

Machine teaching: How people's expertise makes AI even more powerful

April 24 2019, by Jennifer Langston



Microsoft Corporate Vice President for Business AI Gurdeep Pall talks at a recent conference about autonomous systems solutions that employ machine teaching. Credit: Dan DeLong for Microsoft

Most people wouldn't think to teach five-year-olds how to hit a baseball

by handing them a bat and ball, telling them to toss the objects into the air in a zillion different combinations and hoping they figure out how the two things connect.

And yet, this is in some ways how we approach [machine learning](#) today—by showing [machines](#) a lot of data and expecting them to learn associations or find patterns on their own.

For many of the most common applications of AI technologies today, such as simple text or image recognition, this works extremely well.

But as the desire to use AI for more scenarios has grown, Microsoft scientists and product developers have pioneered a complementary approach called machine [teaching](#). This relies on people's expertise to break a problem into easier tasks and give machine learning models important clues about how to find a solution faster. It's like teaching a child to hit a home run by first putting the ball on the tee, then tossing an underhand pitch and eventually moving on to fastballs.

"This feels very natural and intuitive when we talk about this in human terms but when we switch to machine learning, everybody's mindset, whether they realize it or not, is 'let's just throw fastballs at the system,'" said Mark Hammond, Microsoft general manager for Business AI.

"Machine teaching is a set of tools that helps you stop doing that."

Machine teaching seeks to gain knowledge from people rather than extracting knowledge from data alone. A person who understands the task at hand—whether how to decide which department in a company should receive an incoming email or how to automatically position wind turbines to generate more energy—would first decompose that problem into smaller parts. Then they would provide a limited number of examples, or the equivalent of lesson plans, to help the machine learning algorithms solve it.

In supervised learning scenarios, machine teaching is particularly useful when little or no labeled training data exists for the machine learning algorithms because an industry or company's needs are so specific.

In difficult and ambiguous reinforcement learning scenarios—where algorithms have trouble figuring out which of millions of possible actions it should take to master tasks in the physical world—machine teaching can dramatically shortcut the time it takes an intelligent agent to find the solution.

It's also part of larger goal to enable a broader swath of people to use AI in more sophisticated ways. Machine teaching allows developers or subject matter experts with little AI expertise, such as lawyers, accountants, engineers, nurses or forklift operators, to impart important abstract concepts to an [intelligent system](#), which then performs the machine learning mechanics in the background.

[Microsoft researchers](#) began exploring machine teaching principles nearly a decade ago, and those concepts are now working their way into products that help companies build everything from intelligent customer service bots to [autonomous systems](#).

"Even the smartest AI will struggle by itself to learn how to do some of the deeply complex tasks that are common in the real world. So you need an approach like this, with people guiding AI systems to learn the things that we already know," said Gurdeep Pall, Microsoft corporate vice president for Business AI. "Taking this turnkey AI and having non-experts use it to do much more complex tasks is really the sweet spot for machine teaching."



Mark Hammond, Microsoft general manager for Business AI and former Bonsai CEO, developed a platform that uses machine teaching to help deep reinforcement learning algorithms tackle real-world problems. Credit: Dan DeLong for Microsoft

Today, if we are trying to teach a machine learning algorithm to learn what a table is, we could easily find a dataset with pictures of tables, chairs and lamps that have been meticulously labeled. After exposing the algorithm to countless labeled examples, it learns to recognize a table's characteristics.

But if you had to teach a person how to recognize a table, you'd probably start by explaining that it has four legs and a flat top. If you saw the person also putting chairs in that category, you'd further explain that a

chair has a back and a table doesn't. These abstractions and feedback loops are key to how people learn, and they can also augment traditional approaches to machine learning.

"If you can teach something to another person, you should be able to teach it to a machine using language that is very close to how humans learn," said Patrice Simard, Microsoft distinguished engineer who pioneered the company's machine teaching work for Microsoft Research. This month, his team moves to the Experiences and Devices group to continue this work and further integrate machine teaching with conversational AI offerings.

Millions of potential AI users

Simard first started thinking about a new paradigm for building AI systems when he noticed that nearly all the papers at machine learning conferences focused on improving the performance of algorithms on carefully curated benchmarks. But in the real world, he realized, teaching is an equally or arguably more important component to learning, especially for simple tasks where limited data is available.

If you wanted to teach an AI system how to pick the best car but only had a few examples that were labeled "good" and "bad," it might infer from that limited information that a defining characteristic of a good car is that the fourth number of its license plate is a "2." But pointing the AI system to the same characteristics that you would tell your teenager to consider—gas mileage, safety ratings, crash test results, price—enables the algorithms to recognize good and bad cars correctly, despite the limited availability of labeled examples.

In supervised learning scenarios, machine teaching improves models by identifying these high-level meaningful features. As in programming, the art of machine teaching also involves the decomposition of tasks into

simpler tasks. If the necessary features do not exist, they can be created using sub-models that use lower level features and are simple enough to be learned from a few examples. If the system consistently makes the same mistake, errors can be eliminated by adding features or examples.

One of the first Microsoft products to employ machine teaching concepts is Language Understanding, a tool in Azure Cognitive Services that identifies intent and key concepts from short text. It's been used by companies ranging from UPS and Progressive Insurance to Telefonica to develop intelligent customer service bots.

"To know whether a customer has a question about billing or a service plan, you don't have to give us every example of the question. You can provide four or five, along with the features and the keywords that are important in that domain, and Language Understanding takes care of the machinery in the background," said Riham Mansour, principal software engineering manager responsible for Language Understanding.

Microsoft researchers are exploring how to apply machine teaching concepts to more complicated problems, like classifying longer documents, email and even images. They're also working to make the teaching process more intuitive, such as suggesting to users which features might be important to solving the task.

Imagine a company wants to use AI to scan through all its documents and emails from the last year to find out how many quotes were sent out and how many of those resulted in a sale, said Alicia Edelman Pelton, principal program manager for the Microsoft Machine Teaching Group.



Microsoft Corporate Vice President for Business AI Gurdeep Pall talks at a recent conference about autonomous systems solutions that employ machine teaching. Credit: Dan DeLong for Microsoft

As a first step, the system has to know how to identify a quote from a contract or an invoice. Oftentimes, no labeled training data exists for that kind of task, particularly if each salesperson in the company handles it a little differently.

If the system was using traditional machine learning techniques, the company would need to outsource that process, sending thousands of sample documents and detailed instructions so an army of people can attempt to label them correctly—a process that can take months of back and forth to eliminate error and find all the relevant examples. They'll

also need a machine learning expert, who will be in high demand, to build the machine learning model. And if new salespeople start using different formats that the system wasn't trained on, the model gets confused and stops working well.

By contrast, Pelton said, Microsoft's machine teaching approach would use a person inside the company to identify the defining features and structures commonly found in a quote: something sent from a salesperson, an external customer's name, words like "quotation" or "delivery date," "product," "quantity," or "payment terms."

It would translate that person's expertise into language that a machine can understand and use a machine learning algorithm that's been preselected to perform that task. That can help customers build customized AI solutions in a fraction of the time using the expertise that already exists within their organization, Pelton said.

Pelton noted that there are countless people in the world "who understand their businesses and can describe the important concepts—a lawyer who says, 'oh, I know what a contract looks like and I know what a summons looks like and I can give you the clues to tell the difference.'"

Making hard problems truly solvable

More than a decade ago, Hammond was working as a systems programmer in a Yale neuroscience lab and noticed how scientists used a step-by-step approach to train animals to perform tasks for their studies. He had a similar epiphany about borrowing those lessons to teach machines.

That ultimately led him to found [Bonsai](#), which was [acquired by Microsoft](#) last year. It combines machine teaching with deep reinforcement learning and simulation to help companies develop

"brains" that run autonomous systems in applications ranging from robotics and manufacturing to energy and building management. The platform uses a programming language called [Inkling](#) to help developers and even subject matter experts decompose problems and write AI programs.

Deep reinforcement learning, a branch of AI in which algorithms learn by trial and error based on a system of rewards, has successfully outperformed people in video games. But those models have struggled to master more complicated real-world industrial tasks, Hammond said.

Adding a machine teaching layer—or infusing an organization's unique subject matter expertise directly into a deep reinforcement learning model—can dramatically reduce the time it takes to find solutions to these deeply complex real-world problems, Hammond said.

For instance, imagine a manufacturing company wants to train an AI agent to autonomously calibrate a critical piece of equipment that can be thrown out of whack as temperature or humidity fluctuates or after it's been in use for some time. A person would use the Inkling language to create a "lesson plan" that outlines relevant information to perform the task and to monitor whether the system is performing well.

Armed with that information from its machine teaching component, the Bonsai system would select the best reinforcement learning model and create an AI "brain" to reduce expensive downtime by autonomously calibrating the equipment. It would test different actions in a simulated environment and be rewarded or penalized depending on how quickly and precisely it performs the calibration.

Telling that AI brain what's important to focus on at the outset can short circuit a lot of fruitless and time-consuming exploration as it tries to learn in simulation what does and doesn't work, Hammond said.

"The reason machine teaching proves critical is because if you just use reinforcement learning naively and don't give it any information on how to solve the problem, it's going to explore randomly and will maybe hopefully—but frequently not ever—hit on a solution that works," Hammond said. "It makes problems truly solvable whereas without machine teaching they aren't."

Provided by Microsoft

Citation: Machine teaching: How people's expertise makes AI even more powerful (2019, April 24) retrieved 25 April 2024 from <https://phys.org/news/2019-04-machine-people-expertise-ai-powerful.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.