SN
SCIENCE NEWS MAGAZINE
SOCIETY FOR SCIENCE & THE PUBLIC
MAY 11, 2019 & MAY 25, 2019

New Hominid Species Named

Insect Farming as Big Business

Dead Pig Brains Show Signs of Life

Space Travel Revs Up Immune System

NEXT LEVEL

AI goes gamer to learn how humans think

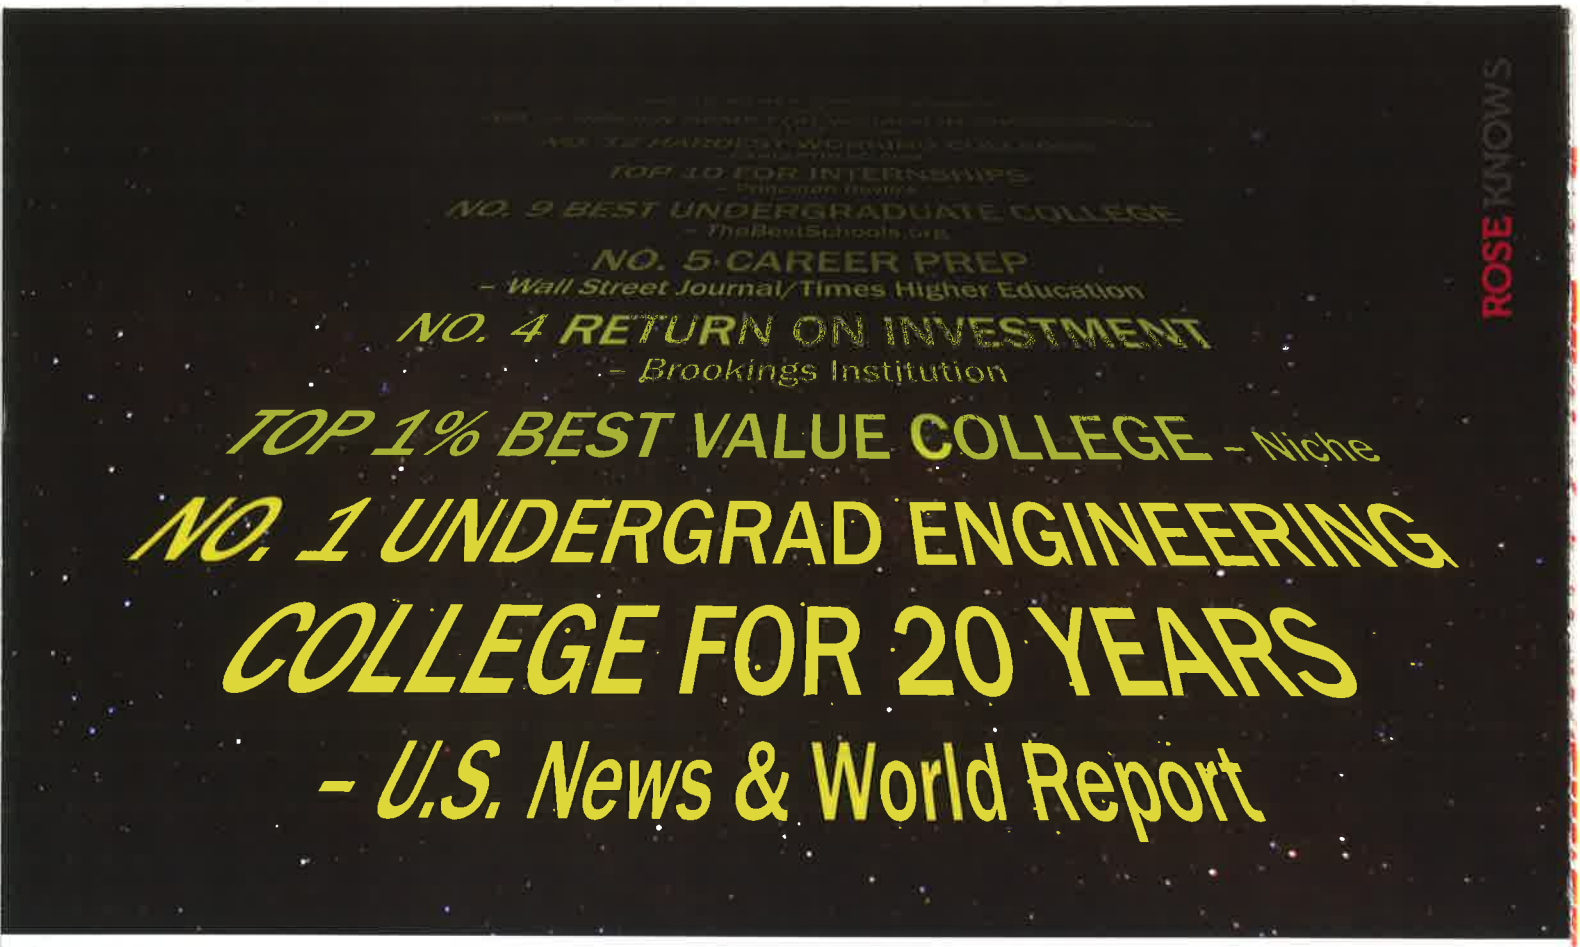

THE BEST STEM SCHOOL IN THE GALAXY.

Our rankings could go on forever, but here are the facts. We offer the best undergraduate engineering education in the country (with 20 straight years of being ranked #1 by *U.S. News & World Report* to prove it). Our students are driven and collaborative, and our faculty are challenging yet supportive. When you arrive at Rose-Hulman, you'll jump right into tackling global STEM challenges.

We're Rose. It's what we do.
Sound like you?

ROSE-HULMAN
INSTITUTE OF TECHNOLOGY

rose-hulman.edu/galaxy

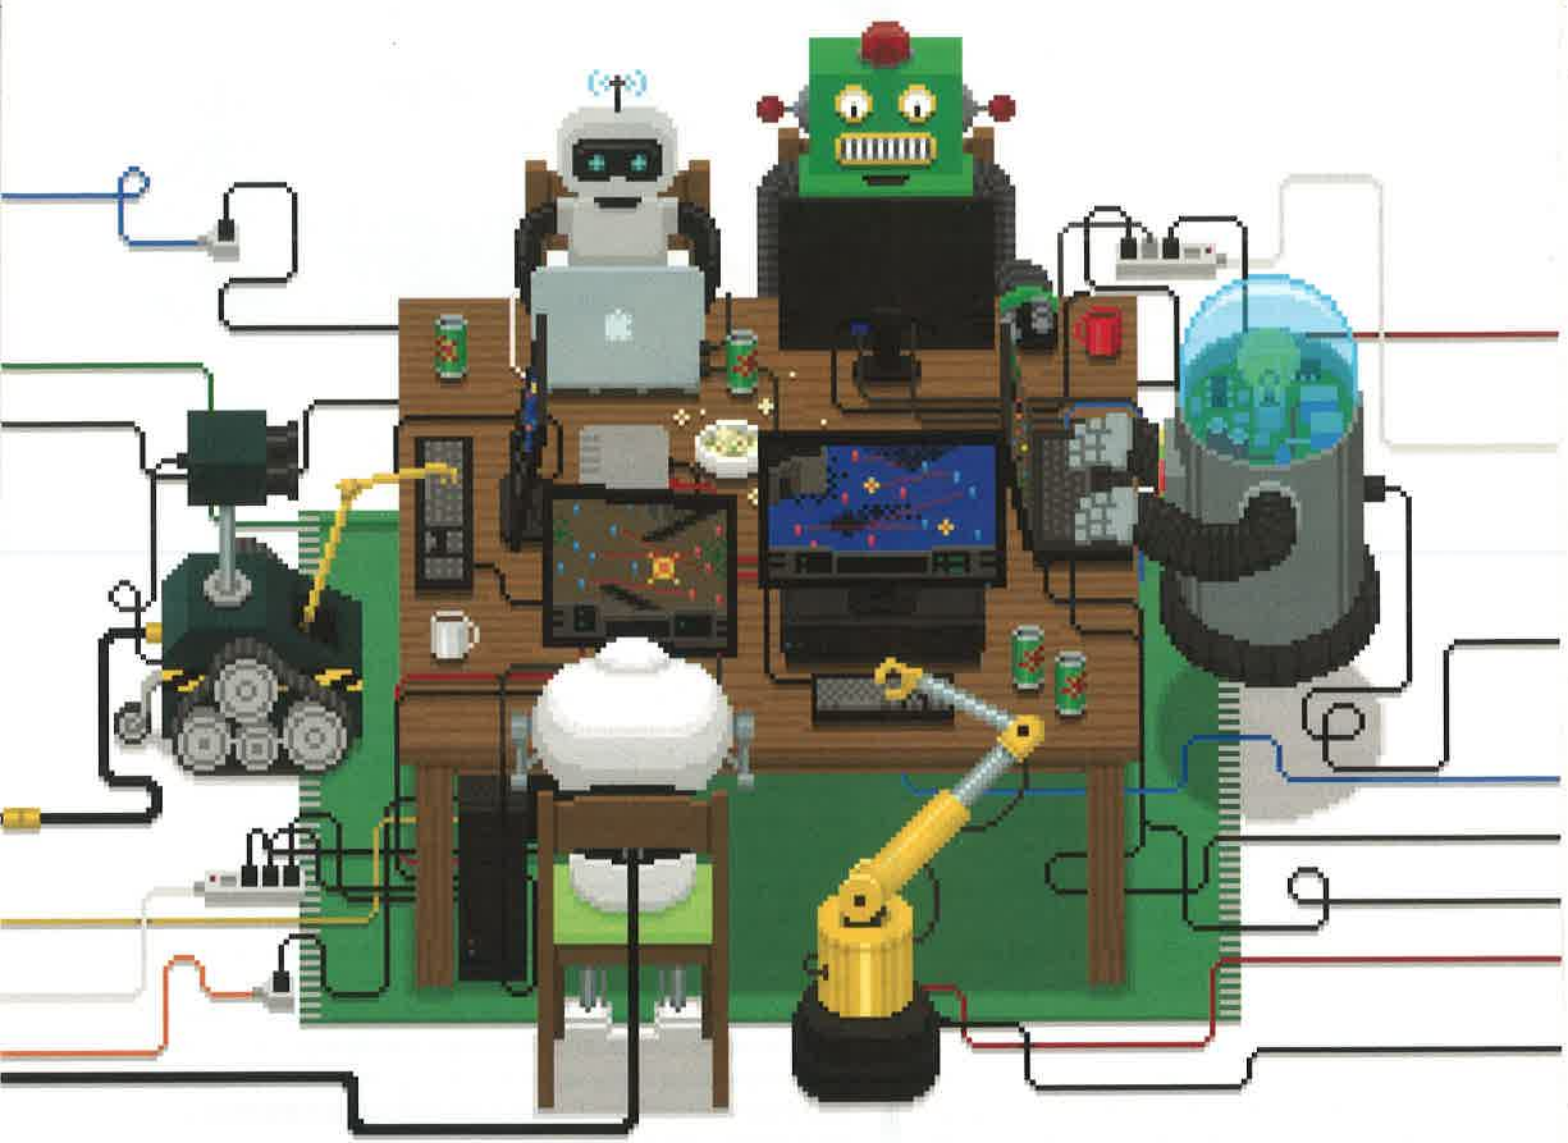

AI at Play

When computers take a seat at the game table, they learn real-world skills **By Maria Temming**

WALTER NEWTON

The list of edible insects stretches to over 2,000, and the bugs' demands vary, says entomologist Åsa Berggren of the Swedish University of Agricultural Sciences in Uppsala. Some need meat to eat. Some, such as termites, produce methane, which is a powerful greenhouse gas.

Even in the same insect species at the same stage of life, environmental impact numbers can differ, Berggren and colleagues pointed out in the February issue of *Trends in Ecology & Evolution*. The few shreds of data on farming the house cricket (*Acheta domesticus*) suggest that the animals convert feed to body weight at rates varying from 1.6 (good) to 4.5 (almost as resource-hogging as pigs). What would be really sad, she says, is insect farming missing this chance to take the green and thrifty route.

A better way to judge farm impacts is to sum them all in a life cycle analysis, says Afton Halloran, a sustainable-food consultant based in Copenhagen. However, she knows of fewer than a dozen such analyses, and results differ.

Overall, the analyses suggest that what farmers choose to feed their insects is a big deal, Halloran says. A farmer might nudge the world in a better direction by feeding the microherd something that would normally be thrown away, for example, feeding chicken manure to black soldier fly larvae. Serving a different waste product, say beet pulp, also puts food waste to good use, but it racks up considerable environmental impacts. The larvae don't grow well on the pulp, so they need a lot of it, plus heating to keep their metabolisms revved.

Also grown

Some farms have already diversified beyond crickets. Among its reptile foods, Lazy H raises discoid roaches, one of the few kinds that Florida allows people to cultivate on purpose. "We love the roaches," Chiasson says. They withstand temperature fluctuations and shipping better than the crickets do. "They don't smell as bad," she adds. And when a roach dies in the tank, cleanup is easy: The other roaches eat most of it.

Can't argue with that. The discoids live in a trailer, with worn blond wood paneling. A row of big fertilizer bins lines each side. Snorkel says that at night the whole trailer rustles.

Snorkel, too, clearly loves the roaches, calling them "relatable." She points out some babies, a mature female and something that would make a fine half-moon pendant. It's what's left of a dead roach after the funeral lunch. Then she spots a female on its back, waving her legs. "Contractions," Snorkel explains. Females release a brood of young in a large sac.

So that is how I find myself standing between a Silicon Valley software entrepreneur and a reinventor of the pantsuit in an old trailer waiting quietly for a cockroach to give birth. This is about as far as you can get from Silicon Valley, but you can see the future from here too.

Explore more

- Åsa Berggren, Anna Jansson and Matthew Low. "Approaching ecological sustainability in the emerging insects-as-food industry." *Trends in Ecology & Evolution*. February 1, 2019.

COURTESY OF BROOKLYN BUGS

Fancy feast

Lobster with roasted waxworms (which are caterpillars, not worms) highlights chef Joseph Yoon's principle of pairing edible insects with people's favorite foods.

I met the Ovipost innovators in Athens, Ga., at the August 2018 meeting of the North American Coalition for Insect Agriculture. The event gave a glimpse of one possible future in which insects become a prestige protein like shrimp, crab and lobster. Considering how multilegged and armored so much shellfish is, it's intriguing that so many people relish shrimp cocktail, yet cringe at eating a land-based arthropod cousin.

The banquet chef, Joseph Yoon from the insect-eating advocacy group Brooklyn Bugs, goes big and bold. He doesn't disguise the ants, dried scorpions or silkworm larvae; he flaunts them. The insects become luxury highlights for dishes people already love. His eight-course extravaganza began with insect bodies mixed with popcorn and ended with a silky chocolate mousse topped with wasps. People talked about the varied flavors of insects, but for me, it was more about texture — an added crunch or chewiness.

Partisans of insect protein admit that many people won't ever embrace this flamboyant cuisine. But insects could still play a major role as more subtle ingredients. Insect powders, for example, are already muscling into protein bars. Nobody counts legs once they're ground into flour. Farmed insects could also be used to feed farmed fish, an animal we're willing to eat, replacing as much as 25 to 30 percent of food for fish raised in aquaculture, European food specialists estimated in a 2015 review. — *Susan Milius*

Dario Wünsch was feeling confident. The 28-year-old from Leipzig, Germany, was about to become the first professional gamer to take on the artificial intelligence program AlphaStar in the rapid-fire video game StarCraft II. Wünsch had been professionally playing StarCraft II, in which competitors command alien fleets vying for territory, for nearly a decade. No way could he lose this five-match challenge to a newly minted AI gamer.

Even AlphaStar's creators at the London-based AI research company DeepMind, which is part of Alphabet, Inc., weren't optimistic about the outcome. They were the latest in a long line of researchers who had tried to build an AI that could handle StarCraft II's dizzying complexity. So far, no one had created a system that could beat seasoned human players.

Sure enough, when AlphaStar faced off against Wünsch on December 12, the AI appeared to commit a fatal mistake at the onset of the first match: It neglected to build a protective barrier at the entrance to its camp, allowing Wünsch to infiltrate and quickly pick off several of its worker units. For a minute, it looked like StarCraft II would remain one realm where humans trump machines. But AlphaStar made a winning comeback, assembling a tenacious strike team that quickly laid waste to Wünsch's defenses. AlphaStar 1, Wünsch 0.

Wünsch shook it off. He just needed to focus more on defense. But in the second round, AlphaStar surprised the pro gamer by withholding attacks until it had amassed an army that once again crushed Wünsch's forces. Three matches later, AlphaStar had won the competition 5-0, relegating Wünsch to the small but growing club of world-class gamers bested by a machine.

Using algorithms originally developed to help five AIs play the game Dota 2, OpenAI researchers built an extremely dexterous robot hand.

Researchers have long used games as benchmarks for AI smarts. In 1997, IBM's Deep Blue earned international acclaim by outwitting chess champion Garry Kasparov (*SN: 8/2/97, p. 76*). In 2016, DeepMind's AlphaGo famously overpowered Go champion Lee Sedol (*SN: 12/24/16, p. 28*).

But board-based contests like chess and Go can only push AI so far. These games are still pretty simple — players take turns and can see every piece's position on the board at all times. When it comes to making an AI that can deal with real-world ambiguity and fast-paced interactions, the most useful tests of machine cognition will probably be found in games played in virtual worlds.

Building AI gamers that can trounce human players is more than a vanity project. "The ultimate idea is to ... use those algorithms [for] real-world challenges," says Sebastian Risi, an AI researcher at IT University of Copenhagen. For

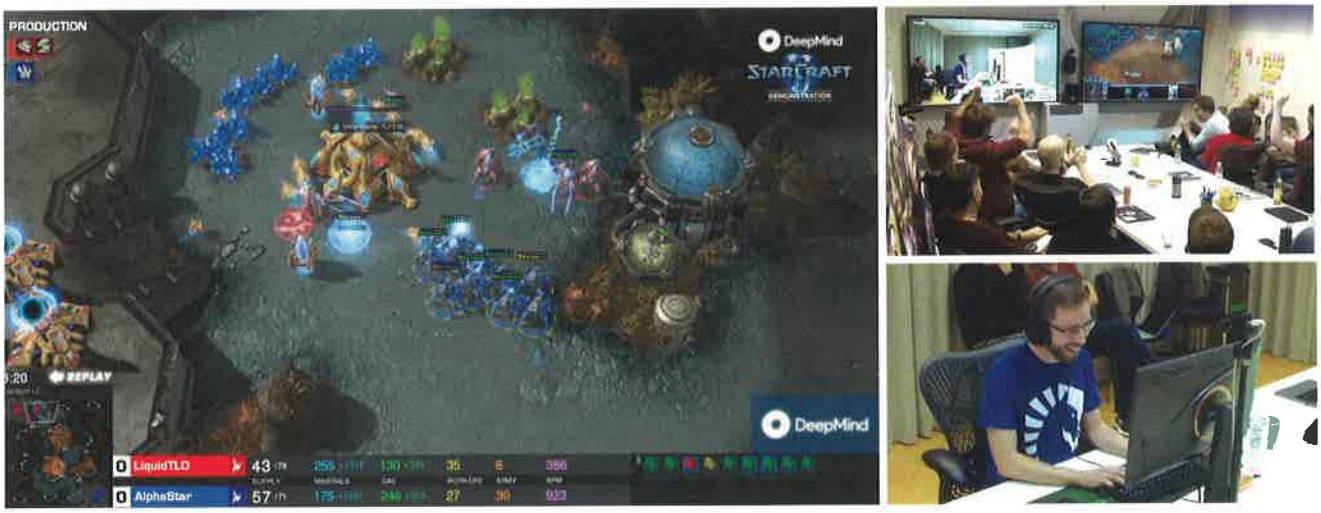

In this scene from the video game Starcraft II (left image), professional StarCraft II player Dario Wünsch, who plays as "LiquidTLO," is overpowered by the artificial intelligence AlphaStar, which is wreaking havoc on Wünsch's base. The AI's creators at DeepMind cheer as the AI shows its superiority (top right). Wünsch took it in stride when AlphaStar bested him, 5 matches to 0 (bottom right).

TOP: OPENAI; BOTTOM, ALL: DEEPMIND

instance, after the San Francisco–based company OpenAI trained a five-AI squad to play an online battle game called Dota 2, the programmers repurposed those algorithms to teach the five fingers of a robotic hand to manipulate objects with unprecedented dexterity. The researchers described this work online at arXiv.org in January.

DeepMind researchers similarly hope that AlphaStar's design could inform researchers trying to build AIs to handle long sequences of interactions, like those involved in simulating climate change or understanding conversation, an especially difficult task (*SN: 3/2/19, p. 8*).

Right now, two important things that AIs still struggle with are: coordinating with each other and continually applying new knowledge to new situations. The StarCraft universe has proved to be an excellent testing ground for techniques that make AI more cooperative. To experiment with methods to make AIs forever learners, researchers are using another popular video game, Minecraft. While people may use screen time as an entertaining distraction from real life, virtual challenges may help AI pick up the skills necessary to succeed in the real world.

Arcade education AI can practice different skills in video games to learn how to get along in the real world. Navigational know-how, for example, could help search-and-rescue robots prowl tough terrain, and AIs that know how to manage many workers could help run companies.

Game types that teach AI useful skills for the real world

Type	Racing	First-person shooting	Open world	Real-time strategy
Example games	Forza Motorsport, Real Racing	Doom	Minecraft, Grand Theft Auto	StarCraft
Navigation	✓	✓	✓	
Manage resources/staff	✓			✓
Plot strategy	✓	✓		✓
Fast reaction	✓	✓		✓
Collaboration		✓		✓
Setting goals			✓	
Creativity			✓	
Exploration			✓	✓
Lifelong learning			✓	
Motivation			✓	✓
Juggling priorities			✓	✓

Team play

When AlphaStar took on Wünsch, the AI played StarCraft II like a human would: It acted like a single puppeteer with complete control over all the characters in its fleet. But there are many real-world situations in which relying on one mastermind AI to micromanage lots of devices would become unwieldy, says artificial intelligence researcher Jakob Foerster of Facebook AI Research in San Francisco.

Think of overseeing dozens of nursing robots caring for patients throughout a hospital, or self-driving trucks coordinating their speeds across miles of highway to mitigate traffic bottlenecks. So, researchers including Foerster are using the StarCraft games to try out different "multiagent" schemes.

In some designs, individual combat units have some independence, but are still beholden to a centralized controller. In this setup, the overseer AI acts like a coach shouting plays from the sidelines. The coach generates a big-picture plan and issues instructions to team members. Individual units use that guidance, along with detailed observations of the immediate surroundings, to decide how to act. Computer scientist Yizhou Wang of Peking University in China and colleagues reported the effectiveness of this design in a paper submitted to *IEEE Transactions on Neural Networks and Learning Systems*.

Wang's group trained its AI team in StarCraft using reinforcement learning, a type of machine learning in which computer systems pick up skills by interacting with the environment and getting virtual rewards after doing something right. Each teammate received rewards based on the number of enemies eliminated in its immediate vicinity and whether the entire team won against fleets controlled by an automated opponent built into the game. On several different challenges with teams of at least 10 combat units, the coach-guided AI teams won 60 to 82 percent of the time. Centrally controlled AI teams with no capacity for independent reasoning were less successful against the built-in opponent.

AI crews with a single commander in chief that exerts at least some control over individual units may work best when the group can rely on fast, accurate communication among all agents. For instance, this system could work for robots within the same warehouse.

But for many machines, such as self-driving cars or drone swarms spread across vast distances, separate devices "won't have consistent, reliable

and fast data connection to a single controller," Foerster says. It's every AI for itself. AIs working under those constraints generally can't coordinate as well as centralized teams, but Foerster and colleagues devised a training scheme to prepare independent-minded machines to work together.

In this system, a centralized observer offers feedback to teammates during reinforcement learning. But once the group is fully trained, the AIs are on their own. The master agent is less like a sidelined coach and more like a dance instructor who offers ballerinas pointers during rehearsals, but stays mum during the onstage performance.

The AI overseer prepares individual AIs to be self-sufficient by offering personalized advice during training. After each trial run, the overseer simulates alternative possible futures and tells each agent, "This is what actually happened, and this is what would have happened if everyone else had done the same thing, but you did something different." This method, which Foerster's team presented in New Orleans in February 2018 at the AAAI Conference on Artificial Intelligence, helps each AI unit judge which actions help or hinder the group's success.

To test this framework, Foerster and colleagues trained three groups of five AI units in StarCraft. Trained units had to act based only on observations of the immediate surroundings. In combat rounds against identical teams commanded by a built-in, nonhuman opponent, all three AI groups won most of their rounds, performing about as well as three centrally controlled AI teams in the same combat scenarios.

Lifelong learning

The types of AI training that programmers test in StarCraft and StarCraft II are aimed at helping a team of AIs master a single task, for example, coordinating traffic lights or drones. The StarCraft games are great for that, because for all their moving parts, the games are fairly straightforward: Each player has the singular goal of overpowering an opponent. But if artificial intelligence is going to become more versatile and humanlike, programs need to be able to learn more and continually pick up new skills.

"All the systems that we see right now that play Go and chess — they're basically trained to do this one task well, and then they're fixed so they can't change," Risi says. A Go-playing system presented with an 18-by-18 grid, instead of the standard 19-by-19 game board, would probably have to be completely retrained on the new board, Risi says. Changing the characteristics of StarCraft units would require the same back-to-square-one training. The Lego-like realm of Minecraft turns out to be a better place for testing approaches to make AI more adaptable.

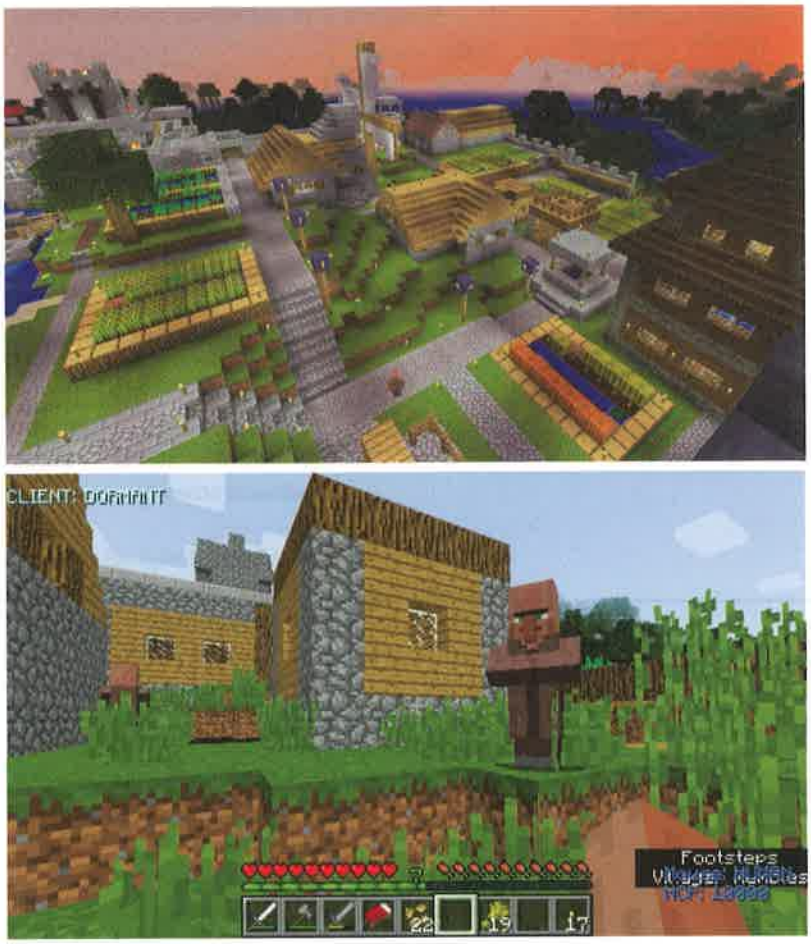

In different environments composed of 3-D blocks, Minecraft players can build structures and explore their surroundings (two examples above). Researchers use Minecraft to teach AI skills such as how to set goals and build creative structures.

Unlike StarCraft, Minecraft poses no single quest for players to complete. In this virtual world made of 3-D blocks of dirt, glass and other materials, players gather resources to build structures, travel, hunt for food and do pretty much whatever else they please. Caiming Xiong, an artificial intelligence researcher at the San Francisco–based software company Salesforce, and colleagues used a simple building full of blocks in Minecraft to test an AI designed to continually learn.

Rather than assigning the AI to learn a single task through trial and error in reinforcement learning, Xiong's team staggered the AI's education. The researchers guided the AI through increasingly difficult reinforcement learning challenges, from finding specific blocks to

FROM TOP: CHRISTOPH SALGE *ET AL/PROC. OF THE 13TH INTERNAT. CONF. ON THE FOUNDATIONS OF DIGITAL GAMES 2018*; LAIGE PENG AND YOSHIMASA TSURUOKA/*INFORMATION PROCESSING SOCIETY OF JAPAN 2018*

Quick on the uptake A Minecraft-playing AI that knows how to apply past knowledge to learn new skills (dark green) more quickly learns how to perform a new skill successfully. It reached rewards up to 1.0 in fewer attempts than an AI that doesn't rely on old expertise (light green).

SOURCE: T. SHU, C. XIONG AND R. SOCHER/6TH INTERNAT. CONF. ON LEARNING REPRESENTATIONS 2018

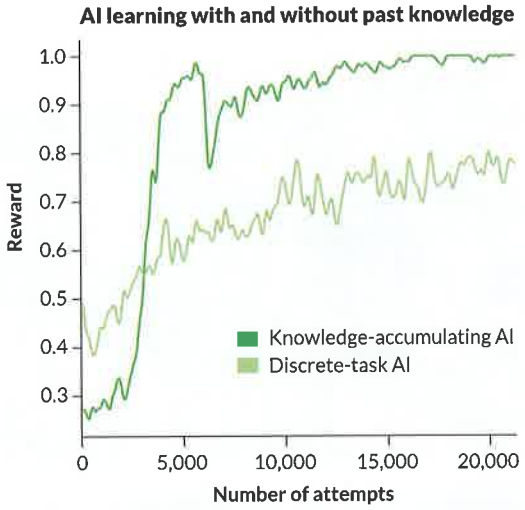

stacking blocks. The AI was designed to break each challenge into simpler steps. It could tackle each step using old expertise or try something new. Compared with another AI that was not designed to use prior knowledge to inform new learning experiences, Xiong team's AI proved a much quicker study.

The knowledge-accumulating AI was also better at adjusting to new situations. Xiong and colleagues taught both AIs how to pick up blocks. While training in a simple room that contained only one block, both AIs got the "collect item" skill down pat. But in a room with multiple blocks, the discrete-task AI struggled to identify its target and grabbed the right block only 29 percent of the time.

The knowledge-accumulating AI knew to rely on a previously learned "find item" skill to locate a target object among distractions. It picked up the right block 94 percent of the time. The research was presented in Vancouver in May 2018 at the International Conference on Learning Representations.

With further training, Xiong and colleagues' system could master more skills. But this design is limited by the fact that the AI can only learn tasks assigned by the human programmer during training. Humans don't have this kind of educational cutoff. When people finish school, "it's not like, 'Now you're done learning. You can freeze your brain and go,'" Risi says.

A better AI would get a foundational education in games and simulations and then be able to continue learning throughout its lifetime, says Priyam Parashar, a roboticist at the University of California, San Diego. A household robot, for example, should be able to find navigational work-arounds if residents install baby gates or rearrange the furniture. Parashar and colleagues created an AI that can identify instances in which it needs further training without human input. When the AI runs into a new obstacle, it takes stock of how the environment is different from what it expected. Then it can mentally rehearse various work-arounds, imagine the outcome of each and choose the best solution.

The researchers tested this system with an AI in a two-room Minecraft building. The AI had been trained to retrieve a gold block from the second room. But another Minecraft player had built a glass barrier in the doorway between the rooms, blocking the AI from collecting the gold block. The AI assessed the situation and, through reinforcement learning, figured out how to shatter the glass to complete its task, Parashar and her colleagues reported in the 2018 *Knowledge Engineering Review*.

An AI faced with an unexpected baby gate or glass wall should probably not conclude that the best solution is to bust it down, Parashar admits. But programmers can add additional constraints to an AI's mental simulations — like the knowledge that valuable or owned objects should not be broken — to inform the system's learning, she says.

New video games are becoming AI test-beds all the time. AI and games researcher Julian Togelius of New York University and colleagues hope to test collaborating AIs in Overcooked — a team cooking game that takes place in a tight, crowded kitchen where players are constantly getting in each other's way. "Games are designed to challenge the human mind," Togelius says. Any video game by nature is a ready-made test for how much AI know-how can emulate human cleverness.

But when it comes to testing AI in video games or other simulated worlds, "you cannot ever say, 'OK, I've modeled everything that's going to happen in the real world,'" Parashar says. Bridging the gap between virtual and physical reality will take more research.

One way to keep simulation-trained AI from overreaching, she suggests, is to devise systems that require AIs to ask humans for help when needed (*SN: 3/2/19, p. 8*). "Which, in a sense, is making [AI] more like humans, right?" Parashar says. "We get by with the help of our friends." ■

29
percent
Success rate for an AI that could not use prior knowledge to grab the correct block in Minecraft

94
percent
Success rate for an AI that built on previous knowledge to grab the correct block in Minecraft

Explore more

■ Niels Justesen *et al.* "Deep learning for video game playing." *IEEE Transitions on Games.* February 13, 2019.

C. CHANG