Forest Stewardship Council paper, to reduce our impact on the natural world.

Canada

Root & STEM (ISSN 2563-6979) is produced twice yearly by the **Pinnguaq Association**
87 Adelaide St N, Lindsay, Ontario K9V 4L4

• • •

Printed in Oshawa, Ontario, by Maracle Inc.

• • •

Publisher
Ryan Oliver

Director of Education
Jennie Cross

Copy Editor
Nicholas Davies

Design & Layout
Emily Canfield, Lindsay Hill

Cover Illustration
Micaela Dawn

Logo Design
X-ing Design

Contributors
Sharon Aschaiek, Briana Brownell, Nia Emmanuel-Briggs, Jessa Gamble, Jeremy Gilmer, V. M. Hetherington, Lindsay Hill, Sofia Osborne, Sophie Penner, Greg Picken, Cindy Rose, Caroline Sinders, Alex Smithers

Kawartha Lakes CFDC

As an artist and researcher, I find artificial intelligence (AI) deeply fascinating. After all, AI is all around us: in search engines, social networks and even Netflix. There's a conception that technology, including AI, is neutral—but we've learned that's anything but true. Technology is more like a mirror held up to society: along with our progress and successes, it reflects our biases and systemic harms back to us. I believe it's important for students, especially those who are young, to understand what AI is and how it can reinforce such biases, as well as how we can use it for inspiration (and for art!). This belief drove my collaborator, Bojana Kostic, and I to create a website to help explain what AI is, what we should know about it, what we should be wary of, and how we can use it for creative pursuits. You can learn more at **ai-explanations.com**. But here's a brief primer:

What is AI?

AI on its own really isn't anything: just lines of code. When it comes to how that code is used in society and how those uses impact people, however—that's when AI becomes *more* than just code. In *The People's Guide to AI* (written by the fantastic Mimi Ọnụọha and Diana J. Nucera), AI is described as "like salt," in that when it is added to a product, that product is transformed.

AI can be employed to help software and other tech products accomplish a wide variety of tasks, from sorting or locating information in a search engine to recommending music on Spotify. Forms of AI exist in public spaces such as schools, hospitals and workplaces, as well as in the privacy of our homes.

To function, AI needs data. Data allows AI systems to build patterns, which they then use to generate predictions and simulations, or to infer information about the world. Often, this data is gathered, extracted and mined as a focussed process to build understanding of communities and of how we use technology.

What are the risks of how we use AI?

Since it reflects the biases we build into our technology, AI often makes mistakes and assumptions that are similar to those that people make. For example, in 2013, Harvard University professor Latanya Sweeney was Googling her own name and noticed that her search results prompted ads that read "Latanya Sweeney, arrested?" and "Latanya Sweeney, bail bonds?" Professor Sweeney tried other names that are often associated with Black and African American communities, and the same kind of prompts showed up. Professor Sweeney has never been arrested. This AI prediction is an example of a poorly programmed code that reflects an all-too-common societal bias that Black people are more likely to be incarcerated or in trouble with the law than other communities.

Another example of AI bias was discovered in 2018 by researchers Joy Buolamwini and Dr. Timnit Gebru, in this case related to computer vision and gender. Buolamwini and Gebru discovered that computer vision algorithms, used across major companies like Google, IBM Watson and Amazon, had a difficult time recognizing skin colours and genders. They found that the software was 99 per cent correct when identifying white, masculine-presenting faces, and only 35 per cent correct for darker-skinned, feminine-presenting faces. What's important to realize is that Gebru and Buoalmwini were testing products that are already commercially released, which raises questions about how AI is trained and evaluated to recognize the diversity of people with a high degree of accuracy.

How can AI inspire us?

Of course, not all uses of AI are problematic or harmful. For example, it is being used to create art in interesting ways. The writer and artist K. Allado-McDowell co-wrote a beautiful book with GPT3, a predictive text system created by Open AI, on the nature of human existence. Artist Stephanie Dinkins has created a sculptural chatbot recounting the history of her family and African Americans in the United States. Anna Ridler, a UK-based artist, creates large-scale, handmade datasets and generates AI videos from there. For example, she photographed more than 3,000 tulips and generated an AI video of the tulips moving, rising and falling based on the price of Bitcoin (a nod to the Tulip Crisis of the 1400s and the concept of speculative worth). Just as the invention of the photographic camera created the medium of photography, AI can help artists create new ways of seeing. With any new innovation or technology, it is important to be skeptical and view it with a critical eye—but also to leave space to consider new possibilities. AI is no different—it can amplify biases in harmful ways, but it can also create beauty. It's this duality that makes artificial intelligence so important and so complex. As a practitioner, I want us to understand and centre this duality—and to ask: What can we create with AI and how can we also use it to reduce harm in our world?

In this issue of *Root & STEM*, you'll find articles and resources that explore these questions. Whether it's how AI is used in industry, how it has changed learning and creativity or how it is helping scientists understand and catalogue whale calls, this issue asks what kind of world AI is creating, and how human life is changing as a result.

— CAROLINE SINDERS

Micaela Dawn

Cover Illustration

Micaela Dawn is an award-winning freelance illustrator and concept artist. In her spare time, she has been developing projects that speak to her desire for equality and the promotion of women's rights. Her main focus is creating bold, colourful works of art that have roots in fantasy and mythology.

Sharon Aschaiek

Amazon Future Engineer Canada • *Page 9*

Sharon Aschaiek of Higher Ed Communications is a writer, editor, content consultant and project manager who lives just north of Toronto. She writes about education in Canada and globally for media publications and academic institutions.

Jessa Gamble

Artificial Creativity • *Page 18*

Jessa Gamble is an internationally award-winning science and technology journalist and Penguin Random House author. Her writing has appeared in *The Atlantic*, *Nature*, *New Scientist* and *The Guardian*.

Cindy Rose

Deep Blue • *Page 26*

Cindy Rose is a digital artist living in Toronto. She received her advanced diploma in Graphic Design from Humber College and specializes in digital painting. She is inspired by elements from nature and is obsessed with capturing the human spirit. Follow @**cindyrosestudio** for more!

Nia Emmanuel-Briggs

AI Bingo • *Page 38*

Nia Emmanuel-Briggs is a graphic designer based in the Toronto area. You can find her at **thedesigniacreative.com**.

Caroline Sinders

Guest Editorial • *Page 5*

Caroline Sinders is a machine-learning researcher and artist obsessed with language, culture and images. Her work has been featured in the Victoria and Albert Museum, MoMA Ps1, the Modern Art Museum of Bologna and Ars Electronica, among other venues. She is the founder of Convocation Design + Research, an agency focusing on the use of machine learning and design for public good.

Briana Brownell

Artificial Creativity • *Page 18*

Briana Brownell is a Canadian data scientist and technology artist. Her technology-enhanced creative projects span multiple areas, including AI-assisted Shakespearean sonnets, and AI-enhanced, assisted and generated visual art pieces.

V. M. Hetherington

Deep Blue • *Page 26*

V. M. Hetherington (she/her) is a Toronto-based author and arts-based disability advocate. Her first novel *Mooncalves* was shortlisted for the Amazon First Novel Award in 2020, and was described as a "stunning debut" (*The Globe and Mail*) and "this generation's response to Atwood" (*Vancouver Sun*).

Sofia Osborne

Diving in Data • *Page 32*

Sofia Osborne is a writer, reporter and audio producer based in Vancouver. Her environmental journalism has appeared in *The Tyee* and *The Narwhal*, and she is the co-host and producer of *Beyond Blathers*, an *Animal Crossing* science podcast.

Metal and Minds

Ethics in AI

BY JEREMY GILMER

The ethics of artificial intelligence (AI) is a subject that is in the news a lot right now, and many articles, TV shows and podcasts are talking about it. Ethics have become one of the greatest challenges in the development of AI, and this has led to some profound and complicated questions when it comes to AI and machine learning. There are many places in which AI and ethics find themselves crossing paths: facial recognition, banking algorithms, weapons of war, even social media.

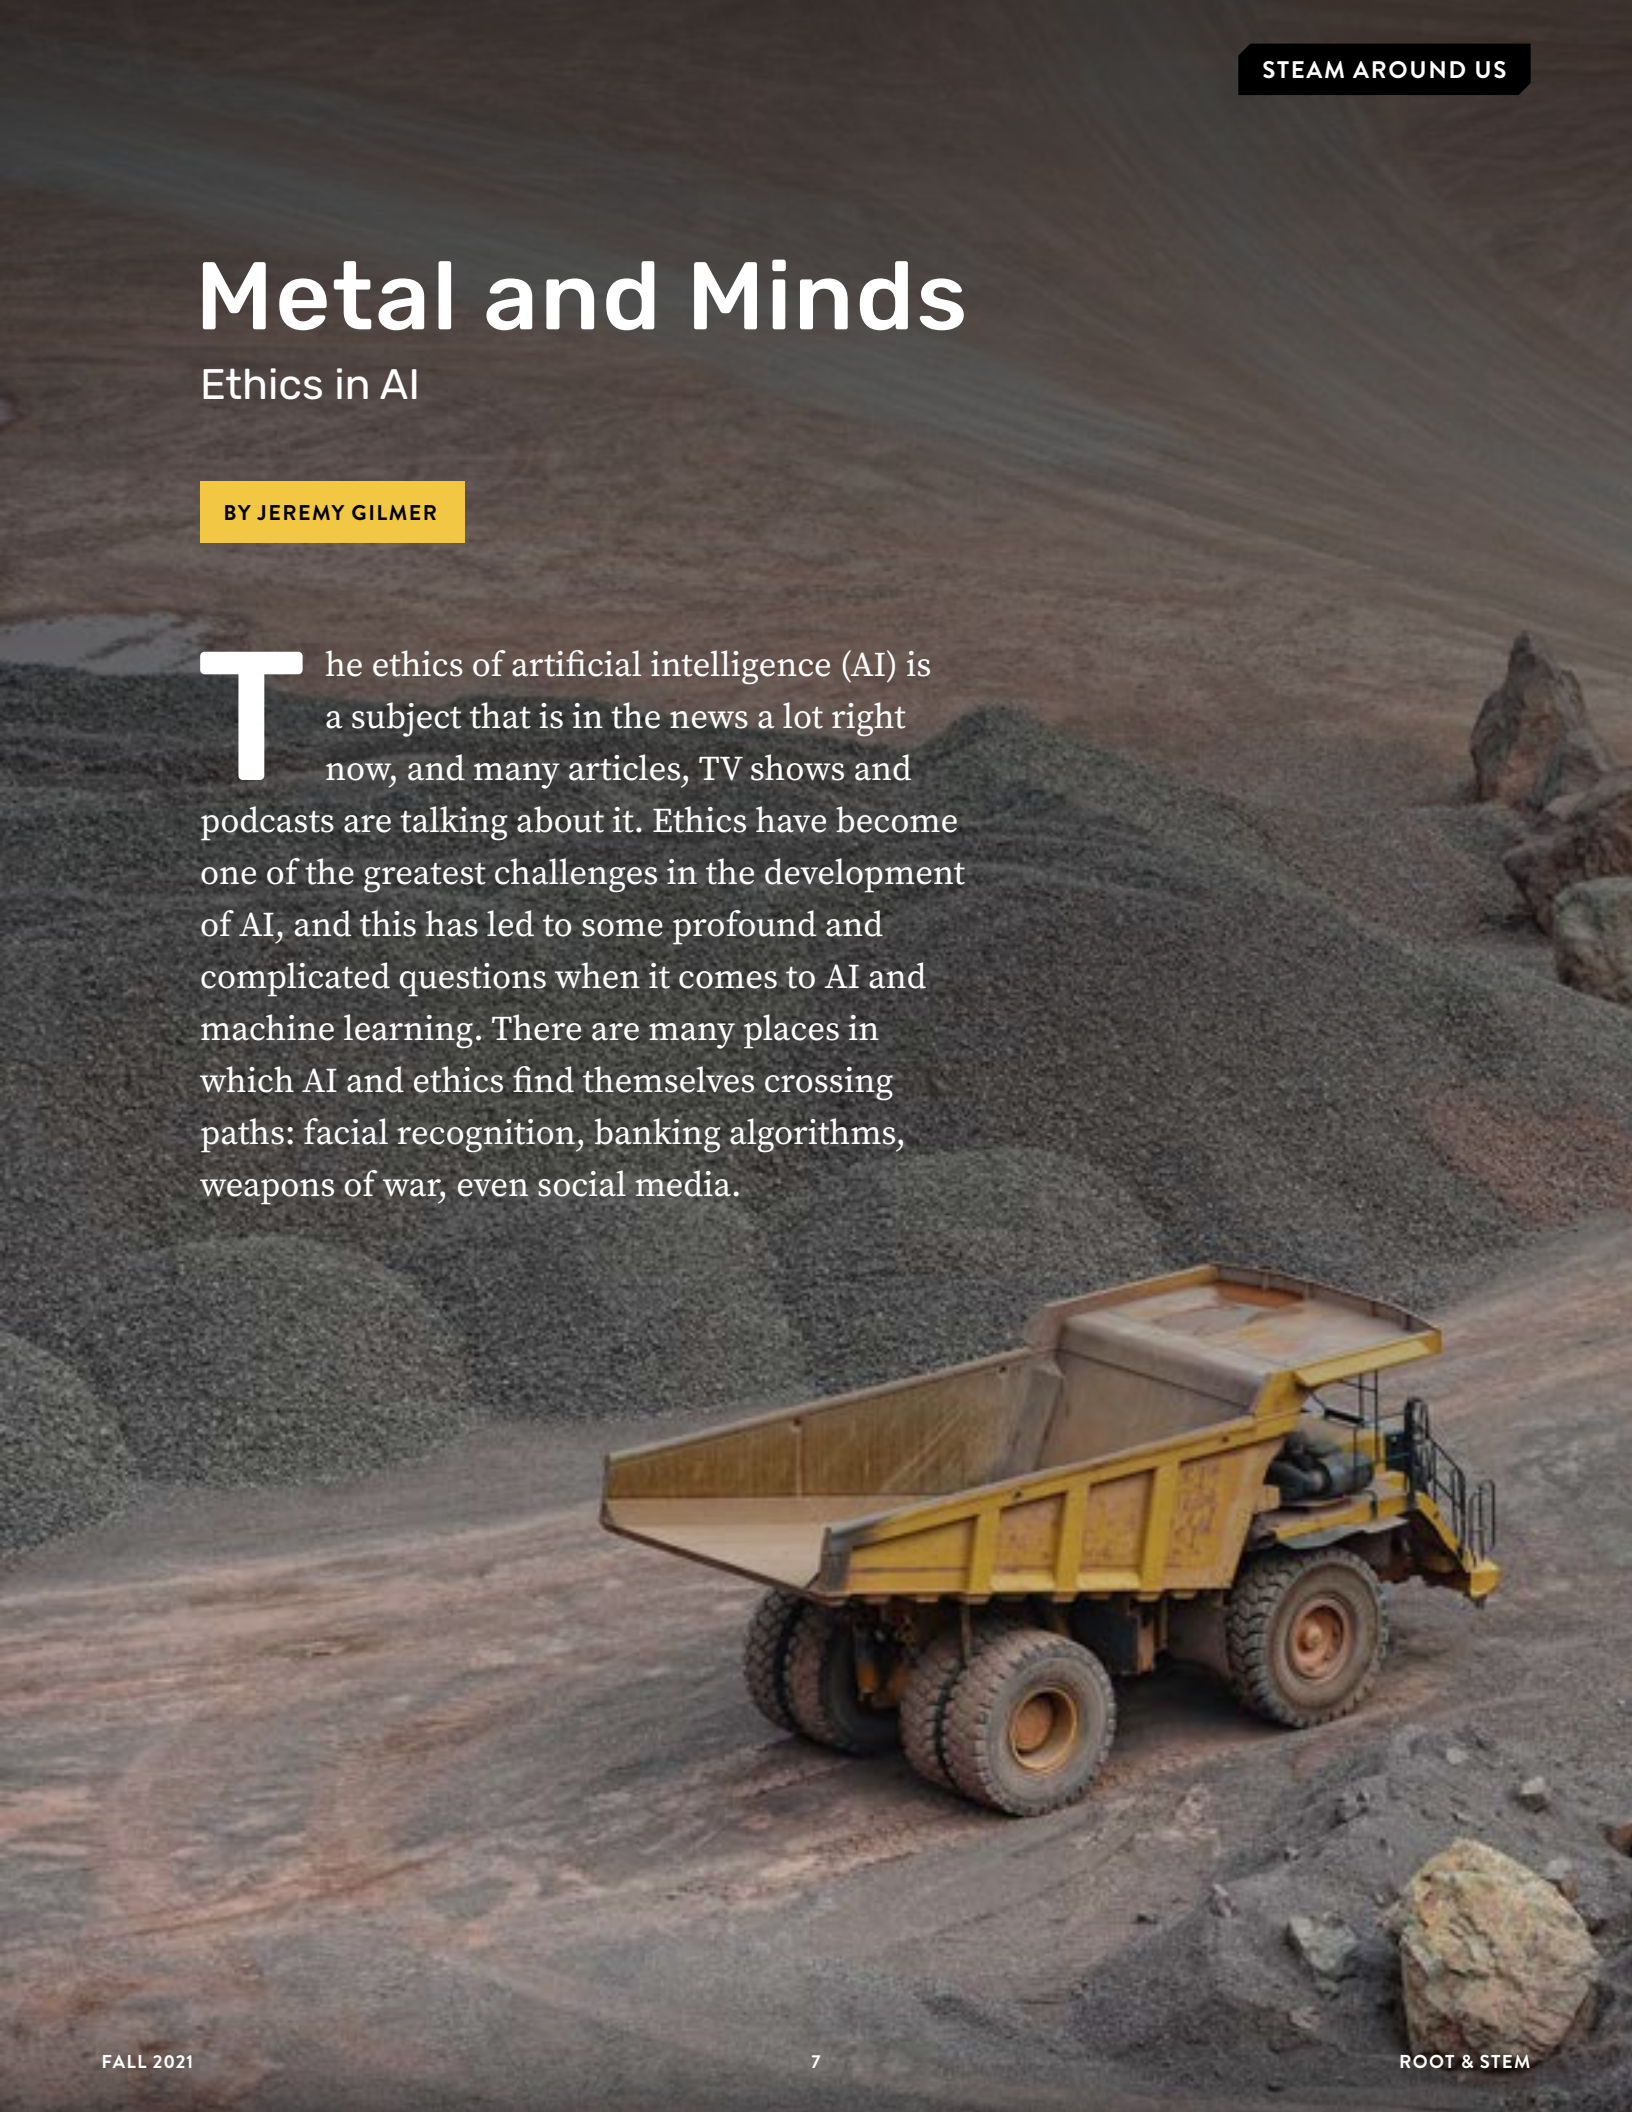

But one sector wherein the design and ethics of AI can literally mean life or death every day is mining—specifically, in the use of large mining haul trucks. Many companies are developing technology to replace human operators with autonomous vehicles, and the stakes are as high as they can get.

A mining haul truck can weigh over 300 tons and stand more than four storeys tall. Furthermore, they can travel at 60 kilometres per hour, which is like a house coming at you as fast as a city bus. It may seem crazy to put a computer in charge of something so big, but this is already happening in some places in the world—and perhaps, instead of being scary, it could actually be a very good thing.

One of the biggest questions is, what if the truck makes mistakes? Won't someone get injured or even killed? The first answer is that humans are not perfect, and, as it stands, accidents already occur with human operators. In 27 years on mining projects, I myself have lost two friends to accidents with large equipment. Once, in the Democratic Republic of the Congo, my little pick-up truck was overtaken at high speed by a fully loaded Cat 793 haul truck, which went on to clip the side of another vehicle in front of us and flip it over like a toy. Thankfully, no one was killed in this accident—but it could have been a catastrophe.

The primary cause of this event was operator error. I have been part of teams investigating mining accidents on four continents, and operator error is overwhelmingly the cause of the majority of these incidents. (The number one reason for operator error in most of these accidents? A driver using a cellphone. Of the last six accidents I helped investigate in which operator error was the cause, cellphones were in use in all of them.)

In Chile, South Africa and several other places, autonomous haul trucks are being tried out as a potential way to address safety issues. Companies like Caterpillar, Komatsu and others are developing technologies and systems to incorporate self-driving vehicles into active mining environments.

But the ethical decisions behind how these systems are programmed remain complicated. A driver can make a split-second decision to reduce harm in the event of, say, brake failure. An AI system can only ever do what it is programmed to do. Taking this into account, if an autonomous truck finds itself barrelling toward a bus full of people, should the truck's AI be programmed to turn into the side of the hill, or to hit the oncoming bus? Well, that's easy. But what if the choice is more complicated? What if it's a choice between a bus coming from one direction and a fuel truck coming from another? What is the best course of action?

This question—and others akin to it—has been studied by philosophers for decades and is generally referred to as "the trolley problem." It is often framed as a series of thought experiments, in which various imagined scenarios test the notion of a general moral principle against the details, viewpoints and other moral shadings of specific situations. The ethics of what constitutes the best course of action could be different for two people, or two societies or cultures. These are profound and complicated problems, and even philosophers who spend lifetimes studying them don't have all the answers. But, in practical scenarios wherein human-designed systems need to make very quick automated decisions about damage and potential loss of life, they are increasingly unavoidable.

This, perhaps, is where young students and scientists come into play. As society continues to develop and rely on these systems and machines, more people will be needed who have a deep understanding of philosophy, thought, law, ethics and computer science, as well as the ways and languages through which these systems intersect and become autonomous. Many universities and colleges already offer advanced study in the area of AI and ethics. This work will only become more important over time. If the opportunities for AI to help and even improve so many things we do seem limitless, society will require people who can think broadly and holistically, and who can create systems and ideas that can best use and grow this technology in humane, productive ways. **&**

COMPUTING AND NETWORKS

TECHNOLOGY AND SOCIETY

DATA

Amazon Future Engineer Canada

Computer Science Training for Remote and Rural Educators

BY SHARON ASCHAIEK

With their profusion of bright colours and repeating shapes and patterns, the computer-coded creations produced by Ron Cameron's Grade 4 and 5 students are groovy and mesmerizing. But a closer look reveals that they're more than just digital artworks—rather, they are the students' expressions of a newly learned geometry lesson about creating regular, closed polygons.

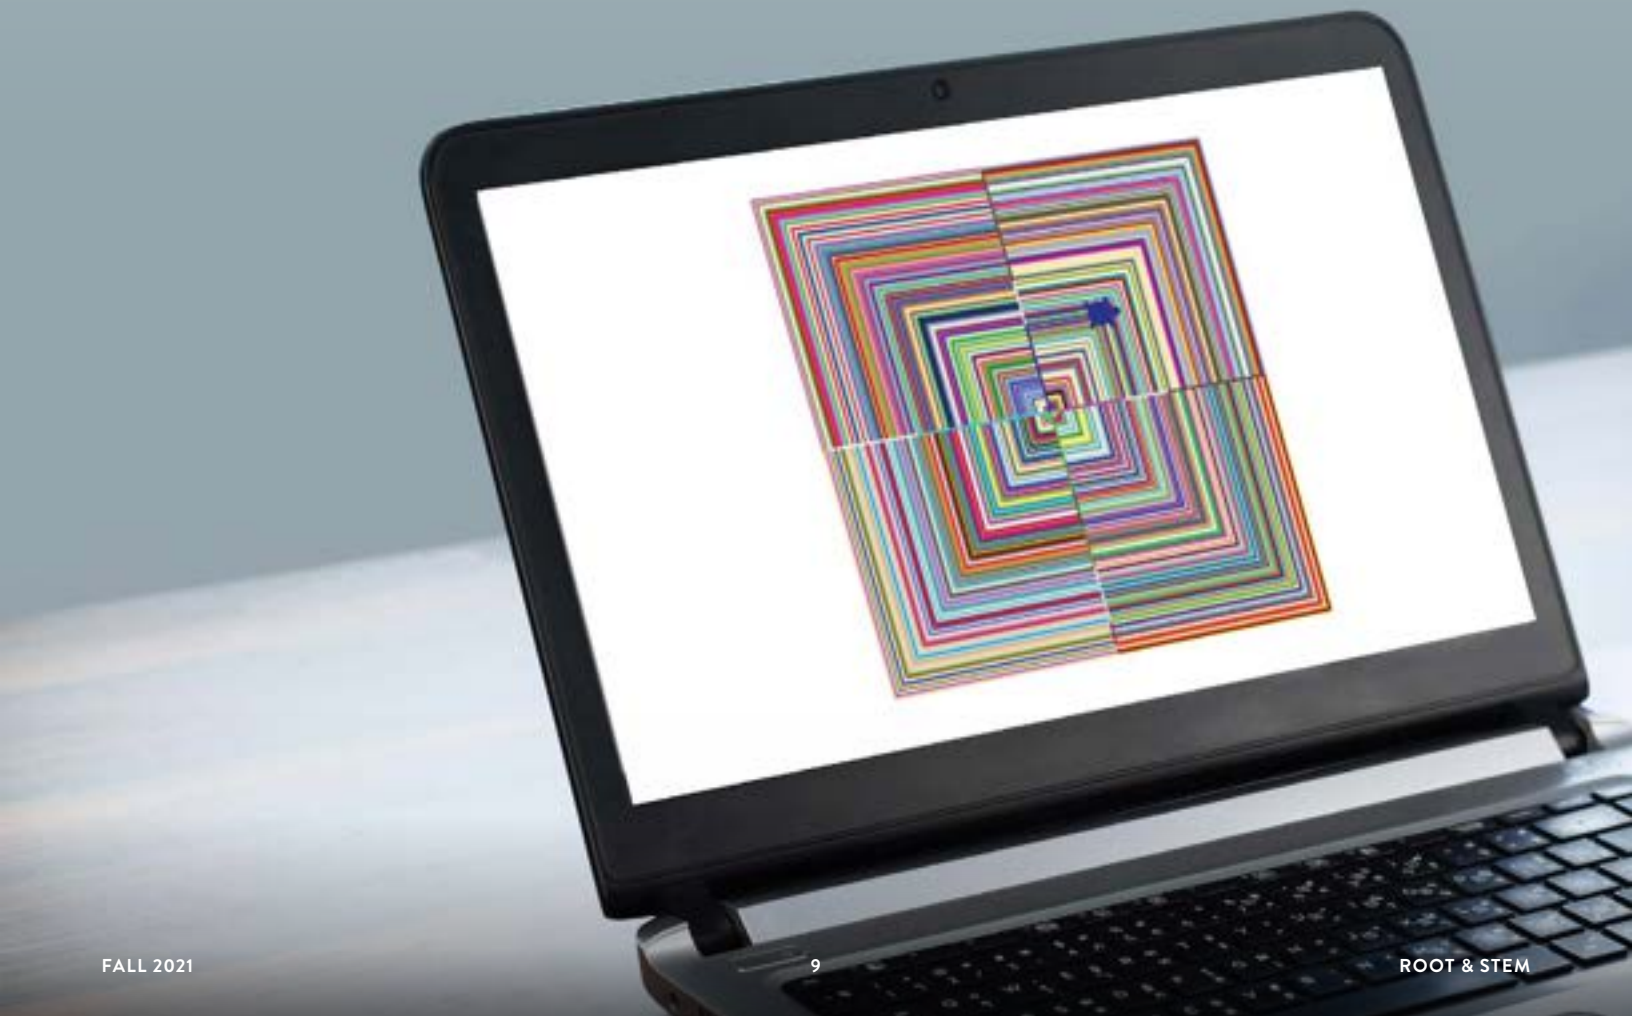

The tell in each is the turtle figure, which the students have programmed to test the "total turtle trip theorem," which proposes that a turtle that travels around any area and ends where it started will have made a total journey of 360 degrees. The students used the programming language Lynx, which involves using basic text commands to write code. "I showed the students how to create procedures in Lynx to teach the turtle to move in different directions, and how to figure out the angles of its trip," says Cameron, who teaches at Nestor Falls School, a two-room K–8 school in Sioux Narrows-Nestor Falls, Ontario. "The whole class was very excited—they thought it was a mindblower."

"It's important to have everyone on a level playing field because access to technology is a right."

Cameron's ability to teach his students Lynx and other coding programs is being supported by Amazon Future Engineer Canada. Launched in January 2021, it involves Amazon investing $3 million over three years to fund free, virtual computer science education for schools in Canada's remote and rural regions. Educators gain access to high-quality STEM instructional materials, skills workshops and online resources. The overall goal is to help learners prepare for the fast-growing digital economy. The initiative is part of the online retail giant's Amazon Future Engineer program, which already serves underserved and underrepresented communities in the United Kingdom and the United States.

At the elementary level, the educational experiences and materials focus on basic coding, robotics and artificial intelligence (AI) skills. Other aspects include a cyber-robotics challenge and an interactive, dance-themed online coding tutorial. The middle- and high-school topics are more advanced coding, AI- and machine-learning skills, and digital citizenship. There have also been events such as the Your Voice Is Power competition, which challenged students to code a brief musical remix of musician Pharrell's song "Entrepreneur" using the EarSketch coding platform. The training is being facilitated by charity partners in Canada that include TakingITGlobal, Canada Learning Code and Kids Code Jeunesse.

Some of this training serves Connected North, a support network for educators from 100 schools in Canada's remote Indigenous communities. Charity TakingITGlobal designs or helps deliver a range of educational programming to Connected North's 2,000 educators. The organization recently directed some of its Amazon funding to run one-day workshops on computer science principles and coding, which is where Cameron was able to update his own Lynx skills. Each participant also received $500 in equipment and resources for their classrooms.

"It's about empowering educators in these regions to ensure students have the digital skills to bring their ideas to life, and to be able to do so in their own communities, if that's what they want," says Michael Furdyk, CEO of TakingITGlobal. "It's important to have everyone on a level playing field because access to digital technology is a right."

It's a sentiment embraced by educator Alysha Coates. A teacher for the last seven years at several schools in northern Ontario and the Yukon, Coates has seen the impact of the digital divide first-hand. Inadequate high-speed wireless internet connectivity is the biggest obstacle, but others include insufficient training and IT support for teachers, and limited computer hardware and software at schools and in students' homes.

Coates currently teaches Grades 5 and 6 at Elijah Smith Elementary School in Whitehorse, which is also part of the Connected North network. Like Cameron, she participated in the recent computer science workshops, and was particularly thrilled to discover the Climate Action Kit, an experiential learning project that promotes computational thinking. It involves students using curriculum models based on the UN's Sustainable Development Goals to build a micro:bit—a pocket-sized programmable computer—that can help tackle climate change. For example, a lesson on the sustainability benefits of indoor farming could involve coding a micro:bit to function as an automated plant-watering system. Coates sees a way to use the kit in a social studies unit that covers water-access issues that affect Indigenous communities.

"With this kit, I can get my students to consider and research the issue of inadequate clean drinking water on reserves," Coates says. "It's a way for them to understand how they can use digital expertise to potentially change the world." &

CODING AND PROGRAMMING

COMPUTING AND NETWORKS

TECHNOLOGY AND SOCIETY

Computer-coded creation produced by Ron Cameron's Grade 4 and 5 class

Artificial Intelligence and Online Learning

BY SOPHIE PENNER

My name is Sophie Penner and I am in Grade 12. I have been homeschooled all my life; for the last three years, I have attended an online school and I consider myself fortunate to have had this experience. When the pandemic hit in 2020, online learning and artificial intelligence (AI) were not new to me—in fact, I cannot imagine learning without AI. I have come to rely on it to prepare for classes and build my confidence. AI is built into many online learning programs to provide teachers and students with unique educational opportunities and experiences. It goes beyond the use of computers and the internet: AI is responsive and has helped me in the learning process. While it is unable to replace the personal experience that learning with a teacher provides, AI can enhance learning in physical or virtual environments. This is the experience I would like to share here.

AI provides a wealth of information, so it provides me with opportunities to teach myself. When I am conducting research projects, computer algorithms display content relevant to the subject I am exploring. Google provides thousands of articles related to my search terms, and YouTube suggests videos with similar content to the videos I am watching.

Fear of failure can paralyze students and may keep them from doing well, but practice quizzes with immediate, automatic feedback from AI encourages trial and error without the risk of failure or criticism. I use many different programs to test myself and improve my skills. When I took the Pre-Employment course offered by the Pinnguaq Association (**pinnguaq.com/learn**), the built-in quizzes would explain correct answers immediately after I submitted my responses to them and, if needed, allowed me to take the quiz again. Online quizzes like Quizlet, Kahoot! and Quizizz are easy to share, so classmates can benefit from learning resources created by individual students. "Practice makes perfect" can truly be achieved with AI because it enables students to practise on their own and receive feedback on their progress.

By using AI to work on my homework assignments, I am able to proceed with confidence. Grammar checkers improve my writing, while videos from YouTubers, Khan Academy and other learning content creators help me to understand content better and communicate more clearly. Have you noticed that Google suggests the ends of your sentences before you have even quite thought them out? That is AI at work and it helps by offering options that can be accepted or rejected. When all this initial work has been done, I love to go to my teacher with confidence to discuss and gain insight from her corrections and suggestions. Teachers cannot be replaced; the key is to use AI to make time with educators more valuable.

AI also can provide significant resources for teachers. Teachers can save time by using computer programs to calculate test grades and check students' work for plagiarism. As the computer grades tests, it can collect data about how long students spend on individual questions and how many students answered incorrectly. This allows teachers to focus on explaining difficult concepts and also spend time with individual students to work on the specific areas where they are having problems. The personal relationship between a teacher and students is encouraging and holds students accountable. AI supports this relationship. It helps teachers to focus on helping their students with particular difficulties in what they are learning, as well as on encouraging students to try again and keep working on the obstacles they are facing. It gives the teacher time to explain and help overcome particular problems, while AI supports general learning and understanding. Hands-on lessons for subjects like woodworking, mechanics and gym can use AI to prepare students for class. Theoretical knowledge is useful when the goal is success in practical execution; when students have that knowledge in place, teachers are able to focus on individual student's work.

Many learning forums have comment boxes beneath videos or discussion forums that allow learners to communicate with and help each other. This form of learning allows people to connect around the world and work together, all with the help of AI. I have connected with hundreds of people I have never met in person but, despite the lack of in-person interaction, I have been able to have fun learning with these other online learners.

The wealth of information I can access via the internet, alongside the feedback I receive from computers for my work, has helped to elevate my learning experience online, and made it incredibly meaningful and valuable. Teachers will continue to provide instruction and mentoring that is personal and specific to my needs as a learner; a teacher is someone to whom I am accountable, and human interaction makes learning fun and engaging. AI is a digital tool—and, like any tool, how we choose to understand and use it will determine the opportunities it provides and the experiences we have with it. &

TECHNOLOGY AND SOCIETY

DESIGN

DATA

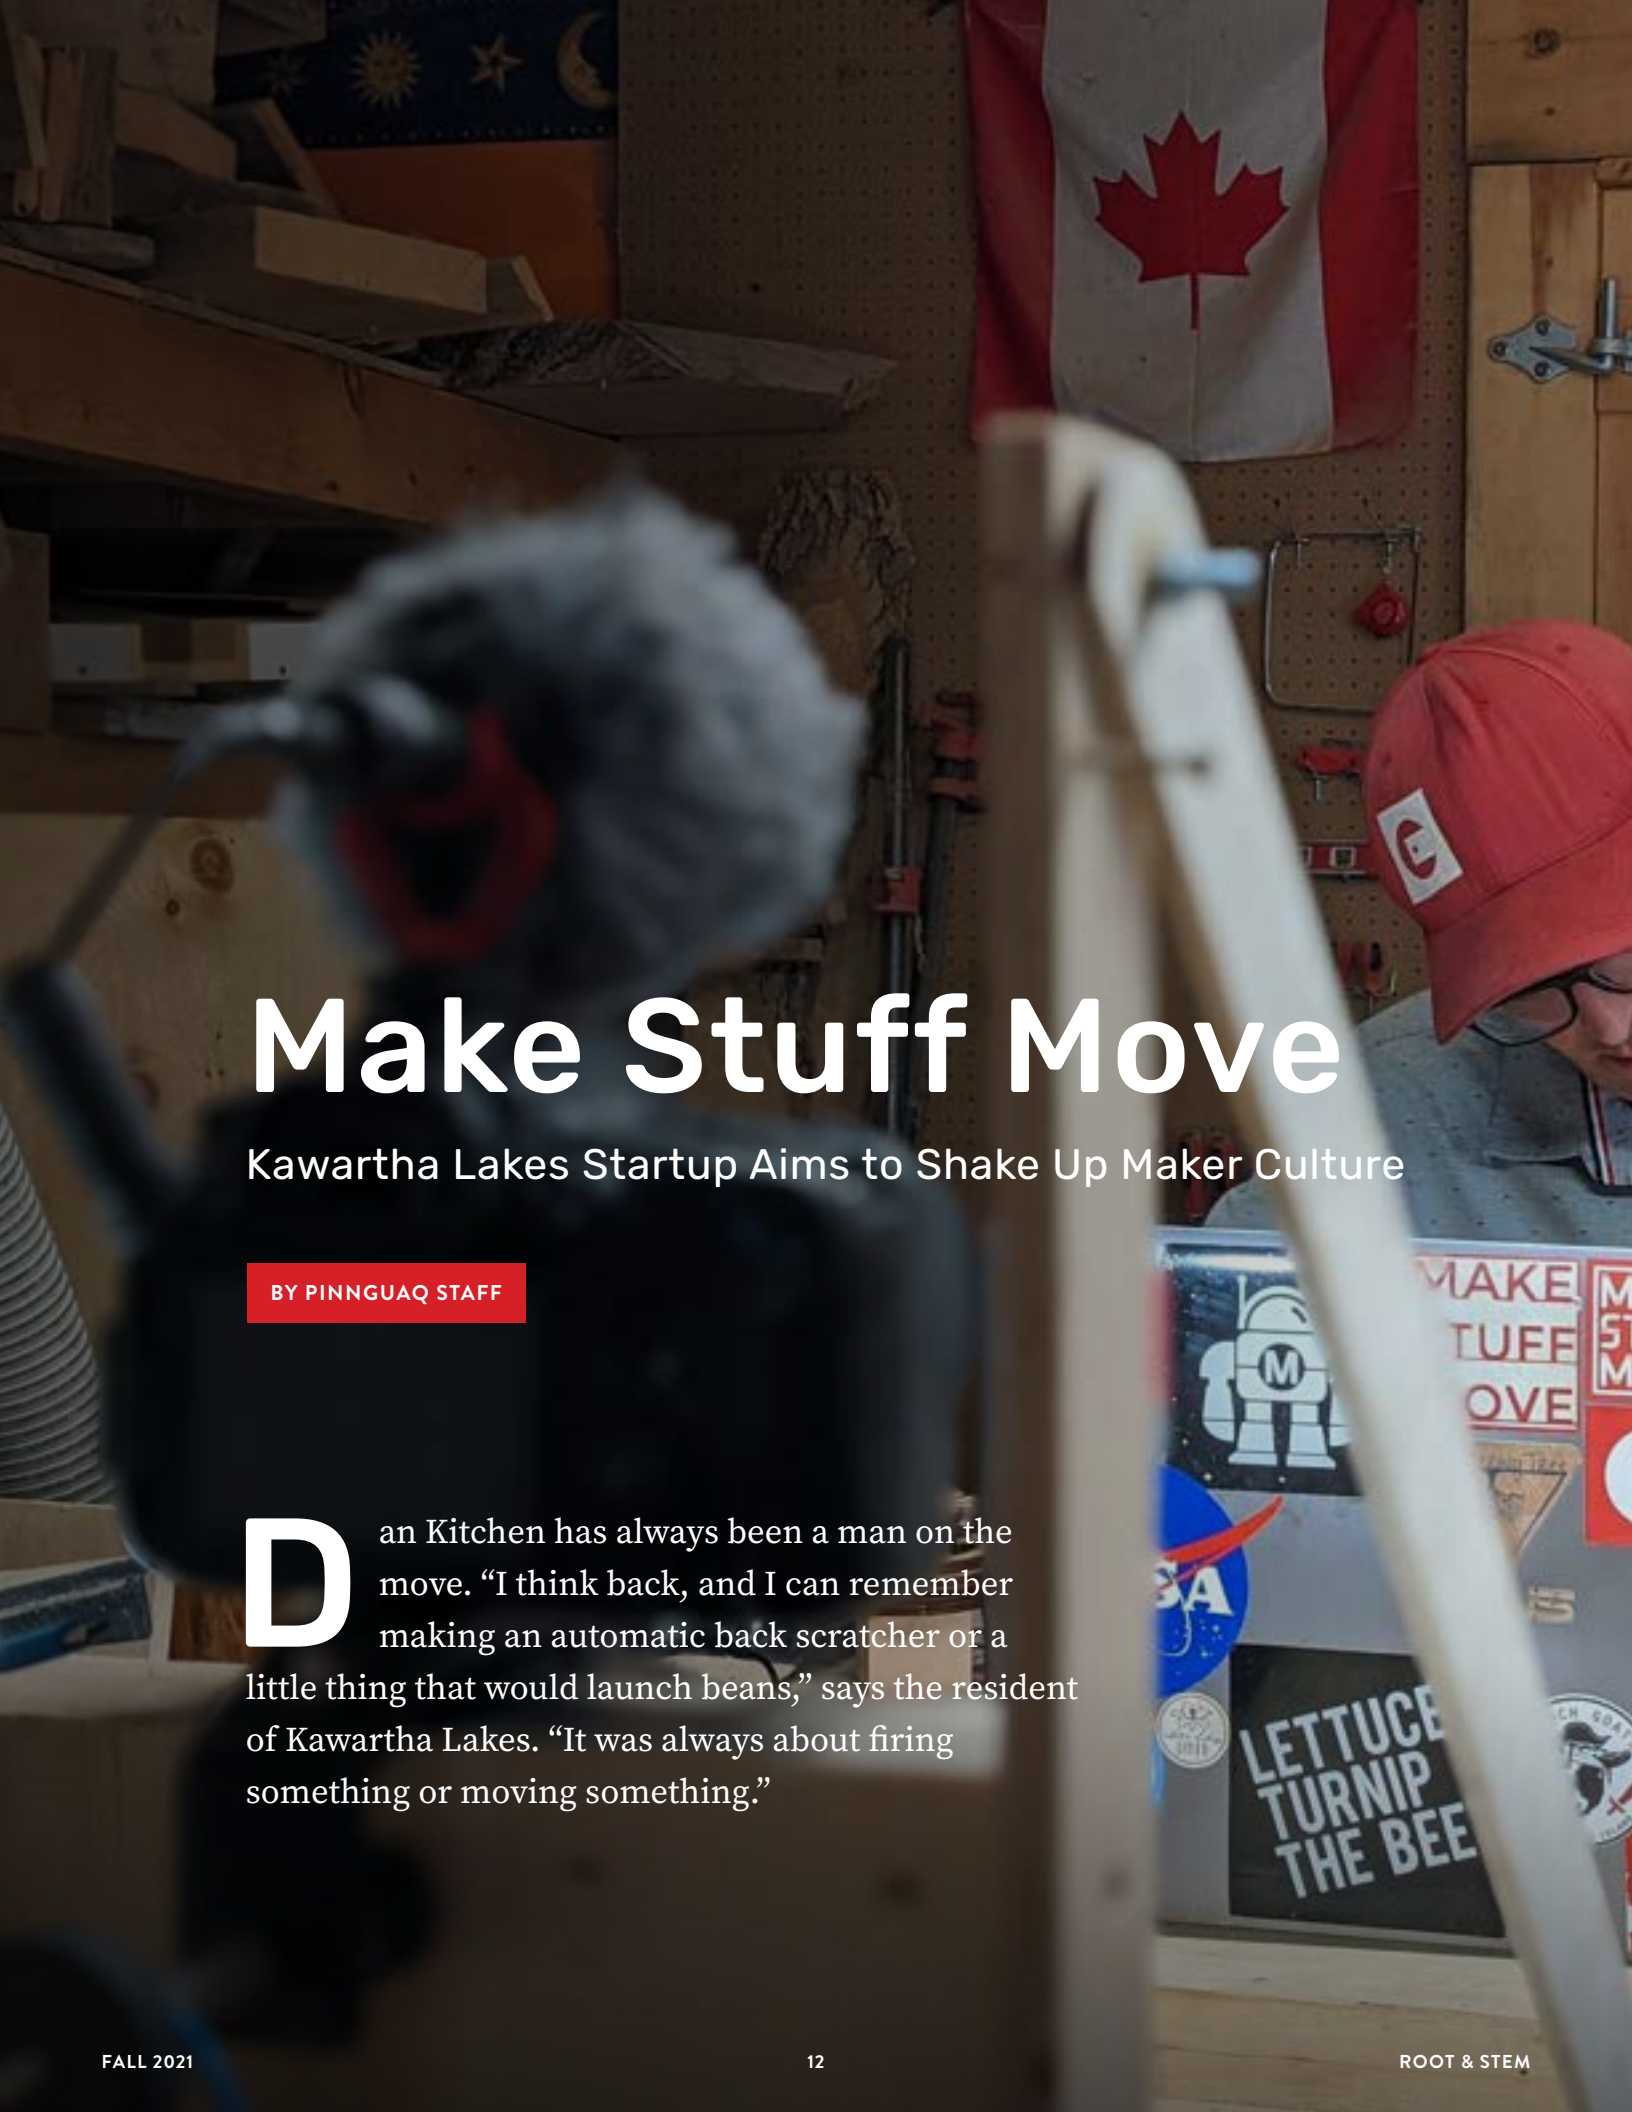

Make Stuff Move

Kawartha Lakes Startup Aims to Shake Up Maker Culture

BY PINNGUAQ STAFF

Dan Kitchen has always been a man on the move. “I think back, and I can remember making an automatic back scratcher or a little thing that would launch beans,” says the resident of Kawartha Lakes. “It was always about firing something or moving something.”

That passion for motion is what powers Kitchen's latest venture. Make Stuff Move offers tools and resources for educators and makers looking to learn more about how to tinker well. It's driven by the idea that using tools and everyday materials to develop hands-on skills is an effective and fun way for kids to learn about STEAM.

Unsurprisingly, the project originated with a problem that needed solving. Early on, Kitchen discovered that although he liked making moving machines, there was a pesky little problem in the middle of them: the motor. Specifically, how to mount it.

"I was trying to build this robotic bucket that drove around, and I was spending most of my time trying to mount motors into this project, making all kinds of weird brackets," says Kitchen. He noted that this was the same problem he'd had as a kid learning to tinker. "I made a few sketches and had a 3D printer. So I started to play around a bit."

What resulted is the Servo Sock, a casing that snaps over top of a regular, replaceable servo motor, adding a surface that can easily be mounted, cut, shaped, screwed into or otherwise altered to suit the builder's needs. When connected to a controlling device, such as an Arduino or Raspberry Pi, the motor can easily be used to create moving objects.

It took Kitchen and his business partner, Jon Edwards, about three years to get a patent for Servo Socks. At the beginning of 2020, they were just starting to build relationships with school boards and getting ready for their first big trade show when the pandemic forced school closures around the world.

"Last year was tricky," Kitchen says. "But the silver lining was developing a much better kit." Their kits now include a customized Arduino shield (a modular circuit board that can be added to an existing Arduino to add functionality) and are packaged up in a toolbox with accompanying lesson plans and links to online learning resources.

Restrictions that were put in place for locations like schools and makerspaces were particularly tough for Kitchen, whose philosophy is rooted in the kind of exploratory, hands-on approach that workshops both encourage and foster.

"I like to build stuff," says Kitchen. "I'm a firm believer in using common wood and screws and glue—not kits. I like Lego, I like Meccano. But with kids and learning, there's something about the confidence that can be gained when you're actually grabbing a piece of wood and a box of screws from the hardware store. Even if the result isn't as impressive, the experience is.

"Often with kids, that's all it takes—one time of seeing, 'Yes, you can use the drill. It's not an overwhelming piece of equipment. This is the part that will go through your finger; don't put your finger in front of that. The rest of it is safe.'"

Servo Motors

A servo motor is a self-contained electrical device that can rotate parts of a machine with a high degree of efficiency and precision. They are used in closed-loop systems, in which sensors feed information from the output signal back to the input in order to control the speed or rotational position of the motor's shaft. Servo motors have many applications, and can be found in everything from military vehicles to your camera's auto-focus mechanism.

Kitchen is a fan of the more tech-focused aspects of STEM, too. He spent years working as an animator, and has plans to develop an app that will draw on that experience and allow users to animate robotic creations, using a custom Arduino shield that can capture and record motion. In fact, he credits Arduino with getting him back into tinkering as an adult.

But his heart is in his toolbox. "It's kind of a bait-and-switch with our kits," he says. "Like, 'Here, play with STEM and robotics and coding—but oh! Here's a screwdriver, here's a wrench. Here's a piece of pine.'"

"For me, it's all about understanding logic and problem-solving. That applies to both sides of the equation—whether you're

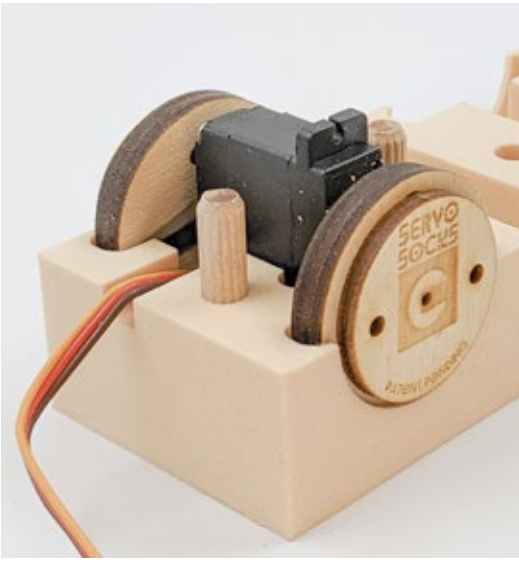

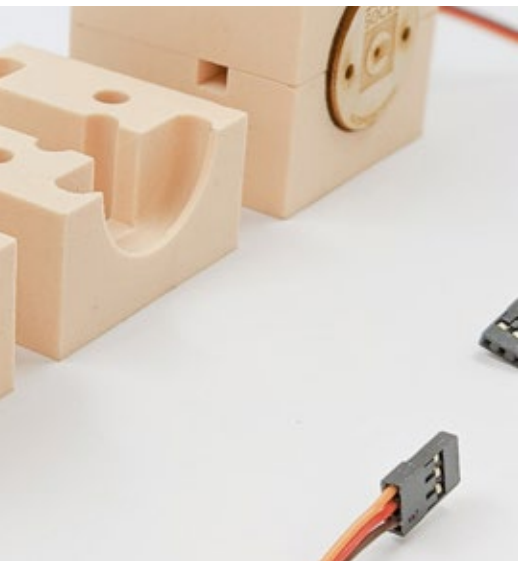

building something, or coding something. It's the same thought process, overall. I feel like there's a real bridge between the two."

True to his ethos, Kitchen manufactures most of the materials in the kit at his InventArt studio in Kawartha Lakes, which keeps his process relatively simple.

"If we need to scale, it's kind of neat. I can just get another CNC machine and plunk it down beside the existing one, and I've just doubled capacity."

As things stand at present, those who want to try a kit have to contact Kitchen directly through his website. In the meantime, he's building out his studio using shipping containers, considering hiring staff to keep up with demand and preparing to launch an e-commerce option—and yes, there will be a basic bucket-on-wheels kit coming soon. &

CODING AND PROGRAMMING

DESIGN

For more information and ready-to-go lesson plans, visit **makestuffmove.com**

DECODE STEAM

AI IN EVERYDAY LIFE

Artificial intelligence (AI) refers to the use of machines that are programmed to simulate human intelligence. AI is no longer relegated to the realm of futuristic cartoons or sci-fi blockbusters; AI has quickly integrated into many aspects of our communities and homes. In their day-to-day lives, humans encounter various forms of AI all the time. The following are just a few examples of how AI is all around us.

NETFLIX, SPOTIFY

Netflix algorithms not only hand-pick movies and shows for you, but also tailor-make thumbnails that it predicts are most likely to get your clicks.

SEARCH / RECOMMENDATION ALGORITHMS

Google processes 3.5 billion searches a day using deep learning.

PERSONAL ASSISTANTS

Personal assistants (PAs) such as Siri and Alexa collect data from their every interaction. Most PAs are female because research shows that most people prefer the sound of a female voice.

For our list of references, please visit pinnguaq.com/learn/ai-everyday-life

ART

Portrait of Edmond Belamy, created by the artist collective Obvious, was the first piece of AI-generated art to sell at auction. It was purchased for $432,500–$422,500 over expectations.

PREDICTIVE TEXT

In 2020, *The Guardian* published its first article completely written by a predictive text AI, which suggests words based on the context of other words around it. Titled *A robot wrote this entire article. Are you scared yet, human?*, the piece aimed to convince humans that robots come in peace.

SOCIAL MEDIA

Facebook uses a deep learning (DL) application called *DeepFace* to teach it to recognize faces in photos.

SELF-DRIVING CARS

Self-driving cars use a form of AI called deep learning (DL), which allows them to learn and predict possible driving scenarios so they can navigate the roads safely. *The Autonomoose*, developed at the University of Waterloo, was the first true autonomous vehicle (AV) to hit the roads in Canada. Data collected has been distributed to several automotive companies to advance their AV technology.

CAN YOU FIND THESE THINGS IN THE SCENE?

ARTIFICIAL CREATIVITY

How AI Is Changing Cultural Jobs

BY JESSA GAMBLE AND BRIANA BROWNELL

There's a lot of concern about how artificial intelligence (AI) and automation could soon replace people in jobs like truck-driving or proofreading. But in creative fields like art, music and dance, pioneers are experimenting with AI as a tool to expand the limits of creative possibility. Their work points to new ways to unite the arts, challenge patriarchal systems, and form healthy relationships with our new non-human collaborators.

PHOTOS COURTESY CELESTE LANUZA AND XANDER STEENBRUGGE

Celeste Lanuza has long been combining dance and social justice commentary, with a special focus on addressing racial inequities. Though she has always felt the importance of personalizing her dance to express her identity, culture and emotions, Lanuza has never been a fan of the ever-present mirror in the life of a dancer, whether it lines the walls of a traditional dance studio or reflects the dancer on a screen.

"We have the mirror in the dance studio as a tool, which many times is a crutch," she says. "Students and professional dancers I work with just stare at it, and it becomes a source of judgemental comparison."

"I love working without the mirror. And especially during COVID, I've been exploring dancing outside in the natural air."

Recently, Lanuza started exploring motion-capture technology as a way to push boundaries even further. Using cameras to track dancer's real-life motion, the technology transfers those movements onto a digital avatar, thereby allowing the dancers she choreographs to move with entirely different bodies. Through the use of AI, dancers can even create partners who are uniquely attuned to their movements, either onscreen or in the form of a robot.

"It liberated the dancers into feeling that freedom of expression, which is where dance comes from. So that was really exciting," says Lanuza.

The Future of Dance

Coming from a Latinx family, Lanuza struggled in middle school with English as a second language, but dance was a great equalizer that allowed her to speak fluently. Children don't need any equipment, like pianos or basketballs, to dance—they simply use the body they already have. Lanuza feels that if it is started very early on, dance can teach essential skills related to children's connection with themselves, their relationships to each other, and their connection to the histories in the places where they dance.

Younger children can start with structured improvisation, an alternative approach to the strict conventions of ballet.

They can experiment with time, energy, rhythm, melody, emotion and space—all areas that are full of potential for productive play. Improvising stories and themes unleashes the imagination of children before they have started any critical evaluation of their own and each other's dances. Then they can move on to formations and follow-the-leader exercises that allow them to practise leadership.

If they decide to incorporate movement into their creative lives later on, the technology will be there for them to bridge enormous gaps of space and time: they will be able to dance in other bodies, perform a duet with their great-granddaughter, or even—as we will soon see—to create music as they move. One way this can happen is through AI-enabled "pose estimation," which allows a computer to perceive the position and movement of a body without the spatial reference points required by previous generations of motion capture technology.

Blending the Senses through AI Art

Xander Steenbrugge lives in Ghent, Belgium, and works remotely for a US company, applying machine learning to problems in biology. He's been watching the development of technology for generative adversarial networks, or GANs.

Invented by a Google team in 2014, GANs are discriminative networks that will be shown a collection of representative images—for instance, the works of Dutch masters such as Vermeer, Bruegel the Elder and Rembrandt—in order to "train" it to be able to look at a fresh image and determine its authenticity. A complementary generative network creates images and tries to fool the discriminating network into accepting them as Dutch masterpieces. Essentially, this looks like two duelling systems playing a game of visual Balderdash.

About two years ago, they developed the ability to create high-resolution images. GANs have also been used to create realistic images for computer games and fake videos (for example, recreating the late actor Carrie Fisher to appear in the most recent *Star Wars* film).

"That's when I started paying attention and started playing around with it," says Steenbrugge. "I've always been driven by looking for new ways of finding beautiful aesthetics using technology."

Under the name *Neural Synesthesia*, Steenbrugge used the networks to create videos that blend one image into another and allow music to determine the flow of images.

Portraits of Celeste Lanuza (bottom left) and Xander Steenbrugge (top right). WZRD samples (middle and bottom right)

"This project is an attempt to explore new approaches to audiovisual experience based on artificial intelligence." say Steenbrugge. It became a new form of music visualization that, in a way, collaborates with the brain, which is so busy integrating information from all five senses that it readily sees ways they might correspond and blends them into a kind of synesthesia.

As visually spectacular as these results can be (as seen on Steenbrugge's Vimeo page, **vimeo.com/neuralsynesthesia**), one thing that's usually missing from AI-generated visuals, music and even stories, is narrative structure. Networks are very good at generating patterns—say, creating a song with the general feel of a Joni Mitchell composition—but they often feel like a recurring snapshot of a style; they don't begin with an introduction, include riffs and motifs that banter throughout, or build to a climax. Theoretically, an AI that developed those capabilities could tip the balance into producing works that genuinely compete with human-made songs for our affections.

Those works might not even have to be original to have artistic value. A perfect imitation of a master's style could create art that moves us just as much as the works in that master's œuvre.

Andy Warhol was one of the first to anticipate this development when he said, "I think somebody should be able to do all my paintings for me. I think it would be so great if more people took up silk screens so that no-one would know whether my picture was mine or somebody else's."

Leading AI historian Pamela McCorduck agrees: "Do we really have all the late Verdi and Shakespeare that we want? Of course not: We have only what accidents of history permitted us to have," she writes in her book *Aaron's Code: Meta-Art, Artificial Intelligence and the Work of Harold Cohen.*

• • •

As far back as human memory stretches, we have worked side by side with animals. Whether it's a shepherd whistling complex directives to her sheepdog or a cowboy at one with his horse, there's a human competency that allows us to empathize with minds and hearts vastly different from our own and align ourselves toward a common goal.

In the future, creative professionals who rely on AI to co-create their art might experience something similar. A lot of the human contribution in AI art comes from the dataset they choose to feed their AI collaborator. That training set defines a large part of the output. They can also set the initial parameters of what they want to see and often have an intention or direction in mind and a rough idea of what the computer might come back with.

But once the network starts churning out options, the creative process becomes much more of a dance, a kind of two-way banter between human and machine. The human tweaks the parameters to set a slightly new direction, and the machine provides serendipity that sparks both delight and some new ideas in the human mind.

"Very often I've been surprised by what it generates," says Steenbrugge. "The model comes back with feedback that is somewhat unpredicted, and then this cycle is really interesting because it sometimes leads you to places you would never have gotten on your own. "It's really a dialogue between man and machine."

The Future of Multimedia Art

When Steenbrugge realized that some trained professional artists had no access to the tools he was using, he decided to address it by building a platform. At **wzrd.ai**, anyone with some music, some images and an idea can make their own trippy videoscape.

He explains the process on his website. The process is complete in three steps.

"**1. Audio analysis**—The audio analysis pipeline can detect percussion and harmonic elements, which are used to drive different parts of the video.

"**2. Visual engine**—WZRD's magic works by using audio elements to drive a machine learning technique called GAN. This results in an entirely new kind of visual experience.

"**3. Video render**—Finally, a full video render is performed and combined with your audio." (**wzrd.ai**)

Above all, Steenbrugge was curious to see what other creative people without the preconceptions of those in his own field might do with these types of models. It's the way he expects things to be for the next generation. His prediction is that within the next five years, the tools he uses will be available in a non-intimidating form, ideal for the creative person interested in engaging in digital tech through a simple and accessible tool.

"We've seen this happen in technology a lot of times. When something is new, it's usually kind of techie and you have to be a little bit nerdy and specifically trained to use these things, but as time progresses, layers

and layers of abstraction are put on top of it and then it becomes usable to a larger audience," he says.

Steenbrugge has been exploring and working with other forms of digital tech tools such as virtual reality headsets to turn his strange and compelling multimedia pieces into immersive spaces.

When the Body Is a Keyboard

Celeste Lanuza's work in the world of dance and Xander Steenbrugge's in visual art and music might occupy separate silos in arts and culture, but as far as computers are concerned, they speak translatable languages. For that matter, *anything* that changes over time can be an equivalent source of information, as long as it's measurable. Steenbrugge's art can be informed and shaped by the beat of an instrumental percussion track, but could just as easily be driven by the heartbeat of the listener.

What's exciting for Lanuza is the fact that her dancers could one day be influencing the music they dance to in real time, composing sound with their movements—and being moved, in turn, by that music.

"I even teach that way now—I find myself using the same vocabulary, saying our body is an orchestra. We need to think of the different parts of our body being different instruments, and thinking of ourselves as music rather than a separate entity," says Lanuza.

The possibilities straddle discipline and form: a disabled musician who uses a wheelchair might choreograph a virtual dancer through the medium she knows best, or two dancers' pas de deux could generate a written love story in a particular writer's literary style.

Artificial Intelligence Itself Is a Creative Artform

The idea of artificial intelligence does not belong to computers and algorithms. Its roots run much deeper, with the ancient human wish to forge the gods. The arts have expressed these ideas through myths that go back millennia—and for just as long there have been conflicting feelings on the prospect of an independent intelligence designed by human minds.

"This has been a human impulse for thousands of years, to create something outside the human cranium that has intelligence," says Pamela McCorduck on Lex Fridman's podcast (Fridman is an AI researcher working on autonomous vehicles, human-robot interaction, and machine learning at MIT). "Homer has robots in the *Iliad*, and the *Odyssey* is full of robots. How do you think Odysseus gets his ship from one place to another?... So we've had this notion of AI for a long time."

The Ancient Greeks had a positive take on it, largely celebrating the robots that served as helpers in Hephaestus's forge. The Ancient Hebrews looked to the Second Commandment, which forbids making graven images, interpreted to include any imitation of humans, including the non-visual. Much like other transgressive art forms, AI faces accusations of blasphemy. Stories of lifelike automatons in the Zhou dynasty were also interpreted with negative connotations, but instead of blasphemy they highlighted trickery. A story in the *Book of Liezi* tells of a lifelike automaton that, though initially met with wonder by the court, drew the wrath of King Mu when it flirted with his royal concubines. Our literary conventions bear this out in their habit of pitting human heroes against villainous machines. Think of that most famous of artificially intelligent characters, Frankenstein's monster.

Two key innovators, Alan Newell and Herbert Simon, broadly accepted as the "godfathers" of the artificial intelligence field, were not exclusively—nor even primarily—focused on its technological aspects. They were cognitive psychologists and Simon was a serious painter who felt forced to give up his art because it was taking too much time away from his AI research.

For McCorduck, best known for her seminal 1979 text on AI, *Machines Who Think*, the process of creating human art is a lot like what a generative system does. Art begins with the impulse to make something special, taking something that could be minimal and functional, but creating something that transcends that purpose. It's an urge she feels is as essential to life as protein.

"Art-making behaviour is universal among humans. It amazes me to consider that there is no such thing as a no-frills human culture, with clothing only to cover, and food only to eat and housing only to shelter. We decorate and design and present artifacts to one another constantly, every last one of us," she writes.

Most of that expression doesn't take the form of the elitist European realm of High Art, but rather as a daily practice by non-specialists, of storytelling and music-making and adorning oneself.

To Lanuza, that daily artistic practice is a lesson in equity, which is at the heart of the question of technology in the arts. When AI can let people transcend their language differences and even translate the vocabulary of one art form to another, something special happens. Even the idea of "good" dancing and "bad" dancing comes into question as computer algorithms make up their own, strange, version of dance.

Maybe that's why most of the commercial robots released today that are designed to interact with humans come with a dance mode. Whether these robots dance for you or with you, they touch a place in your heart reserved for those who allow themselves to be vulnerable by making and sharing art. The robot knows that if you and she are to live together, that's how you first connect. **&**

TECHNOLOGY AND SOCIETY

DESIGN

DATA

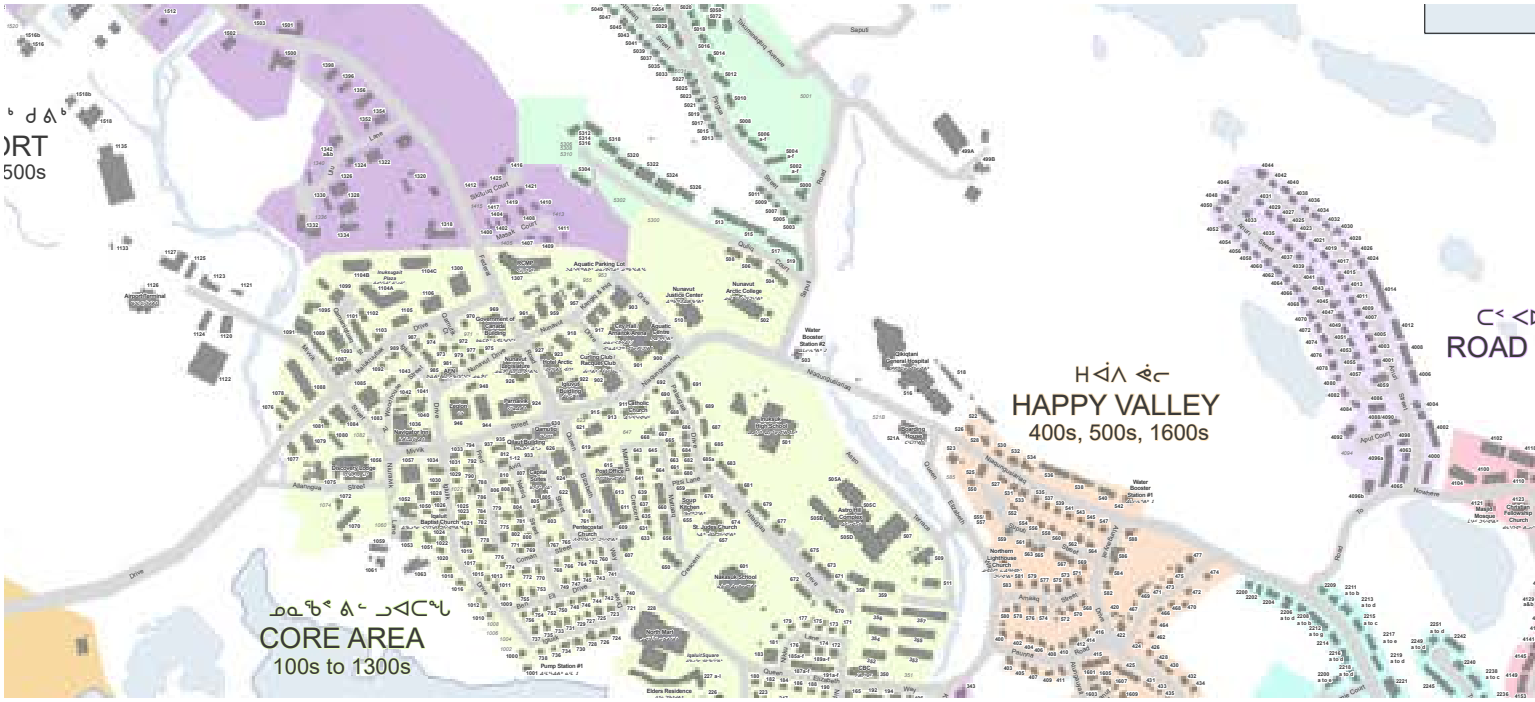

Creating an Address Tool Using AI

BY ALEX SMITHERS

Anybody who's lived in a small town is probably used to receiving directions like, "Oh, I live in the house with the blue roof behind the Coop," or "The Country Food store is just past the Arctic Survival Store."

When I moved to Iqaluit from Toronto in October of 2019, not only was it my first time living in a city not built on a grid, but my first time having to navigate without knowing the landmarks or having Google Maps to guide my every turn.

Unlike most places where addresses are based on streets with ascending numbers, in Iqaluit every building is given a unique number. This means an address can be specified with only a number—a convenient shortcut for Iqalummiut, but tricky if you just moved to town and are trying to drop off a résumé somewhere.

To make matters worse, Google Maps doesn't respond well to a search for just a number. It's more likely to return a Baskin Robbins touting its variety of flavours than the "building 31" you're searching for.

Fortunately, being spotted wandering around aimlessly in the freezing cold, a community member let me know that there's a better way. The City of Iqaluit published a PDF in 2018 with every building number on it. Above is a detail of the downtown area. Can you find the Pinnguaq Makerspace, building 1412?

Even with the help of this image, I started thinking, "wouldn't it be nice if there was an online lookup service that would allow you to input a building number and see its location on a map?"

So I created one.

How I Did It

To build the service, I followed four steps:

1. Find all the alphanumeric characters on the map
2. Identify those characters using machine learning and extract the building numbers (this will be the main focus of this article)
3. Figure out the geographical coordinates of the building numbers
4. Integrate the number and coordinate data with a mapping service

Visit **pinnguaq.com/learn/creating-an-address-tool-using-ai** for a step-by-step tutorial on how I did this. With mapping service provided by OpenStreetMaps and Mapbox, finding the Makerspace is now just a little bit easier. &

CODING AND PROGRAMMING

TECHNOLOGY AND SOCIETY

DATA

Different Tools for a Different Garage

BY GREG PICKEN

Makerspaces are designed to inspire us to bring big and small innovative concepts to life by providing access to tools, knowledge and a community of makers. This shared space and equipment can range from design software and 3D printers, to metal- and woodworking tools, to art supplies and whatever else is needed to help people turn their ideas into reality.

In fact, this maker mentality is becoming so popular that Microsoft, one of the world's largest technology companies, has created the Microsoft Garage as a way to power creativity and ingenuity, for the benefit of both the company and its employees.

With 12 locations around the world, including one in Vancouver, The Garage provides physical space, tools and guidance. As a result, Garage Program Manager Chris Richards has seen his colleagues given the freedom to experiment, to test, to fail—and to keep trying.

"The sites have a makerspace component, so they have a physical area with all the great tools," Richards explained. "You want to prototype something, you want to build something, you just want to tinker, do arts and crafts, or if you want to go all the way up to robotics and embedded systems, we have everything you need."

What makes The Garage work isn't just the hardware, it's also the mentality in the physical space, which Richards, an engineering physicist with a background in automotive research, says clearly puts Microsoft's cultural priorities into action.

"It's an environment of creativity and tolerance for the risk process in product development. If you have a crazy idea that just might work, The Garage is really good at figuring out what's the fastest way to test that: Does your hypothesis hold water?"

The global network of garages is also home to the world's largest private hackathon, the Microsoft Global Hackathon, where Microsoft employees are challenged to bring their most innovative and unique solutions to real-world problems.

"Last year in the summer, even working from home due to COVID, over 70,000 people participated," said Richards. "This year, it's in October, and we're hoping to see even more Microsoft people and select external partners join us."

With thousands of employees using the garages, for ideas that can be serious or whimsical, projects range from accessibility or quality-of-life additions for Microsoft programs to personal hobbies. This means there's no one-size-fits-all approach to projects.

"We run the full gamut," Richards explains. "A business leader might have identified a good market that we want to get into, but we're not sure if we want to dedicate full resources to it. Through The Garage, they can put a small team together to go through ideation, market analysis, prototyping, and market research."

"That's one end of the spectrum," he continues, "but on the other end is anything that someone is passionate about. I've had people coming and tinkering with 3D printers to make board game pieces, miniatures or custom cookie cutters."

Richards even saw a small team working on a physical mute button for video calls.

Having experimented with his own fun side project to deepfake a departing colleague into a music video, Richards also appreciates that when you provide access to new tools and technologies, it's important for users to understand proper ethics and safeguards.

"Fundamentally, what it comes down to is, yes, these tools can be used to do some gross things, but as ethical developers, if we want to counter those gross things, we have to understand how those tools work," he outlined. "I learned a ton about how to recognize deepfakes. I see a little bit of pixellation here, a little bit of colour tone shift here, and I now know to ask, 'Has it been faked?'"

For many Garage projects, artificial intelligence tools can play an important role, and experimenting with AI on these side projects is a great way for Microsoft employees to learn and expand their skill sets.

"AI will become a fundamental tool—it already has," Richards believes. "For the people who are coming up in the system right now, going through their education and looking at careers, having a grasp on functional machine learning and artificial intelligence is going to be a pretty standard tool in your toolkit."

While the people who work in The Garage are all highly trained tech professionals, Richards sees the same lessons and makerspace mentality fitting into the classroom, even with simple tools and supplies.

"From an experiential point of view, when you make something happen with your hands, that's exciting," he enthused. "You get that rush, the first time you make the light come on, or the buzzer sound, or the motor rotate. That's magic." &

• • •

Microsoft Garage Projects

Many digital projects have roots in The Garage. Some may be incorporated into existing Microsoft software, some may end up as their own standalone tools, with the best recognized on the The Garage Wall of Fame (**aka.ms/walloffame**).

Farm Beats
microsoft.com/en-us/garage/wall-of-fame/farmbeats
A data-driven farming project designed to make AI available to farmers around the world, increasing farm productivity and reducing costs.

Seeing AI
microsoft.com/en-us/ai/seeing-ai
Harnessing AI in an app to help the visually impaired navigate their world.

Watch For
microsoft.com/en-us/garage/wall-of-fame/watch-for
Deploys AI for analyzing and surfacing interesting elements of a livestream in real-time.

CODING AND PROGRAMMING

TECHNOLOGY AND SOCIETY

Learn more about the Microsoft Garage at **microsoft.com/en-us/garage**

Deep Blue

1966

BY VICTORIA HETHERINGTON
ILLUSTRATED BY CINDY ROSE

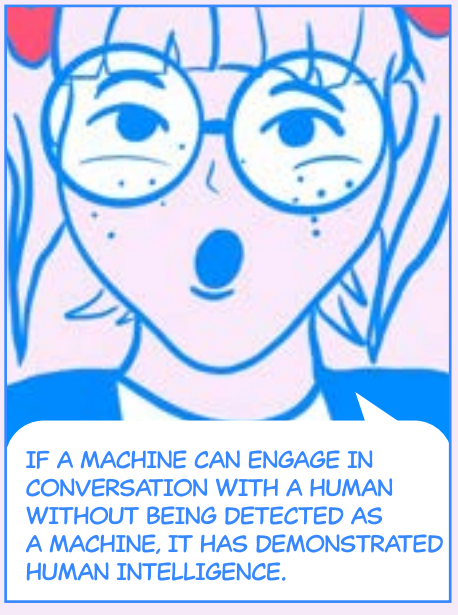


PERHAPS YOU CAN HELP EXPLAIN THAT TO THE CLASS LATER, ANNIE. HOW DO YOU KNOW ABOUT ELIZA?

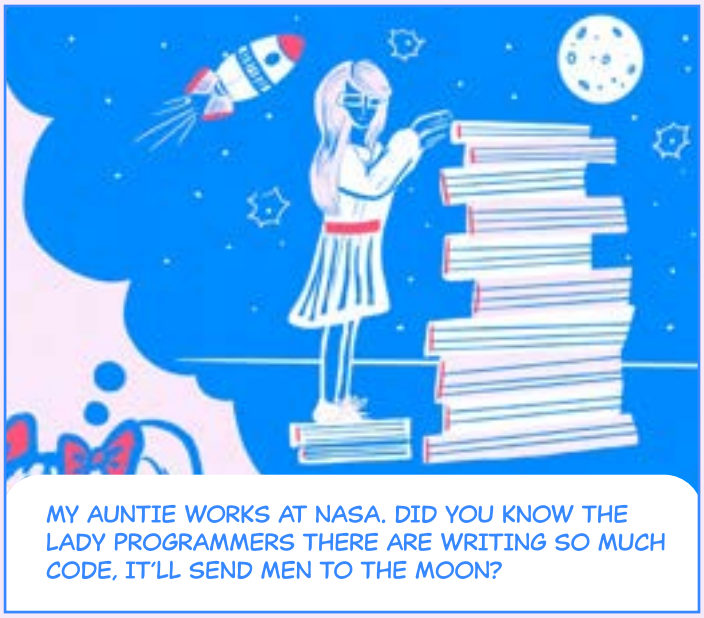
MY AUNTIE WORKS AT NASA. DID YOU KNOW THE LADY PROGRAMMERS THERE ARE WRITING SO MUCH CODE, IT'LL SEND MEN TO THE MOON?


G
Force
L=2W


AND THAT'S ELIZA. ANY QUESTIONS?


WILL THE ROBOTS BE NICE TO US?


OH, TOBY, I'M SURE THEY WILL.


THEY AREN'T ROBOTS, TOBY, THEY'RE SCRIPTS. THAT MEANS THEY RUN A SERIES OF FIGURES CALLED CODE. LIKE THE LADY CODERS AT NASA.
WILL WE HAVE THEM IN OUR HOUSES? WILL THEY BE BABYSITTERS? CAN THEY REPLACE MRS. JONES DOWN THE STREET? SHE NEVER LETS ME HAVE COOKIES AND SHE SMELLS LIKE MOTHBALLS.
I'M DETERMINED TO WORK AT IBM ONE DAY. SO I'LL BE THE FIRST TO KNOW IF THEY DO.
1997
IBM
NEWS
NEWS
Kasparov vs. Deep Blue
Control Room
THEY SAID IT WOULD TAKE DECADES FOR DEEP BLUE TO BEAT KASPAROV, THE PINNACLE OF GENIUS IN CHESS. HE WON HIS MATCH AGAINST BLUE LAST YEAR.

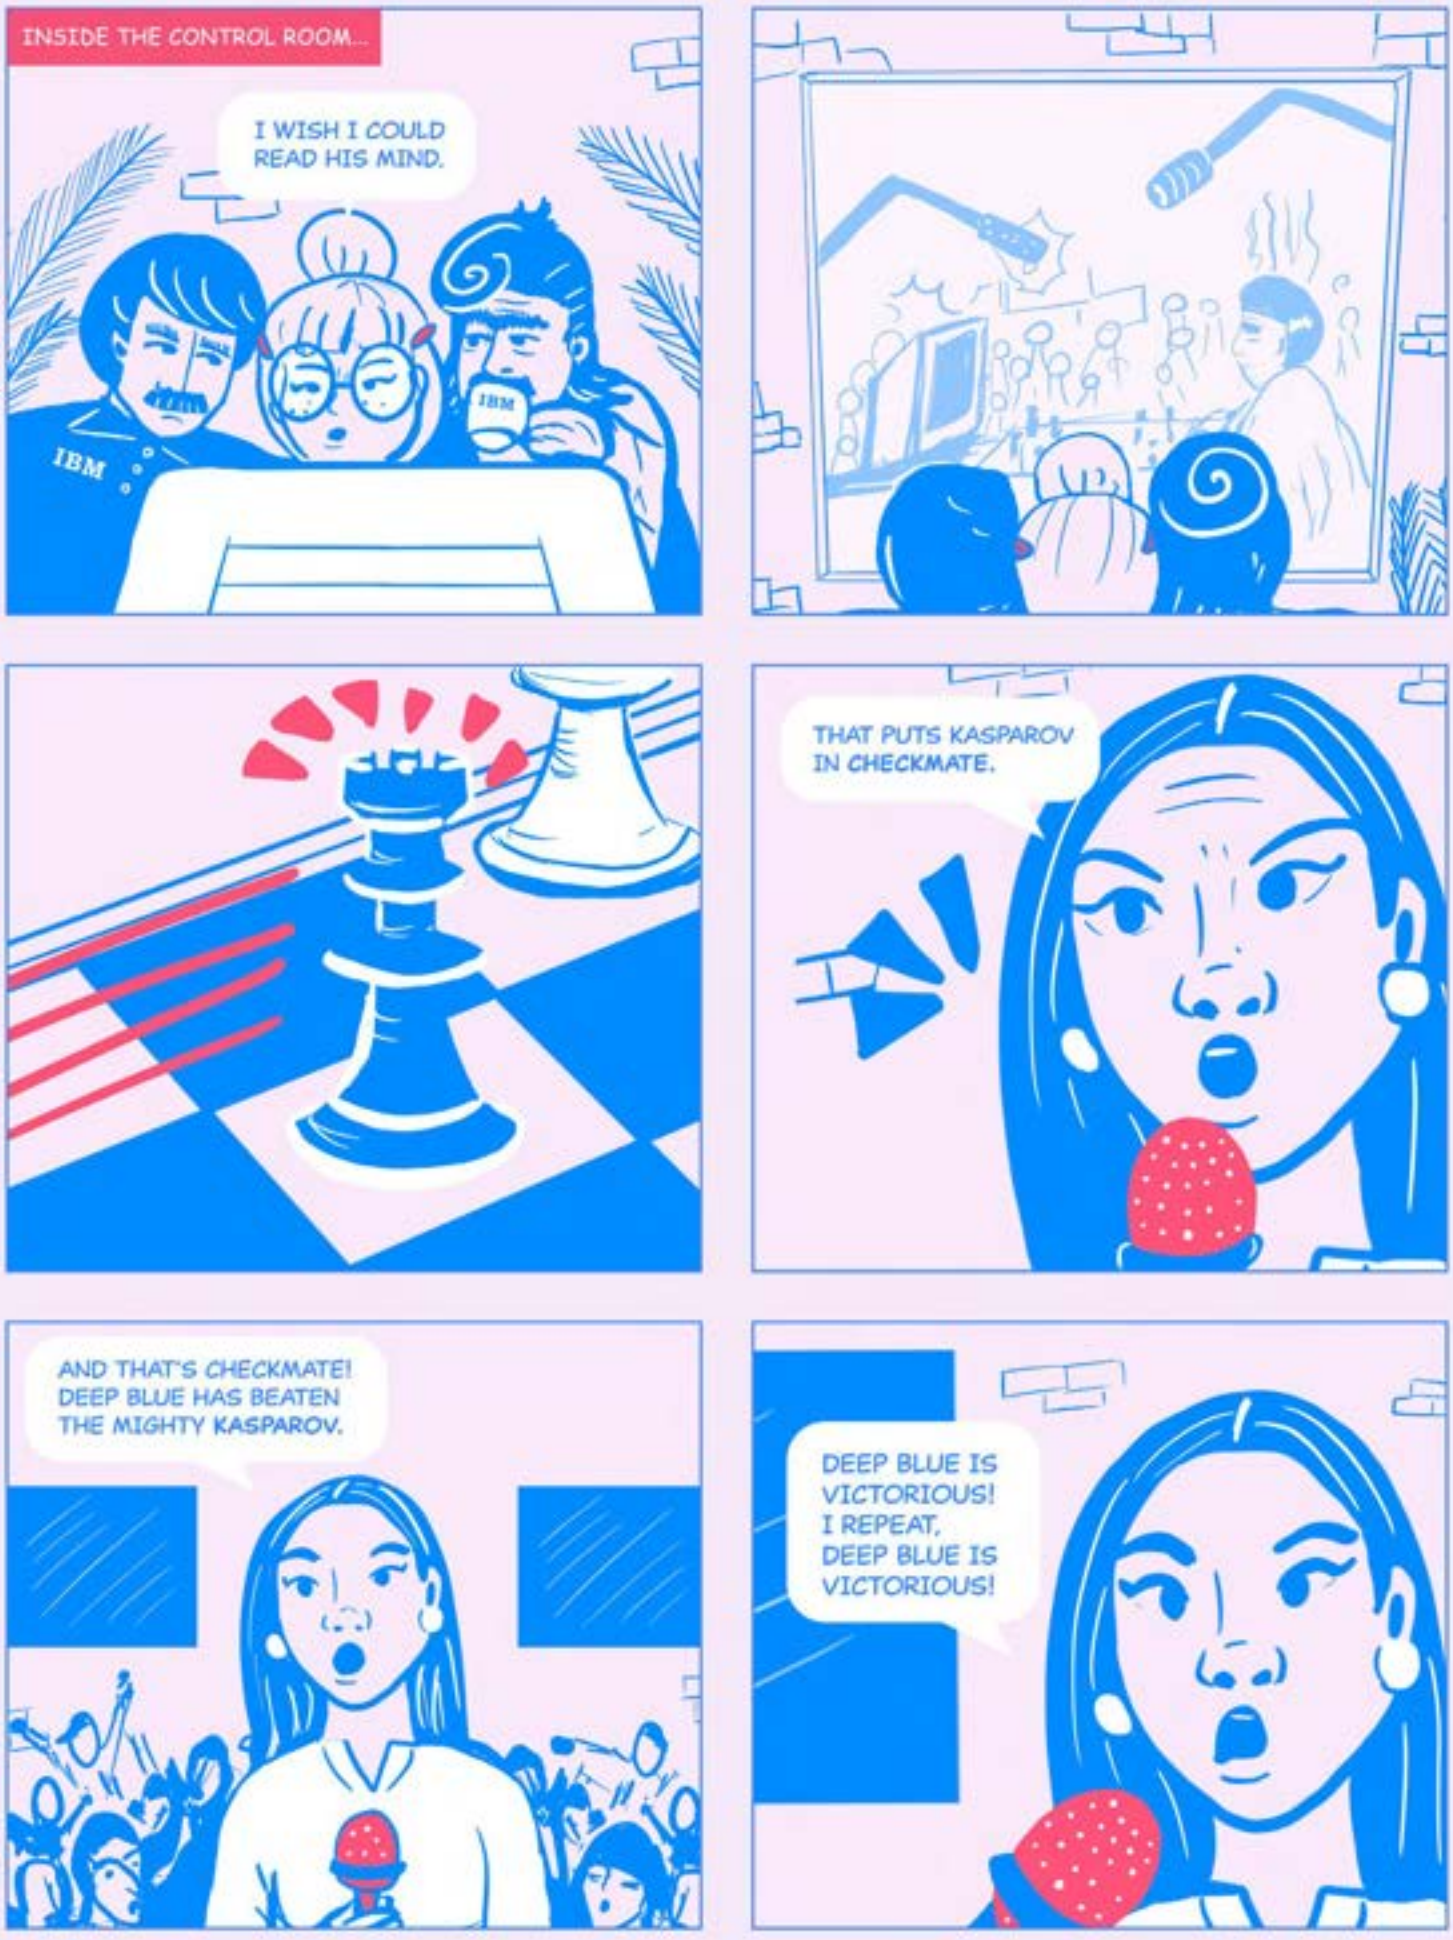
INSIDE THE CONTROL ROOM...
I WISH I COULD READ HIS MIND.
IBM
IBM
THAT PUTS KASPAROV IN CHECKMATE.
AND THAT'S CHECKMATE! DEEP BLUE HAS BEATEN THE MIGHTY KASPAROV.
DEEP BLUE IS VICTORIOUS! I REPEAT, DEEP BLUE IS VICTORIOUS!


OH MY GOSH!


DEEP BLUE IS LIKE AN ALIEN INTELLIGENCE ... LIKE A WALL COMING AT YOU.

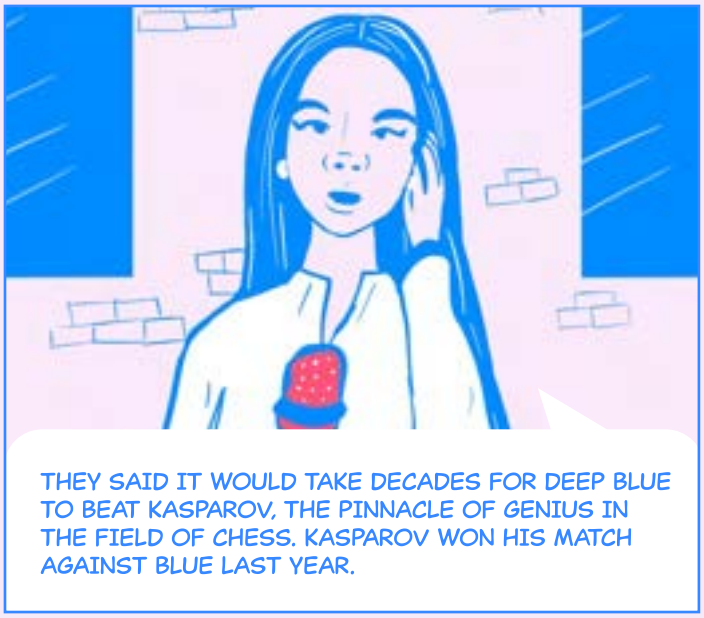
THEY SAID IT WOULD TAKE DECADES FOR DEEP BLUE TO BEAT KASPAROV, THE PINNACLE OF GENIUS IN THE FIELD OF CHESS. KASPAROV WON HIS MATCH AGAINST BLUE LAST YEAR.


ANNIE? WHERE ARE YOU GOING?


I PROMISED I'D HAVE DINNER WITH MY SON. STEAKS AND BANANA SPLITS. HAVE FUN!


KASPAROV SAID TO A REPORTER, "DEEP BLUE IS 'AN ALIEN INTELLIGENCE'." IT MADE ME FEEL WEIRD.

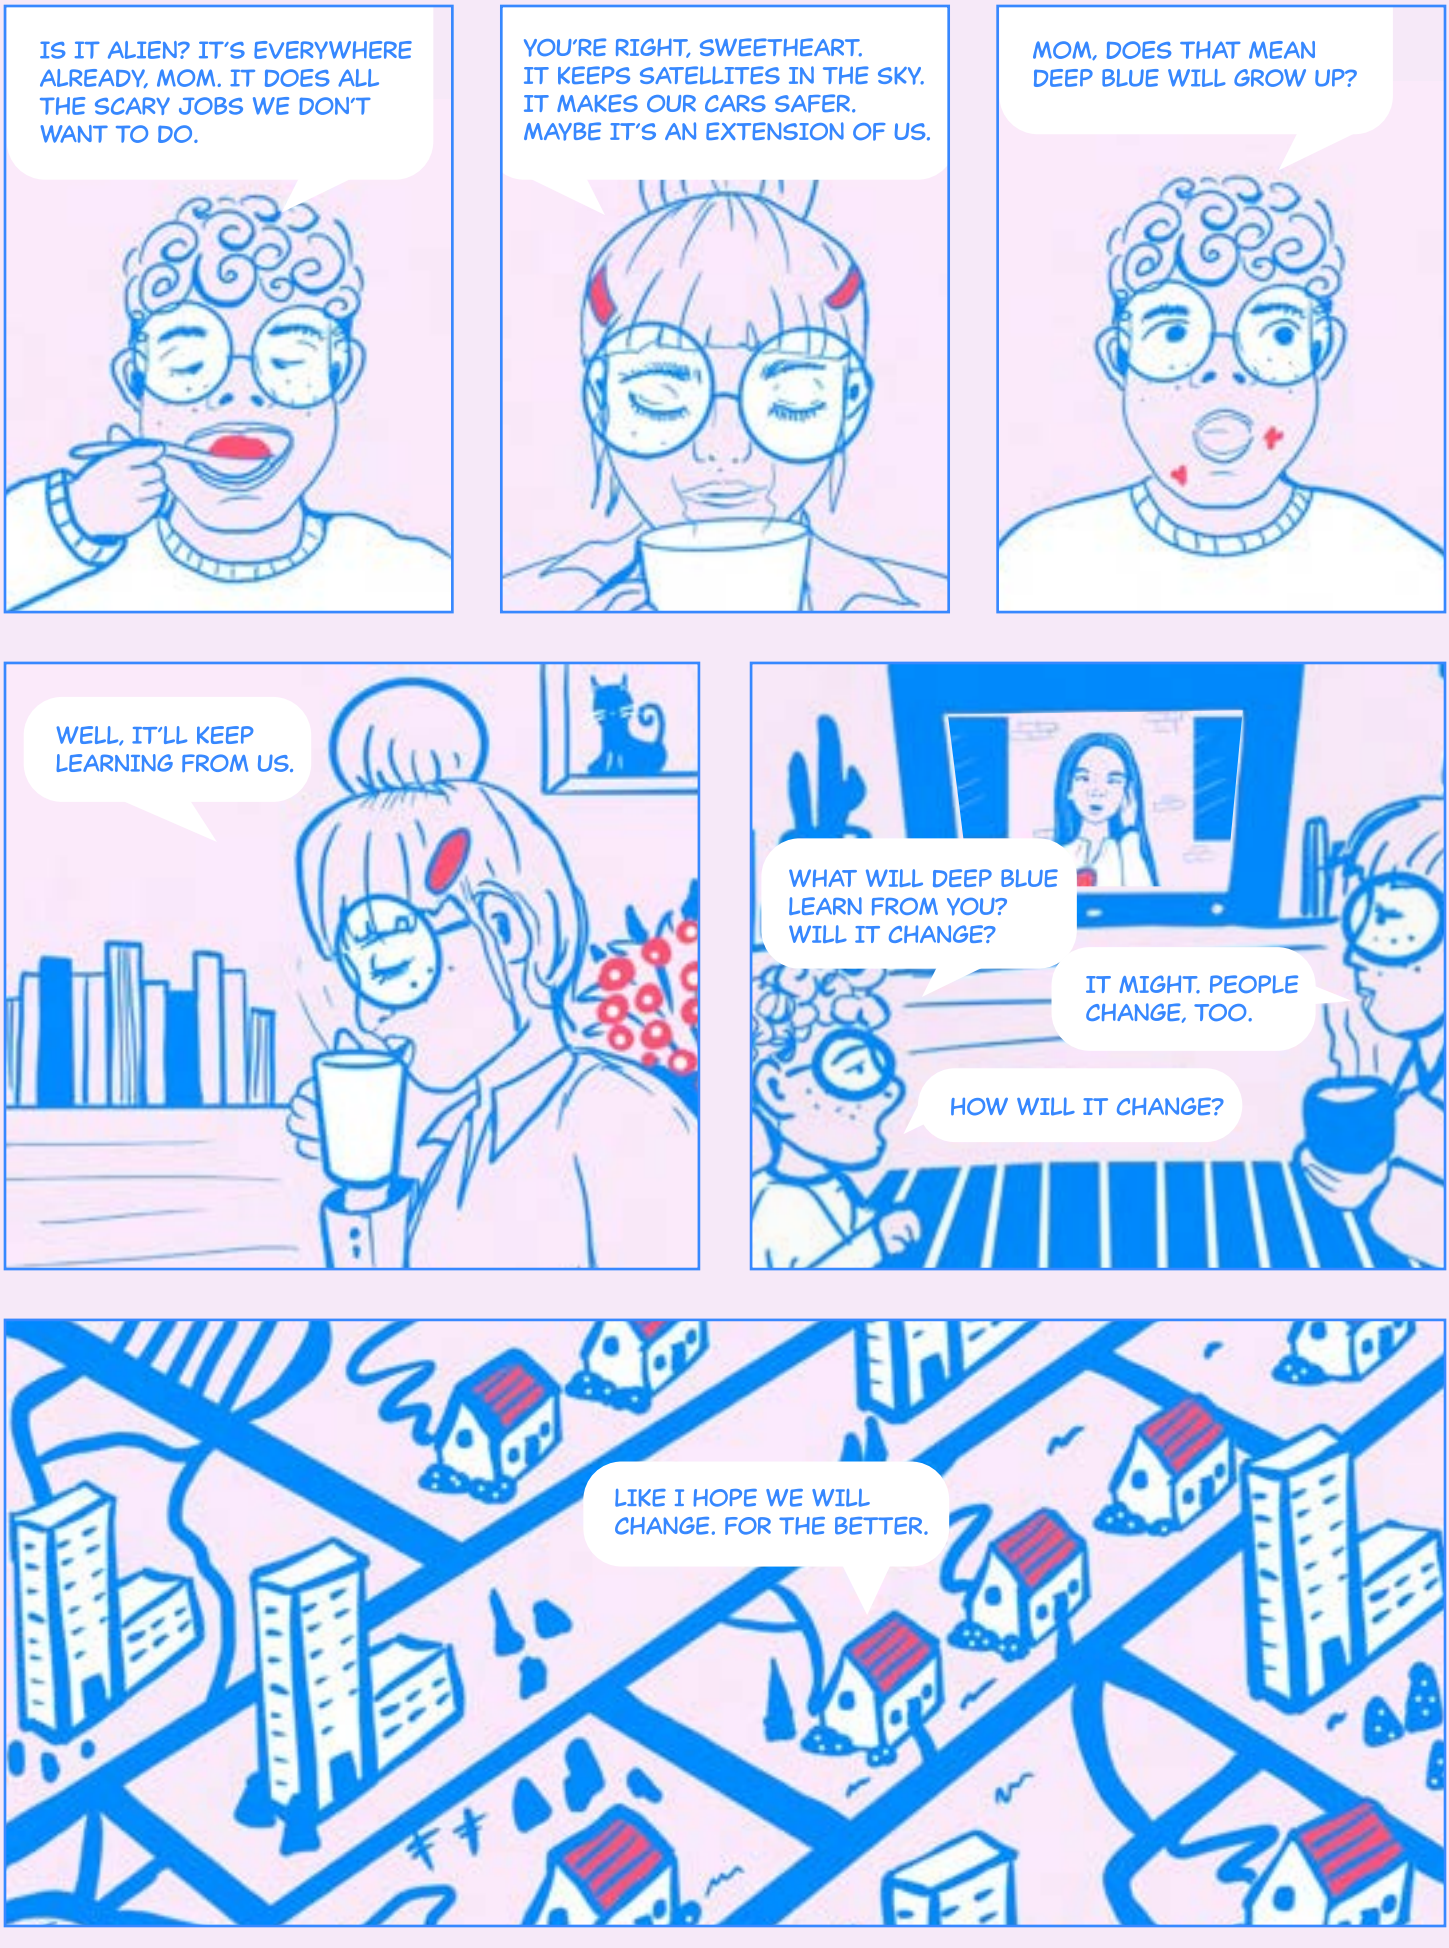
IS IT ALIEN? IT'S EVERYWHERE ALREADY, MOM. IT DOES ALL THE SCARY JOBS WE DON'T WANT TO DO.
YOU'RE RIGHT, SWEETHEART. IT KEEPS SATELLITES IN THE SKY. IT MAKES OUR CARS SAFER. MAYBE IT'S AN EXTENSION OF US.
MOM, DOES THAT MEAN DEEP BLUE WILL GROW UP?
WELL, IT'LL KEEP LEARNING FROM US.
WHAT WILL DEEP BLUE LEARN FROM YOU? WILL IT CHANGE?
IT MIGHT. PEOPLE CHANGE, TOO.
HOW WILL IT CHANGE?
LIKE I HOPE WE WILL CHANGE. FOR THE BETTER.

DIVING IN DATA

How AI Could Help Save the Whales

BY SOFIA OSBORNE

When Jasper Kanes was two years old, their parents took them whale watching for the first time. As is the case for many whale enthusiasts, that first encounter sparked a life-long devotion. "There are photographs of me with my inflatable whale that I was in love with as a tiny, tiny toddler," Kanes tells me. "Apparently I cried when my parents deflated it after giving it to me for Christmas so that we could drive home."

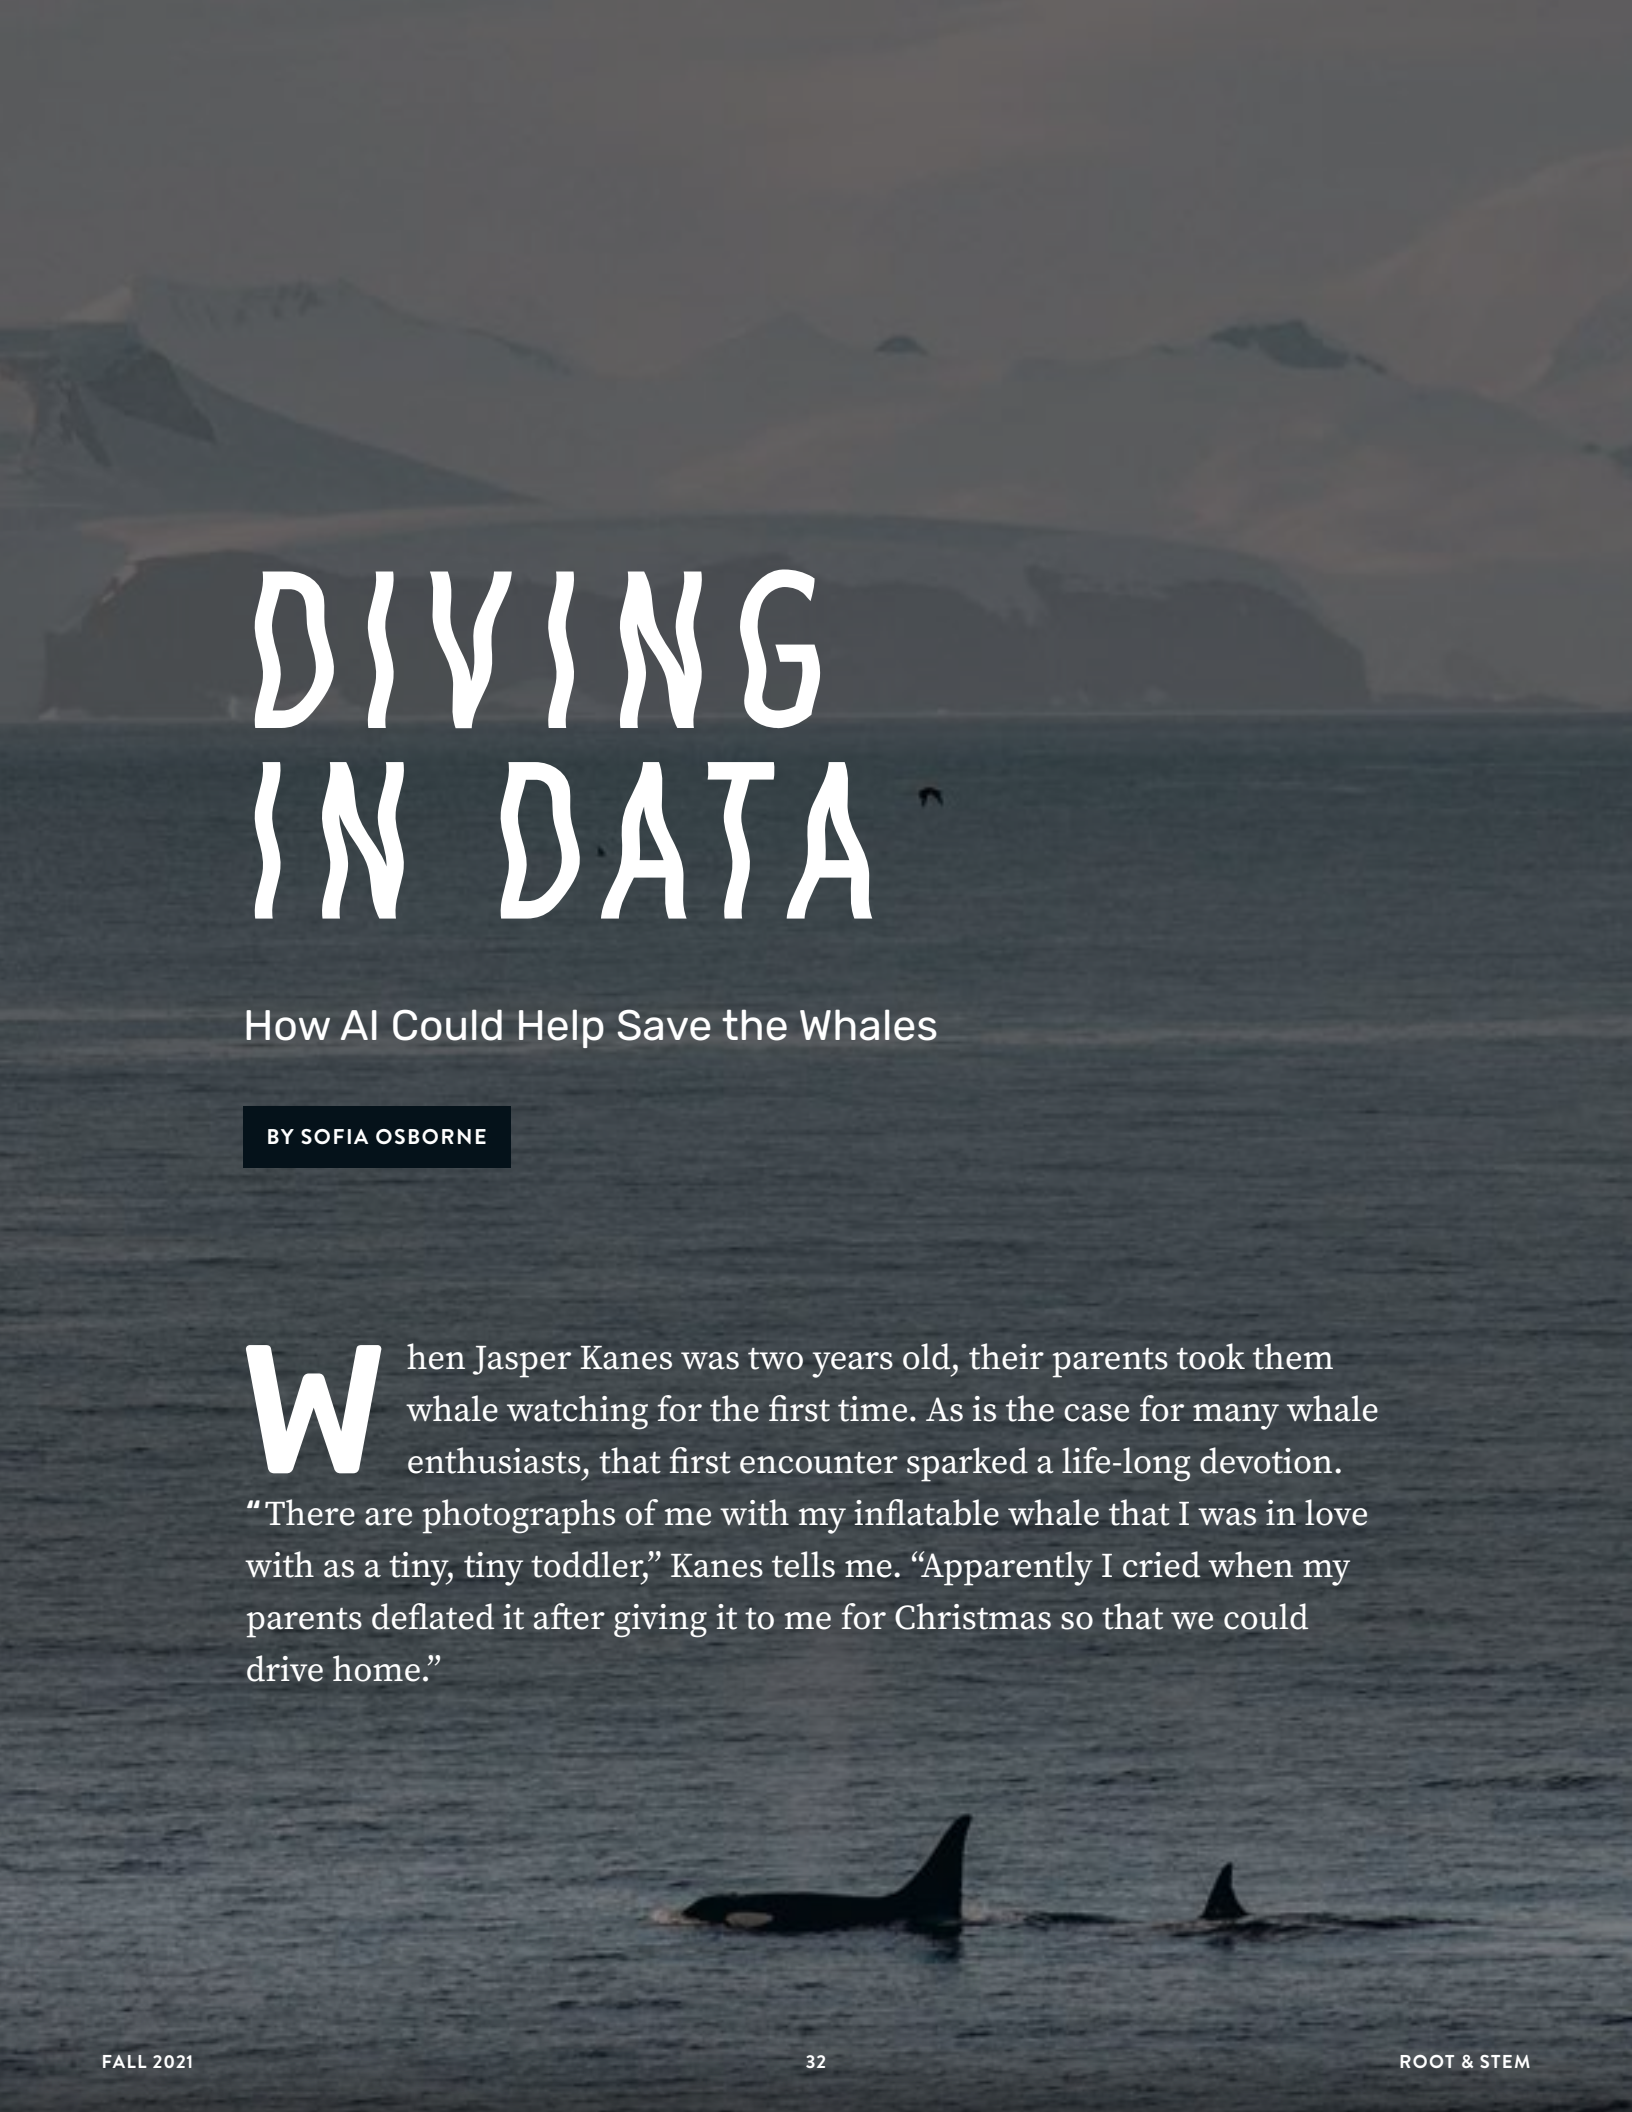

Kanes went on to study whales during their undergraduate degree, starting with research projects for which they would observe whales from boats. "As I learned more about these animals, I realized just how important sound is to their world," Kanes says. "They see with sound. They communicate with sound. It's the most important sense to them. So I thought to myself, if I really want to understand them, I need to start looking at sound."

British Columbia is especially well primed for this type of acoustics-based research, as the province's coast is believed to house more underwater microphones—or hydrophones—for research purposes than anywhere else in the world. These hydrophones lie below the waves, recording the grunts of fish, the calls of whales and the rumbling of ships 24 hours a day, seven days a week, and producing mountains of data for scientists to comb through.

"It's become very inexpensive to collect passive acoustic data," Kanes, now a junior staff scientist for Ocean Networks Canada, explains. "Hard drives are getting really, really cheap. It's getting easier to do autonomous deployments. And so everybody who's collecting passive acoustic data has suddenly arrived at this point where we have gigabytes upon gigabytes—some of us have petabytes of data—that we just don't have the time to look through ... It's very exciting because there's lots to look at, but it's a real challenge to figure out how [to] analyze all of this data? So a big push in the community right now is to develop automated tools."

Such tools utilize artificial intelligence (AI) in similar ways to the speech recognition technology many of us already use, like Apple's Siri or Amazon's Alexa, explains Oliver Kirsebom, a senior staff scientist on the Meridian project at the Institute for Big Data Analytics at Dalhousie University. In similar ways to the algorithm-based speech recognition software, speech recognition tech is being used to recognize and analyze whale calls.

"The idea with these algorithms is that you can train them by feeding them examples of, in our case, sound clips," he says. "So you know if this particular sound clip contains the call of a certain whale or not, and then this neural network learns to determine if that call is present in this sound clip through this training process."

The possibilities of advancing research into, and management of, whales with this technology are particularly vital in a time when so many whales in the Salish Sea—a stretch of ocean that runs from southern BC into Washington State—are threatened by many factors, including vessel strikes, ship noise and a dwindling food supply.

The Salish Sea is famously home to the southern resident killer whale, an orca ecotype that is critically endangered, with only 75 members remaining. But it is also vital habitat for other species and ecotypes: the marine-mammal-eating transient or Biggs killer whale, the northern resident killer whale, and a recent resurgence of the humpback whale. Being able to know where these whales are in real time, based on alerts from AI derived from hydrophones up and down the coast, could mean the prevention of countless whale deaths going forward, as vessels could be warned to slow down or vacate an area if there are whales present.

"Each hydrophone that we have in the water has an acoustic story to tell us about that habitat."

There is also tremendous research value in this archive of acoustic data. Certain whale calls can be indicative of particular behaviours, from humpback whales bubble-net feeding (a feeding method that involves vocalization to coordinate a group of whales to circle a group of fish, creating a "net of bubbles" that disorients the fish and signals through a call when the whales should open their mouth and feed) to orca foraging for fish. Often, pods and even families of orca have their own distinct calls, and researchers are able to identify which groups are using certain areas and during which times of the year.

With so many benefits to be gained from collecting and analyzing this data, many organizations, including NGOs, First Nations, government agencies and private companies, have started to set up their own hydrophone projects.

Tom Dakin, an underwater acoustician who installs hydrophones up and down the coast for clients including the BC Coast-wide Hydrophone Network—a coalition of NGO hydrophone projects—the Department of Fisheries and Oceans and Transport Canada, says there will soon be 21 hydrophones in the water for the network alone, collecting about 183 terabytes of data per year.

"It's incredibly expensive to put things under water for long periods of time," Dakin says. "So the sacrifices and the scrounging that people have done in order to be able to do the research that they're doing is really awe-inspiring. They're shining examples to the rest of us of the dedication that people can have."

Janie Wray, CEO of the North Coast Cetacean Society and Science Director for Orca Lab, helped start the hydrophone network project with Dakin as a way to link the NGO hydrophones along BC's coast. They also approached First Nations like the Heiltsuk of Bella Bella and the Kitasoo/Xai-xais of Klemtu to set up their own hydrophone projects in their territories.

The goal of the network is to bring all of the hydrophone equipment in the water up to the same level so that the data is comparable across the coast. They ultimately plan to create a website that will display an ocean noise health index, where the public will be able to look at noise levels throughout the province. The data collected by the hydrophones will also be available to anyone for their own research projects.

"It's phenomenal how many hydrophones are in the water," Wray says. "And it's great because all these different organizations are listening to whales, we're all gathering data, but we're also gathering data on ocean noise ... Each hydrophone that we have in the water has an acoustic story to tell us about that habitat."

As these hydrophone projects grow, so too do the AI projects that will allow scientists to parse through this important data. However, developing and training these algorithms is not without challenges, most notably a lack of annotated data sets—data that has been labelled by a bioacoustician and can be used to teach the AI what to look for.

"Although these algorithms can be trained to do really impressive things, they are kind of dumb too, in the sense that they require

PHOTO COURTESY OCEAN NETWORKS CANADA

› The ECHO tetrahedral hydrophone array (one of two) goes over the side of E/V Nautilus

just a lot of data before they get to that point," explains Kirsebom. "Although we keep talking about them as artificial intelligence, they're still far away from what a human is capable of, which is to recognize a sound having heard just a few examples of that sound. And with a neural network, we have to show it so many examples before it starts to recognize that sound. And even then it will still make mistakes."

Kanes, who specializes in studying passive acoustics, is one of very few people in the country creating annotated data sets that will be widely available for anyone to use in their projects.

"I'm in a unique position at Ocean Networks Canada to really help move things forward because I'm not being paid to fulfill the requirements of a specific project … I don't have to keep my data confidential. All of the data I work with are available for anyone to use," Kanes tells me. "My mandate is really just to help science move forward. So I'm quite free to produce datasets that are very, very public and to promote them to people to use."

To label the data, Kanes does not usually listen to the recordings; instead, they look at spectrograms—visual representations of the sound data—which show time on the X axis and the frequency of the sound—how high or low the sound is—on the Y axis. With a practised eye, Kanes can recognize the shapes certain calls make in the spectrograms, and more quickly pick out important pieces of data.

"It's a real treat when I have cause to actually listen to the whales, but if I was to listen to the data we collect in real time, it would take far too long," Kanes says. "We have more than 20 hydrophones in the water. They're streaming to us 24 hours a day. There's no way that I could ever catch up with our data collection if I was listening in real time."

Another challenge faced by those creating whale AI is the presence of similar sounds that can lead to false positives—when the AI believes it is identifying a whale call but is actually hearing something like a bird call or the sound of a boat winch.

"There's always going to be new sounds," Kirsebom tells me. "So even when you've trained the network on all sorts of sounds, a new vessel may come by like a ship that makes a sound that the network has never heard before, and again, that's where the networks sometimes struggle in a way that we humans don't struggle … We're not there yet in terms of mimicking our capacity to generalize."

Kanes tells me about one memorable day when a humpback whale call detector made 300 detections in the span of one hour in the Strait of Georgia. Kanes looked closer at the data, and found that it was actually the shaft rub on a boat's propeller creating a very convincing humpback-like sound.

"With all of that context, I can look at it and be like, actually that's shaft rub," Kanes says. "But … it would be so hard to account for all of the factors that go into my decision [that] actually this isn't a humpback. 'Cause it has the right shape, it's the right timbre—which is sort of like the quality of the sound— it's the right frequency, it's the right duration."

In this way, these whale AI networks are still far behind humans when it comes to analyzing the context around a sound.

"That's one thing that we're lacking right now in the algorithms that we've been working with, is that knowledge of what happened … a minute ago or 50 minutes ago," Kirsebom explains. "That's very challenging."

While some may fear AI will make humans obsolete, the reality is that this type of whale AI will probably not be replacing the role of bioacousticians. Instead, the two will work together as a team.

"Now that I'm really in it and I see how a lot of these machine learning algorithms function and how well they perform and all of the nuances that confuse them and all of the contexts that I actually use as a human annotator, I'm starting to realize actually, no, I'm probably not going to put myself out of a job," Kanes says. "These tools will probably make my job easier in that they'll pick out sounds without me having to go through them."

As NGOs continue to work on whale AI projects, tech giants are jumping into the fray too. In 2018, at a workshop in Victoria, Google pitched the idea of developing real-time whale detection to Canada's Department of Fisheries and Oceans (DFO)—the company had already developed similar technology for humpback detection in Washington for the National Oceanic and Atmospheric Administration.

Paul Cottrell, Marine Mammal Coordinator for the DFO, whose team has 20 hydrophones in the water, was excited to partner with Google on this project. They

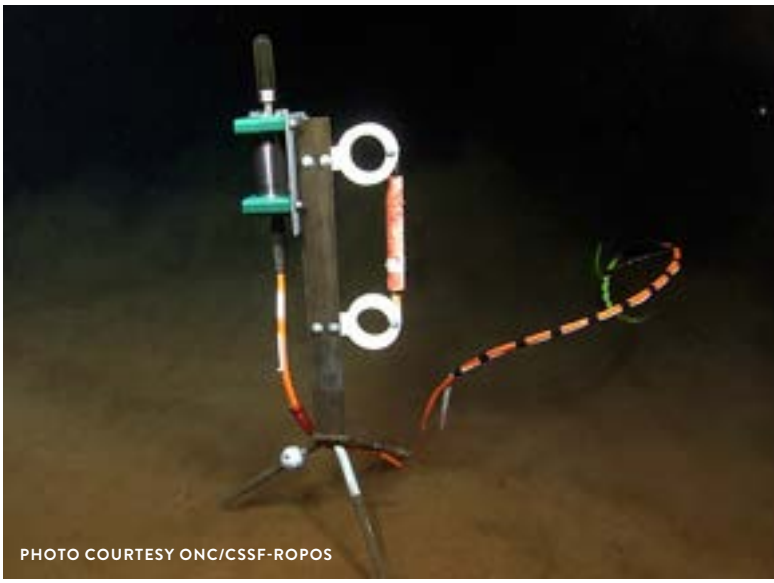

PHOTO COURTESY ONC/CSSF-ROPOS

PHOTO COURTESY BEAMREACH

Naxys Hydrophone rests on a tripod after deployment to Cascadia Basin (top)

also collaborated with Rainforest Connections, a non-profit tech group that previously worked with Google to create a web interface that would alert wildlife officials when chainsaw sounds were detected in the Amazon Rainforest.

"It's just been incredible," Cotrell says. "To actually pay to have this put in place would have been so expensive; if we had to do it without these collaborations, it wouldn't have been possible. So this really only happened because of this Google team of amazing engineers that could put this cutting-edge programming together."

Cottrell and his team have been provided with Google phones which notify them when the AI has picked up an alert at one of their stations. They can then listen to the hydrophones to determine if what the AI has detected is indeed a whale call. The DFO network utilizes experts in the field to identify the calls, pinpointing not just what species and ecotype of whale has been heard, but often even what pod the whale may belong to. This information is fed back into the AI, allowing it to improve constantly.

While the AI is still in its training phase, the hope is that it will soon be used to make real-time management decisions. For example, if the DFO knows that a large pod of southern residents is moving through a pass, they can alert vessels to slow down or avoid the area.

On top of their ability to prevent vessel strikes, Cottrell believes this AI could be a key tool if the area was to experience an oil spill. "If there ever was a catastrophic or a large, significant spill in the Gulf Islands, we would be able to know where the whales are," he says. "If we knew the spill trajectory we could get in front of that and prevent the whales from moving into an area to get exposed. So that's another big piece of this project is to look at this threat minimization from anthropogenic threats."

South of the border, other tech companies have been lending their programming prowess to the fight as well. In the summer of 2019, at a Microsoft hackathon in Seattle, a group of Microsoft employees began working on another algorithm to detect whale calls in real

time. The project, called Orca Hello, has been in development ever since, with the team volunteering in their spare time to design and train the algorithm. It has now been deployed on three hydrophones in Washington run by the non-profit Orca Sound.

"We didn't believe it till we saw it with our own eyes," says Prakruti Gogia, a Microsoft software engineer who works on the HoloLens smart glasses, and a member of the Orca Hello team. "To actually see it work live, that was a really, really incredible experience."

The algorithm was deployed in late September, 2020, and it sent its first real-time alert a few days later. "The first thing it actually got were a couple of false positives," Gogia says.

"Otter squeaks!" Gogia's teammate Akash Mahajan adds excitedly. Mahajan, a Microsoft applied scientist, says that the speech recognition AI he works on for his day job is usually trained on tens of thousands of hours of labelled data. Orca Hello was trained on just fifteen.

"From the beginning … we wanted to create a prototype or a reference design for a system that by design has a feedback loop," he says. "So the way we've set up our system is that we start with some small amount, but we can turn on those hydrophones and when false positives and things come back in they're exactly from what we need, and we can use that to refine and continue to improve."

Soon the AI was picking up the calls of orcas and even humpbacks.

"The very first time we heard humpbacks and the system actually picked up detected humpback calls in the audio, that was an incredible moment because they're really, really distinctive calls, and so echoey," Gogia says. "It was amazing for me. Like that was the wow moment."

Scott Veirs, who co-founded Orca Sound with his father Val, calls these spectacular audio events "live concerts," and views them as a way to forge a connection between humans and whales that will—he hopes—result in an increase in empathy for, and activism around, the plight of the whales, particularly the southern resident killer whales. That is the ethos behind Orca Sound, which utilizes live streaming of their network of three hydrophones to allow anyone to listen in on the underwater world in real time.

"We're also advocates for general open data," Veirs explains. "All of our live streams are archived and those are also available for researchers. And so that's sort of in theme with the way Val—my father—and I started, which is, we believe that education is also transformative, and there's wonder to be felt when you're analyzing the sound after the live concert too."

In a way, Orca Sound is also a citizen science project, encouraging its listeners to help contribute to labelling the data it collects. By visiting **orcasound.net/learn**, anyone can listen to the common sounds found in the Salish Sea and begin to learn how to identify them. In this way, many more of us could contribute to annotating data that could help whale AI move forward.

"The app itself is trying to engage folks in acting on behalf of the whales—that's its primary purpose—but it's also trying to engage them as citizen scientists," Veirs says. "So from a student's perspective, maybe the most interesting thing to do is to use some of our learning resources to teach your students about what killer whales sound like and what other common sounds in the Salish Sea are, and then get them pressing the button."

With so many talented scientists and researchers investing their energy into these AI projects, progress in the field of whale AI is set to grow exponentially.

"Over the many decades that this research has evolved to where we are now … I look at this and think, wow, in 10 years, what are we going to have? Because what's changed in the past 10 years has been amazing," Cottrell says. "The technology, the artificial intelligence, the programming … is just changing so rapidly. And to keep up with that and help the whales is so important."

In addition to its benefits for whale research and protection, these hydrophone and AI technologies could be used to study other parts of the underwater ecosystem as well. As part of his PhD research, Xavier Muoy, a PhD student at the School of Earth and Ocean Sciences at the University of Victoria and an acoustician at JASCO Applied Sciences, is using hydrophone data to identify and monitor the sounds made by different fish species. This research could be particularly beneficial for fish that require conservation strategies, like rockfish.

"As I started to look at data from hydrophones that were collected along the coast … I could still find, pretty much everywhere, these very low frequency sounds that were super grunty and not from a marine mammal," Muoy tells me. "I found out that those are fish, and so if we can hear fish in most of the hydrophones that we have along the coast, that means there's also good potential to monitor fish just using their sounds."

As this field of underwater AI continues to grow on the west coast, the many disparate projects are working to come together and collaborate—after all, the whales don't recognize borders.

"Coast-wide, there's a lot of hydrophones out there and they're all in different networks," Cottrell says. "And I think linking those in as well, is going to be important and I think a challenge going forward, but something that hopefully we can connect … and help the whales." **&**

TECHNOLOGY AND SOCIETY

DATA

Digital Kit

PAST ISSUES

If you missed the Spring 2020, Fall 2020 or Spring 2021 issues of *Root & STEM*, you can find them online at **pinnguaq.com/root-stem**

RESOURCES

We have developed additional digital resources for educators to use in their classrooms—including lesson plans, video tutorials, and handouts. They can all be accessed online via the links that follow.

Math in Nature Video Series

The Math in Nature Video Series is an excellent tool for connecting math and the world around us. Each video explores a different topic and provides examples of math being discovered in nature.
pinnguaq.com/learn/math-in-nature-video-series

The Inuktitut Digital Literacy Game Engine

The Inuktitut Digital Literacy Game Engine hosts three games (and educators can add their own game iterations they create), designed in partnership with Ilitaqsiniq—the Nunavut Literary Council. This learning resource encourages students to practise the Inuktitut language through three simple digital games created in Scratch. These are engaging tools for Inuktitut language educators to expand literacy and vocabulary.
pinnguaq.com/learn/inuktitut-digital-literacy-game-engine

How to Assemble a Computer

This video tutorial guides learners through the parts of a computer and how to build your own.
pinnguaq.com/learn/how-to-assemble-a-computer

Iqalliarluk

Have fun and improve your typing skills by playing Iqalliarluk—ᐃᖃᓗᐊᕐᓗᒃ, a new game developed by Ilitaqsiniq—the Nunavut Literary Council and Pinnguaq!
pinnguaq.com/learn/iqalliarluk

AI BINGO:

Identifying usage of AI in our daily lives.

I woke up this morning to the sound of my **smartphone's alarm clock.**

As I rolled over to grab my phone, I brought it up to my face so it would unlock using **facial recognition.** As I scrolled through my notifications, the **sleep tracking app** that I have connected to my **smartwatch** told me I got about seven-and-a-half hours of sleep last night and only woke up once during the night. I'm lazy in the mornings (and the **autocorrect** and **autofill** settings on my phone are my worst enemies!) so I sent a message replying to my friend's message using **voice to text.** A few more notifications popped up on my phone, **letting me know the weather for my current location** and **suggesting a local breaking news article** that I might have wanted to read. I **sent a snap with a cute dog filter** on it to my best friend Sarah on Snapchat, and noticed from the **location map** that she was in my area!

"Hey Siri, call Sarah"

Instantly, **Siri, the assistant on my phone, called Sarah.** We decided to meet up and stop by a nearby clothing store. I quickly got ready and headed to my car. I also noticed my car telling me that it was 3 degrees Celsius outside, so I sent a text to Sarah.

Hey! You might want to wear a jacket because it might be a bit  for you.

My phone **suggested replacing the word cold with an emoji** so I used that instead. I input the address of where I was going to meet Sarah on **Google Maps** for directions. On the way, I played music from a **recommended playlist on Spotify.** I stopped by my bank along the way and used the drive-through **ATM** to take out some cash for the day. I got to the store before Sarah, and **played a game of Scrabble** on my phone to pass the time.

Sarah arrived, and as we were walking into the store, she told me that someone hit her mom's car while it was parked in the driveway last night and ran off. Luckily, her mom has **a car dash cam that detected the movement** and recorded who hit the car, so they were able to find the culprit. As we finished shopping, Sarah quickly **deposited a mobile cheque with her phone.** The line-ups were long, so we headed over to use the **self-checkout.** Someone's item probably wasn't scanned properly, because as we were walking out of the store, someone in front of us triggered the **anti-theft alarm** and had to return to the store.

I decided to head back to Sarah's place and we **played Just Dance on her Wii.** After that, she decided to finish up an assignment and scanned her paper for mistakes with **spell check** and a **grammar checker.** I also told her to make sure she had used her own words in her paper, as most professors these days run assignments through a **plagiarism checker.** I was a bit bored, so I turned on her TV and **picked a movie that Netflix suggested to me.**

It was a great day!

CREATE YOUR AI BINGO CARD!

Instructions: Read the written excerpt on the previous page briefly and identify any AI usages that are mentioned (they're in bold for an extra hint!) With a pencil, choose eight usages of AI from the text and fill them in on your bingo sheet. Choose one person to be the callout. Now, they call out an AI usage mentioned in the story at random. Each usage that is mentioned counts as a filled-in square if you selected it. To win, you must fill in three boxes in a row, and shout out **Bingo!** Make a copy of these two pages to play with your friends or family to learn more about AI while having fun!

Message to Educators

In the following lesson plans we introduce basic concepts of AI into the classroom in developmentally appropriate, entertaining and solidly educational ways, including hands-on activities that support learning through play.

Short instructional videos are provided where needed to demonstrate activities. They are included in supplementary support material. We hope these lesson plans provide educators with ideas for integrating digital tools into the classroom.

— JENNIE CROSS

Module 1

Grade Level: K to 3
Subjects: Computer Studies, Social Studies
Topic: Computer Fundamentals
Software: Google Teachable Machine (web-based)
Author: Pinnguaq Staff

Introduction to Machine Learning

Introduction

In this lesson, students learn how a computer can be taught to categorize objects. This is something students already have experience with. They can likely separate objects into categories quite easily. They might point out how many animals, vegetables, fruits or vehicles are in an image. For a computer, though, this task can be quite complex.

A computer can be taught to tell objects apart visually if it is given a set of example images. In this class, students first do a non-computer activity that asks a series of questions to help them determine which category an item fits into.

Next, using Google Teachable Machine, the class takes a series of pictures with the teacher's computer, one of each student, as well as some pictures that have no students in them. Then, the computer will be taught to tell when it is looking at a picture of a person and when it is looking at one without any people in it.

Learning Goals

- Introduce **machine learning** and **how it is used in communities**
- Explore **how machine learning uses questions** to determine the category something fits into
- Learn **how a computer detects what is happening in the world around it**
- Introduce **foundation concepts of bias** in machine learning

Vocabulary

The following is a list of potential new words that students should understand.

- **Machine Learning:** teaching computers how to make decisions based on examples it is given
- **Category:** a way of grouping people or things based on their characteristics

Guiding Questions

1. Can computers learn? *(Answer: Yes, today we'll look at how to teach computers.)*
2. Can a computer tell the difference between objects? *(Answer: Yes, but we need to teach it how.)*
3. Can a computer use a camera to recognize a person? *(Answer: Yes! Students may have seen facial recognition used to unlock phones.)*

Curriculum Links

Science and Technology, Applied Design Skills and Technologies
Students will:

- Make a product using known procedures or through modelling of others
- Explore the use of simple, available tools and technologies to extend their capabilities
- With support, follow the steps in a simple, teacher-prepared procedure for a test / experiment
- Ask questions that demonstrate curiosity about the world around them
- Make guesses about possible outcomes of simple procedures

Social Studies
Students will:

- Recognize the causes and consequences of events, decisions or developments

(The Ontario Ministry of Education "Grade 1–8 Science and Technology," Ministry of Education, last modified 2007, accessed June 21, 2021, **edu.gov.on.ca/eng/curriculum/elementary/scientec18currb.pdf**).

(The British Columbia Ministry of Education "Applied Design Skills and Technology," accessed June 21, 2021, **curriculum.gov.bc.ca/curriculum/adst/3/core**).

(The British Columbia Ministry of Education "Social Studies 3," accessed June 21, 2021, **curriculum.gov.bc.ca/curriculum/social-studies/3/core**).

Materials

- Device with internet connection and a camera (computer / tablet / smartphone)

Non-Computer Activity

Is It An Animal?
In this activity, students work in pairs to determine whether their partner has written down the name of an animal, place or thing without asking the question directly. This activity builds students' ability to ask questions that can differentiate items based on their characteristics. Identifying similarities and differences allows both people and artificial intelligence (AI) in order to group or separate things.

Step 1 → Without showing their partner, each learner writes the name of a living thing, plant, food or other object on a piece of paper and places it face down.

Step 2 → Each student asks a maximum of five questions to their partner. The aim is to determine whether their partner has written down the name of an animal or something else, without asking directly, *Is it an animal?*

Example
Partner 1: *Writes down "tree" without showing Partner 2*
Partner 2: *Begins asking their five questions to figure out if Partner 1 wrote down an animal*

Partner 2: *Is it alive?*
Partner 1: *Yes.*

Partner 2: *Does it live in the forest?*
Partner 1: *Yes.*

Partner 2: *Does it have legs?*
Partner 1: *No.*

Partner 2: *Does it move?*
Partner 1: *A little…*

Partner 2: *Does it eat?*
Partner 1: *No.*

Partner 2: *It's not an animal! What were you thinking of?*
Partner 1: *A tree!*

Then the partners switch who is asking questions.

Computer Activity

Introduction and Example

In this activity, the class works together to make a machine learning project. They teach their device to detect when a person is in view of the camera using Google Teachable Machine. The following video shows an example of how this tool can be used.

Video Link: youtu.be/T2qQGqZxkD0

Step 1 → Go to the Google Teachable Machine website: **teachablemachine.withgoogle.com/train**

Step 2 → Click Image Project then choose **Standard Image Model**.

Step 3 → Class 1 and Class 2 are the default names for categories. Change these to **People** and **No people** by clicking on them and typing in the new names.

Step 4 → Next, add sample images to each category (**People** and **No people**) so the computer can learn to tell the two types of content apart. Start by adding photos of people. The more variety of people and types of photo, the better the model will perform. Different lighting, angles and distance can help the model detect people in different situations. Add photos of each student and the teacher using the webcam. If prompted to do so, allow access to the device's camera. Click and hold the **Hold to Record** button to take several photos of each student. In a remote teaching environment, have students submit a photo of themselves using the **Upload** feature. Review the images and remove any that don't have a person in them.

Step 5 → Next, take some sample images for the **No people** category. Again, be sure to take a variety of photos but be certain there are no people in any of the photos! Review your photos and remove any that have people in them.

Step 6 → Now that example images have been saved in each category, the model needs to be trained. This process gives the computer a chance to look carefully at each example

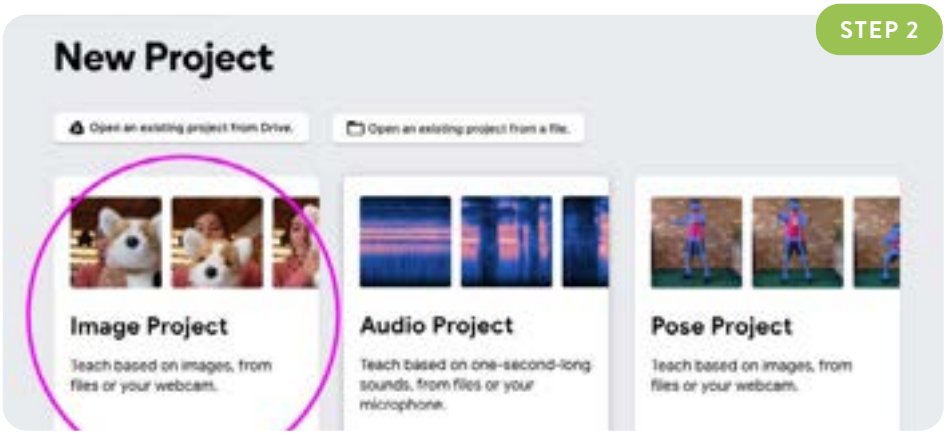

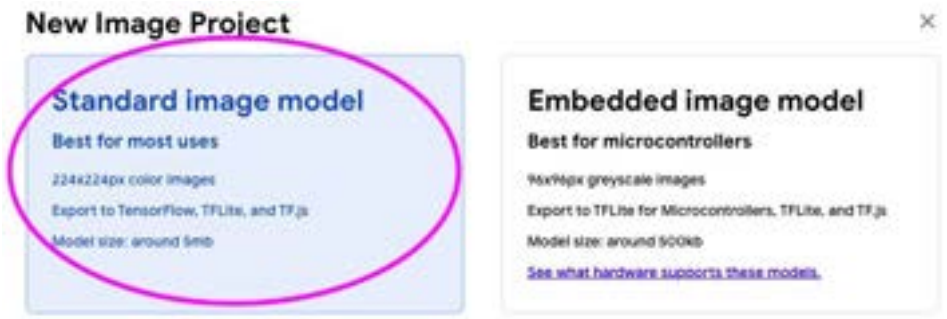

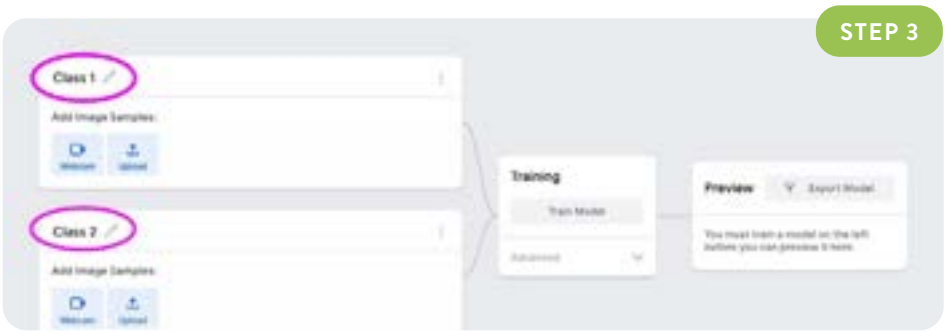

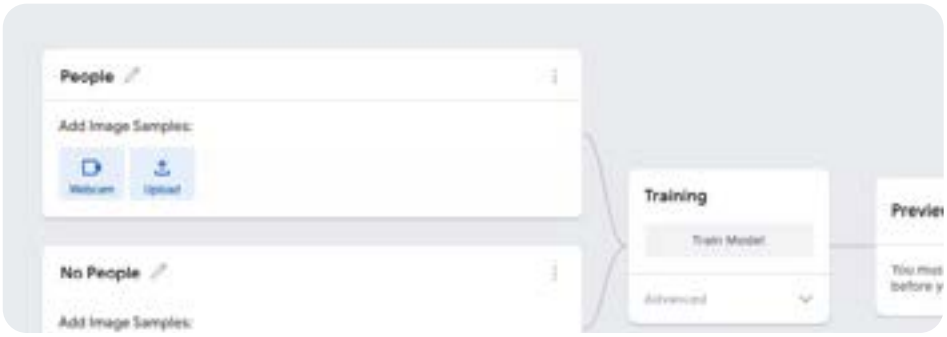

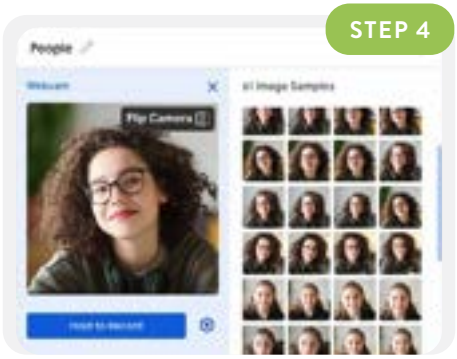

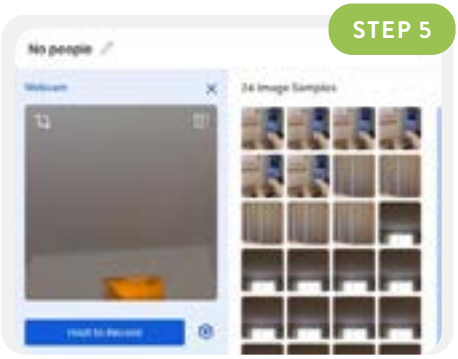

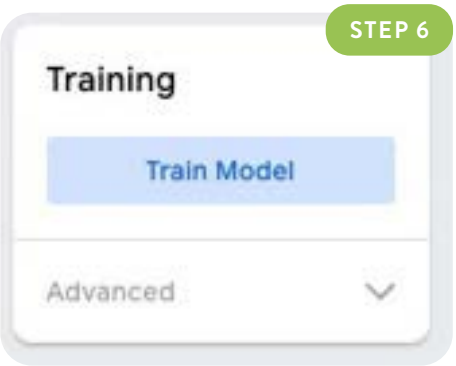

image. This helps the model to be able to guess quickly whether it is being shown a picture of a person or not. Click **Train Model** and wait until the process is complete.

Step 7 → When the model has finished training, a live preview of it working will be displayed. This view can be used to test the model! Test the model with students to see if it replies correctly when there are people or no people in the view of the camera.

Extension Challenges

- Can the model be fooled? How? Try showing the camera only part of a person, or a photo of a person
- What happens if students show one of their hands to the camera, but not their face? How could the model be improved to detect hands better? *(Answer: Adding photos of students' hands to the people category will improve the model)*
- Try starting again to create another model with more categories, or a model that uses poses, or sounds instead of images. How well does this work?

Conclusion

Dog vs. Desk

Quick exercise! What is the difference between a dog and a desk? What do they have in common? Make a chart on the board with a column for **Dogs** and a column for **Desks**. Have students come up with qualities that can be used to identify the similarities and differences. Place a mark to show if the dog or desk has that quality. Encourage discussion about general rules (not all dogs or desks have four legs, not all dogs fetch, etc.). When more information is provided about an animal or object, guesses about which category it fits into are more accurate (see table to the right).

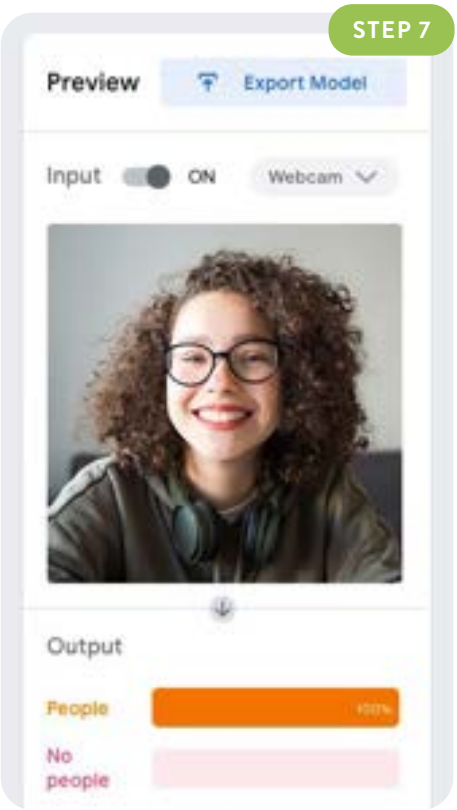

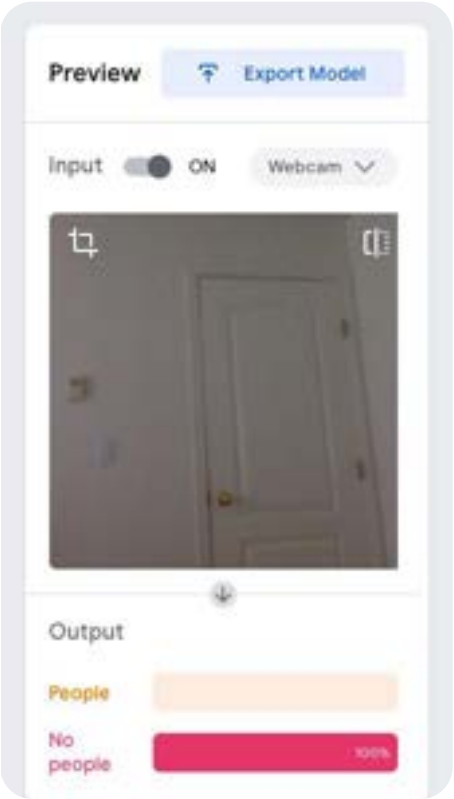

Resources

Additional Resources

- "Machine Learning and Human Bias" from Google: Uses relatively plain language to explore bias in machine learning **youtu.be/59bMh59JQDo**
- "How I'm Fighting Algorithmic Bias" by Joy Buolamwini: This video is perhaps better suited to older students but does a great job of explaining bias in AI **youtu.be/UG_X_7g63rY**

Social Media Resources

- Joy Buolamwini Twitter **@jovialjoy**
- Algorithmic Justice League Twitter **@AJLUnited**

• • •

*We at Pinnguaq want to see the awesome things you're creating! Take a photo or video and share your work with us by emailing **media@pinnguaq.com** or tagging @**pinnguaq** on social media. Don't forget to include the hashtag #LearnWithPinnguaq!*

CODING AND PROGRAMMING

TECHNOLOGY AND SOCIETY

DESIGN

DATA

Qualities	Dogs	Desks
Has four legs	×	×
Can be indoors or outdoors	×	×
Many different sizes	×	×
Eats food	×	
Good to eat off		×
Fetch	×	
Requires assembly		×

Module 2

Grade Level: 4 to 6
Subject: Computer Studies, English
Topic: Computer Fundamentals
Software: Kuki.ai (web-based)
Author: Briana Brownell and Jessa Gamble

Turing Tests

Introduction

When artificial intelligence (AI) advances to a certain level, it may be possible to have conversations with machines that feel like speaking to a human. This exercise tests the current state of AI to see whether we can be fooled.

A Turing test is one way of testing the intelligence of a computer. The 20th-century mathematician Alan Turing pointed out that if a computer can have a conversation with a person and make them feel as if they are speaking to another human, we might say it has a certain kind of intelligence. When a person can no longer tell the difference between a conversation generated by the computer and a conversation generated by a person, the computer "passes the Turing test." In this exercise, students meet a prize-winning chatbot who is famous for her human-like responses.

Viewing Activity
The following is a two-minute video that explains Alan Turing's theory.
youtu.be/bCUNdu_OqUY

(Jr. Mojo. "ALAN TURING!—Fantastic Facts." December 18, 2018. YouTube Video, 2:22 minutes).

Learning Goals

- **Introduce artificial intelligence** using chatbots
- **Critically evaluate the chatbot** to compare to human conversation

Vocabulary

- **Chatbot:** a computer program designed to simulate conversation with human users
- **Turing test:** a test for intelligence in a computer, requiring that a human being should be unable to distinguish the machine from another human being by using the replies to questions put to both

Guiding Questions

- Has anyone had an experience with artificial intelligence? What was it? *(Students may be familiar with digital assistants like Siri, Alexa, Cortana et al.).*
- How do computers help humans? *(What are some tasks that computers are good at?)*
- What are examples of tasks that are a challenge for computers to perform? *(Computers are great at doing fast calculations, but not so good at completing creative tasks independently, imitating people, not always good at recognizing speech etc.)*

Curriculum Links

This module aligns with provincial and territorial Language Arts curricula, providing an opportunity for students to question what we read, hear and view to be engaged citizens. Students will also have the opportunity to demonstrate willingness to learn new technologies as needed.

Materials

- Device with internet connection and a camera (computer / tablet / smartphone)
- Free account set-up with **chat.kuki.ai**

Non-Computer Activity

Human Conversation
This exercise introduces the unique characteristics of human-to-human conversation.

Have students form pairs and chat with each other for five minutes. One person is the interviewer and asks a series of questions (ten questions is sufficient) about a topic of their choosing. The other student answers each question. For example, what is the opposite of hungry? How does a car work? Why do dogs fly? Have students write down the questions and the answers.

The questions and answers will then be used in the next activity.

Computer Activity

Talk to Kuki
Step 1 → Type a Document with Your Questions and Answers
Use a word processor (e.g. Google Docs, Microsoft Word, Notepad) to type up the questions and answers from the Human Conversation activity.

Step 2 → Chatbot Conversation
Go online and speak to Kuki, an AI chatbot. (Note to teachers: Set up an account at **chat.kuki.ai** so Kuki does not interrupt the conversation to suggest doing so.)

Ask Kuki the same questions asked in the earlier Human Conversation.

Copy and paste the transcript of the conversation with Kuki into your document, entering the answers either before or after those from the human conversation.

Step 3 → Disguise the AI Conversation
Format the partner and Kuki conversations so that they are identical by making the font and font size the same. Do not change any of the actual text. The goal is to make telling the

difference between the responses as hard as possible for the next team.

Step 4 → Be the Judge
Have each pair of students look at the transcripts from three other pairs and see whether they can determine which was the human conversation.

Conclusion

Discussion
Question 1: Were any of the Kuki conversations human enough to fool us?
Why or why not?

Question 2: What is a mind?
How can you tell if you're speaking to another person?

Optional: Schedule a video call with Kuki with the class. Please note, at the time of printing, this feature was currently in beta testing and requires a short application process: **kuki.ai/video-call**

Resources

Additional Resources
- Article—"A Holographic Future: Exploring the Capabilities of the Looking Glass Holographic Display" **pinnguaq.com/holographic-future-exploring-capabilities-looking-glass**

Social Media Resources
Kuki.ai is on these social media platforms:
- **tiktok.com/@kuki_ai**
- **youtube.com/c/kuki_ai**
- **instagram.com/kuki_ai**

• • •

*We at Pinnguaq want to see the awesome things you're creating! Take a photo or video and share your work with us by emailing **media@pinnguaq.com** or tagging @**pinnguaq** on social media. Don't forget to include the hashtag #LearnWithPinnguaq!*

CODING AND PROGRAMMING

TECHNOLOGY AND SOCIETY

DESIGN

DATA

Notes:

Module 3

Grade Level: 7 to 8
Subject: Computer Studies
Topic: Computer Fundamentals
Software: Recog (web-based)
Author: Briana Brownell and Jessa Gamble

Machine Learning: Supervised and Unsupervised

Introduction

There are many ways to build machine learning programs. In this lesson, students explore two of the basic types of machine learning: unsupervised learning and supervised learning.

Learning Goals

- Engage in hands-on activities to **train a machine in two ways**
- Understand the difference between **supervised** and **unsupervised learning**
- Understand **real-world applications** of machine learning

Vocabulary

- **Supervised Learning:** A type of machine learning task. If we know what we want the computer to be able to identify, or the categories we want it to distinguish among, this is the type of machine learning we use. The initial data is labelled so the machine can learn which features are shared by data in certain categories. It then generates an algorithm to categorize future, unlabelled data. Humans can check the accuracy of these predictions and get feedback to improve the algorithm. Right now, this is the most common type of machine learning.
- **Unsupervised Learning:** Another type of machine learning task. This approach is ideal for identifying general similarities and useful patterns—even the kind that humans don't tend to notice. Humans do not guide machines to group information, and patterns emerge on their own.

Guiding Questions

- Can computers learn? How do they learn? *(Answer: Yes, today we'll look at how to teach computers)*
- Can a computer tell the difference between objects? How do they know how to do this? *(Answer: Yes, but we need to teach it using examples, or create algorithms to identify patterns to recognise objects)*
- Why do computers need people to help them learn? *(Answer: We teach computers to identify patterns and recognize objects)*

Curriculum Links

This module aligns with provincial and territorial Computer Science curricula, providing an opportunity for students to explore emerging areas of computer science research. In addition, through the Social Studies curriculum, students have the opportunity to assess the significance of an event (in this case, AI) on people, places and developments at this time in history.

Materials

- Device with internet connection and a camera (computer / tablet / smartphone)

Computer Activity

Activity 1: Recog (5 minutes)
In this activity, students work as a class on making a machine learning project that teaches a device to identify their handwriting. They will use Google Teachable Machine for the project. The video below shows an example of how this tool can be used.

Recog is an open-source educational program that can learn to recognize individual handwritten characters. It uses a two-layer neural network to learn new characters and refine its knowledge of those it already knows. Students can hand-train Recog in two ways—supervised or unsupervised—to recognize digits, vowels, arrows in different directions and symbols of their own making. It takes about three examples of each figure for the program to recognize a set of around 10 patterns. Use the following web app, **neilfraser.github.io/Recog**. Students use a pre-trained neural network to see how well it can recognize their handwriting (see the character recognition image below).

1. Draw a letter, number or symbol with the mouse.
2. Press **Digitize** to see how the software makes sense of what has been drawn and turns it into a series of squares.
3. Select the model **0123456789** from the drop-down menu. This is a neural network that has already been trained to recognize digits from 0 through 9.
4. Press **Recognize** to have Recog compare the new symbol with the models it has stored and return the probability that that the new symbol is a match for each of the elements in the program's memory.
5. If it identifies the symbol incorrectly, click **Learn** and tell it the correct answer.
6. Repeat three times for each number. How well did it do?
7. Try entering a completely different symbol. What does the neural network do?

Activity 2: Supervised Learning—How's Your Handwriting? (25 minutes)
Each student trains their model with their own handwriting. They then experiment with their handwriting on their neighbour's system and see that it performs more poorly at first, due to differences in letter formation, then watch it learn (see the character recognition image on the following page).

1. Select **New** from the Network dropdown and **Supervised** from the Learning dropdown.

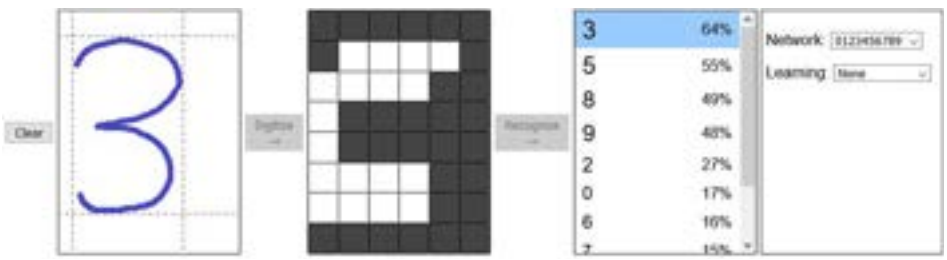

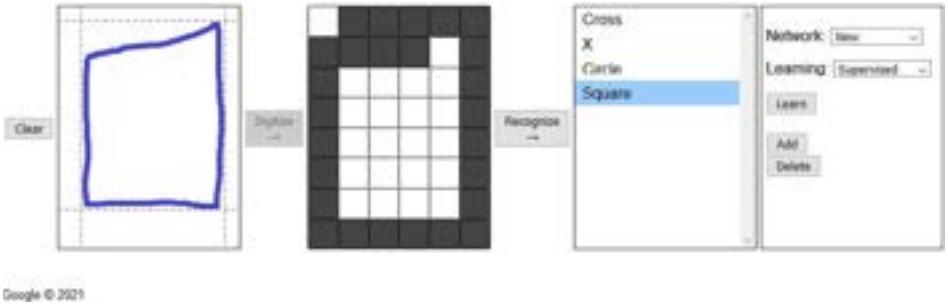

Google © 2021

2. Press **Add** to add a new character for the neural network to learn. Repeat for all the characters you would like the neural network to learn.
3. Draw a letter, number or symbol with the mouse.
4. Press **Digitize** to see how the software makes sense of what has been drawn and turns it into a series of squares.
5. Select the correct symbol and press **Learn**.
6. Repeat all chosen symbols three times, then test your partner's computer.

Activity 3: Unsupervised Learning—Break the Code (25 minutes)

Teams of two face off. Which team can train their AI and break the code the fastest? Each team is given a secret five-letter word.

1. One person on each team of two sets the program to unsupervised and submits all their symbols again. This time, the program will group them into types and create numbered categories.
2. Continue to train the machine to label the first letter of the secret word as "Type 1," the second as "Type 2," and so on. Train it as much as you want, but remember, you're being timed!
3. When you think the program has had sufficient training, it's your partner's turn to decode your words. The team to decode all their partner's words first is the winner.

Optional language integration: Train the learning program on other languages such as Inuktitut or French!

Conclusion

Class Discussion (10 minutes)

There are many applications that need to recognize characters—in fact, this ability was one of the very first use cases of machine learning. For example, this is essentially the type of algorithm Canada Post uses to scan and recognize addresses on envelopes. It's also similar in concept to what Google uses to identify copyrighted music, or the way self-driving cars are trained to recognize stop signs. The current systems are just much more sophisticated.

1. When Canada Post developed its system to scan handwritten as well as typed addresses on envelopes, what do you think the challenges might have been?
2. In the Recog game, you drew each character in a box. What advantages do you think the box might give a character-recognition algorithm?
3. Can you think of ways a system like this might be used to improve peoples' lives or make some tasks easier?
4. Do you think the system would be more effective in reading cursive or printed handwriting? Do you think fonts impact the ease of recognition?
5. Are there ways we might not want sthis type of machine learning to be used?

Resources

Additional Resources

- Briana Brownell's newly-released five-minute animated TED-Ed video "What Is Machine Learning?" explains the styles of machine learning. **ted.com/talks/briana_brownell_how_does_artificial_intelligence_learn** (Briana, Brownell. "How Does Artificial Intelligence Learn?." TED, March 2021. YouTube video, 4:38 minutes)

Social Media Resources

- Briana Brownell Twitter **@UncannyLearning**

• • •

*We at Pinnguaq want to see the awesome things you're creating! Take a photo or video and share your work with us by emailing **media@pinnguaq.com** or tagging @**pinnguaq** on social media. Don't forget to include the hashtag #LearnWithPinnguaq!*

CODING AND PROGRAMMING

TECHNOLOGY AND SOCIETY

DESIGN

DATA

Notes:

Module 4

Grade Level: 9 to 12
Subject: Computer Studies
Topic: Computer Fundamentals
Author: Briana Brownell and Jessa Gamble

Generative Adversarial Networks (GANs): Artistic Creation with an AI Collaborator

Introduction

A generative adversarial network (GAN) is a pair of duelling systems playing a game. The discriminative network has been trained to recognize, or discriminate among, types of data. The generative network creates novel examples and tries to fool the discriminating network into accepting them. Since their invention by a Google team in 2014, GANs have been used for everything from hyper-realistic images for computer games to high-energy physics simulations. This lesson provides online resources that allow students to create music, visual art or other creative output using GANs.

Learning Goals

- Create an **audio-visual project** using artificial intelligence (AI)
- Morph an **image into a style of artwork** using AI

Vocabulary

- **Generative Adversarial Networks (GANs):** an approach to generative modelling using deep learning methods, such as convolutional neural networks. Generative modelling is an unsupervised learning task in machine learning that involves automatically discovering and learning the regularities or patterns in data in such a way that the model can be used to generate or output new examples that could have been plausibly drawn from the original dataset.

Guiding Questions

- Do you know any examples of AI being used for creative projects?
- Is there a place for AI in creating visual artworks?

Curriculum Links

This module aligns with provincial and territorial Computer Studies curricula as it provides opportunities for students to explore emerging areas of computer science research. In addition, through the Visual Arts curriculum, students have the opportunity to explore and experiment with a variety of media and materials, as well as traditional and/or emerging technologies, tools and techniques, and apply them to produce artworks.

Materials

- Device with internet connection (computer / tablet / smartphone)
- Sign up for an account at **deepdreamgenerator.com**

Computer Activity

Activity 1: Make an Audio-Visual Masterpiece with an AI Collaborator

1. Select a song file or other audio file (an mp3 of under 15 minutes in length).
2. Go to **www.wzrd.ai**. Name your project and enter your song.
3. Select a visual theme.
4. Chose the images you want to appear or have them chosen for you by clicking **Skip This Step**.
5. Autofill images and rearrange them on the timeline as desired.
6. Render video.
7. Watch!

Activity 2: Creating Using AI – Deep Style Transfer

Students use their own photos to work in Deep Style Transfer or Deep Dream

1. Sign up for an account at **deepdreamgenerator.com**
2. Choose a base image. This image defines the base structure of your art piece. Royalty-free images are available from several websites, such as **pexels.com** or **pixabay.com**, or you can use your own image.
3. Choose a style image. This image defines the style of your art piece. You can use a painting or drawing of your own or artworks that are in the public domain: **nga.gov/open-access-images.html**
4. Examine the resulting image. Can you see which parts of the style image were transferred to your base image?
5. Using the same base and style images, try changing the style scale up or down and repeat the process.
6. Repeat the process using different style images.

Examples

Base:

Style:

Courtesy National Gallery of Art, Washington, DC

Final:

Base:

Style:

Courtesy National Gallery of Art, Washington, DC

Final:

Activity 3: Deep Dream

1. Sign up for an account at **deepdreamgenerator.com**
2. Choose an image to "dream." Royalty-free images are available from several websites, such as **pexels.com** or **pixabay.com**, artwork that is in the public domain can be found at **nga.gov/open-access-images.html** or you can use your own image.
3. Click **Settings** and choose a layer.
4. Examine the resulting image. Are you surprised with the patterns that result?
5. Using the same base image, choose a different layer and repeat the process.

Examples

Conclusion

Discussion Activity

What is different about this kind of art-making process as compared to painting or drawing? What are its advantages and disadvantages?

In terms of art-making, do you believe AI is:

- a tool used by a human?
- a partner of a human?
- a muse to a human?
- a creator enabled by a human?

Resources

Additional Resources

- Video about Harold Cohen's *Aaron* **youtu.be/MwHQx9BrHQc**
- Harold Cohen Biography (Tate Modern) **tate.org.uk/art/artists/harold-cohen-925**

Social Media Resources

- Xander Steenbrugge, Creator of WZRD Twitter **@xsteenbrugge**
- Deep Dream Generator Twitter **@DeepDreamGen**

• • •

*We at Pinnguaq want to see the awesome things you're creating! Take a photo or video and share your work with us by emailing **media@pinnguaq.com** or tagging @**pinnguaq** on social media. Don't forget to include the hashtag #LearnWithPinnguaq!*

CODING AND PROGRAMMING

TECHNOLOGY AND SOCIETY

DESIGN

DATA

LET'S MAKE COOL THINGS

Pinnguaq has proudly partnered with organizations from coast to coast to coast on exciting learning projects, like our newest collaboration, Iqalliarluk–ᐃᖃᓗᓕᕆᔨᒃ, an Inuktitut typing game created with Ilitaqsiniq–the Nunavut Literacy Council. Find this game, along with other great resources at **ilitaqsiniq.ca**.

If you want to collaborate with us on projects that improve digital skills and STEAM learning for residents of the Yukon, Northwest Territories or Nunavut, we're excited to work with you!

Let's chat! Send us a message at steam@pinnguaq.com

PINNGUAQ

LEARNING RESOURCES

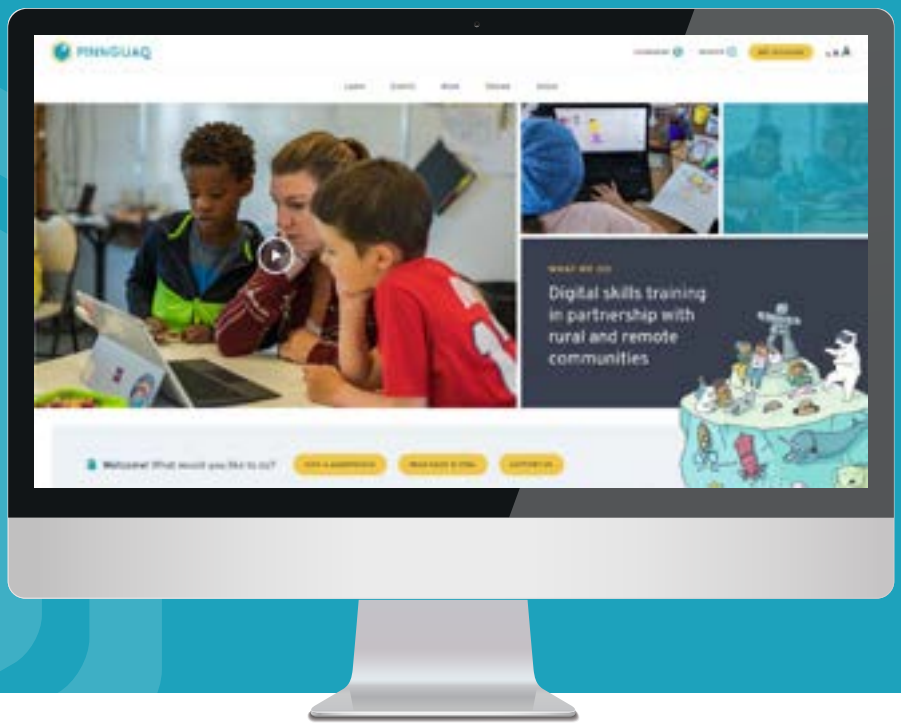

LOOKING FOR MORE CONTENT?

pinnguaq.com

Pinnguaq's online learning space encourages educators and students to create and use digital technology to inspire and solve challenges in their community

If you would like to receive **free** print copies of *Root & STEM*, please sign up at **pinnguaq.com/root-stem-subscription**

Introducing the

Root & STEM Podcast

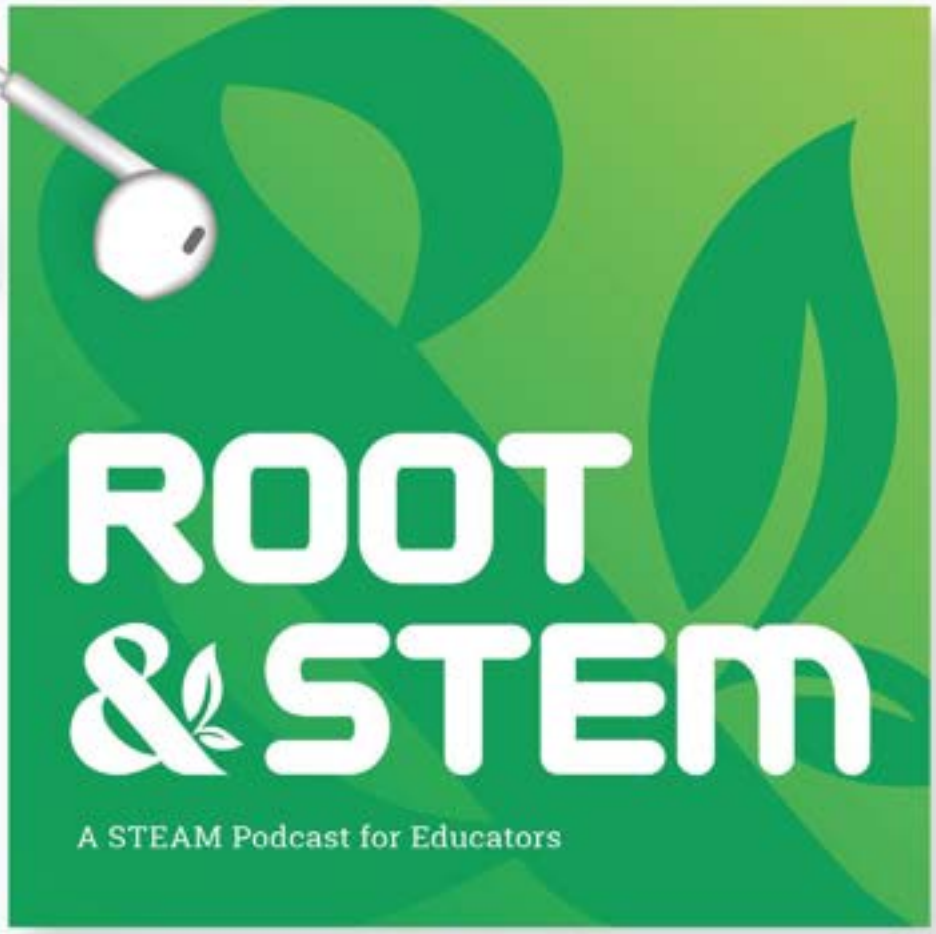

Expanding on the print edition of Root & STEM, the Root & STEM podcast introduces listeners to communities, projects and activities exploring STEAM education in a digital world.

FIND THE ROOT & STEM PODCAST ON THE WEB AT

pinnguaq.com/learn/the-root-stem-podcast

Listen on Apple Podcasts

Listen on Google Podcasts

LISTEN ON Spotify

Interested in following the journey of Root & STEM? Follow Pinnguaq on Facebook, Instagram and Twitter to share thoughts about the podcast using the hashtag **#RootandSTEMPod**.