

The big new idea for making self-driving cars that can go anywhere

The mainstream approach to driverless cars is slow and difficult. These startups think going all-in on AI will get there faster.

[Will Douglas Heaven](#)

Four years ago, Alex Kendall sat in a car on a small road in the British countryside and took his hands off the wheel. The car, equipped with a few cheap cameras and a massive neural network, veered to the side. When it did, Kendall grabbed the wheel for a few seconds to correct it. The car veered again; Kendall corrected it. It took less than 20 minutes for the car to learn to stay on the road by itself, he says.

This was the first time that reinforcement learning—an AI technique that trains a neural network to perform a task via trial and error—had been used to [teach a car to drive from scratch](#) on a real road. It was a small step in a new direction—one that a new generation of startups believes just might be the breakthrough that makes driverless cars an everyday

reality.

Reinforcement learning has had enormous success producing computer programs that can [play video games](#) and Go with superhuman skill; it has even been used to [control a nuclear fusion reactor](#). But driving was thought to be too complicated. “We were laughed at,” says Kendall, founder and CEO of the UK-based driverless-car firm Wayve.

Wayve now trains its cars in rush-hour London. Last year, it showed that it could take a car trained on London streets and have it [drive in five different cities](#)—Cambridge (UK), Coventry, Leeds, Liverpool, and Manchester—without additional training. That’s something that industry leaders like Cruise and Waymo have struggled to do. This month Wayve announced it is teaming up with Microsoft to train its neural network on Azure, the tech giant’s cloud-based supercomputer.

Investors have sunk [more than \\$100 billion](#) into building cars that can drive by themselves. That’s a third of what NASA spent getting humans to the moon. Yet despite a decade and a half of development and untold miles of road testing, driverless tech is stuck in the pilot phase. “We are seeing extraordinary amounts of spending to get very limited results,” says Kendall.

That’s why Wayve and other autonomous-vehicle startups

like Waabi and Ghost, both in the US, and Autobrains, based in Israel, are going all in on AI. Branding themselves AV2.0, they're betting that smarter, cheaper tech will let them overtake current market leaders.

Hype machines

Wayve says it wants to be the first company to deploy driverless cars in 100 cities. But is that yet more hype from an industry that's been getting high on its own supply for years?

"There is way too much overselling in this field," says Raquel Urtasun, who led Uber's self-driving team for four years before leaving to found Waabi in 2021. "There's also a lack of acknowledgment of how difficult the task is in the first place. But I don't believe that the mainstream approach to self-driving is going to get us to where we need to be to deploy the technology safely."

That mainstream approach dates back at least to 2007 and the [DARPA Urban Challenge](#), when six teams of researchers managed to get their robotic vehicles to navigate a small-town mock-up on a disused US Air Force base.

Waymo and Cruise launched on the back of that success, and the robotics approach taken by the winning teams stuck. That approach treats perception, decision-making, and

vehicle control as different problems, with different modules for each. But this can make the overall system hard to build and maintain, with errors in one module bubbling over into others, says Urtasun. "We need an AI mindset, not a robotics mindset," she says.

Here's the new idea. Instead of building a system with multiple neural networks and wiring these together by hand, Wayve, Waabi, and others are each building one large neural network that figures out the details by itself. Throw enough data at the AI and it learns to convert input (camera or lidar data about the road ahead) into output (turning the wheel or hitting the brakes), much like a kid learning to ride a bike.

Going straight from input to output like this is known as end-to-end learning, and it's what GPT-3 did for [natural-language processing](#) and AlphaZero did for Go and chess. "In the last 10 years it's caused so many seemingly unsolvable problems to get solved," says Kendall. "End-to-end learning pushed us forward to superhuman capabilities. Driving will be no different."

Like Wayve, Waabi is using end-to-end learning. It isn't (yet) using real vehicles, however. It is developing its AI almost fully inside a [super-realistic driving simulation](#), itself controlled by an AI driving instructor. Ghost also adopts an AI-first approach, building driverless tech that not only navigates roads but learns to react to other drivers.

200,000 small problems

Autobrain is betting on an end-to-end approach too, but does something different with it. Instead of training one large neural network to handle everything a car might encounter, it is training many smaller networks—hundreds of thousands, in fact—to handle a very specific scenario each.

“We’re translating the hard AV problem into hundreds of thousands of smaller AI problems,” says Igal Raichelgauz, the company’s CEO. Using one large model makes the problem more complex than it actually is, he says: “When I’m driving, I’m not trying to understand every pixel on the road. It’s about extracting contextual cues.”

Autobrain takes the sensor data from a car and runs it through an AI that matches the scene to one of many possible scenarios: rain, pedestrian crossing, traffic light, bicycle turning right, car behind, and so on. By watching a million miles of driving data, Autobrain says, its AI has identified around 200,000 unique scenarios, and the company is training individual neural networks to handle each of them.

The firm has been partnering with car manufacturers to test its technology and has just got hold of a small fleet of its own vehicles.

Kendall thinks that what Autobrains is doing might work well for advanced driver-assist systems, but he does not see it having an advantage over his own approach. "When tackling the full self-driving problem, I'd expect that they would be just as challenged by the complexity faced in the real world," he says.

Cruise control

Either way, should we count on this new wave of firms to chase down the front-runners? Unsurprisingly, Mo ElShenawy, executive vice president of engineering at Cruise, isn't convinced. "The state of the art as it exists today is not sufficient to get us to the stage where Cruise is at," he says.

Cruise is one of the most advanced driverless-car firms in the world. Since November it has been running a live robotaxi service in San Francisco. Its vehicles operate in a limited area, but anyone can now hail a car with the Cruise app and have it pull up to the curb with nobody inside. "We see a real spectrum of reactions from our customers," says ElShenawy. "It's super exciting."

Cruise has built a vast virtual factory to support its software, with hundreds of engineers working on different parts of the pipeline. ElShenawy argues that the mainstream modular approach is an advantage because it lets the company swap

in new tech as it comes along.

He also dismisses the idea that Cruise's approach won't generalize to other cities. "We could have launched in a suburb somewhere years ago, and that would have painted us into a corner," he says. "The reason we've picked a complex urban environment, such as San Francisco—where we see hundreds of thousands of cyclists and pedestrians and emergency vehicles and cars that cut you up—was very deliberate. It forces us to build something that scales easily."

But before Cruise drives in a new city, it first has to map its streets in centimeter-level detail. Most driverless car companies use these kinds of high-definition 3D maps. They provide extra information to the vehicle on top of the raw sensor data it gets on the go, typically including hints like the location of lane boundaries and traffic lights, or whether there are curbs on a particular stretch of street.

These so-called HD maps are created by combining road data collected by cameras and lidar with satellite imagery. Hundreds of millions of miles of roads have been mapped in this way in the US, Europe, and Asia. But road layouts change every day, which means map-making is an endless process.

Many driverless-car companies use HD maps created and maintained by specialist firms, but Cruise makes its own.

"We can re-create cities—all the driving conditions, street layouts, and everything," says ElShenawy.

This gives Cruise an edge against mainstream competitors, but newcomers like Wayve and Autobrains have ditched HD maps entirely. Wayve's cars have GPS, but they otherwise learn to read the road using sensor data alone. It may be harder, but it means they are not tied to a particular location.

For Kendall, this is the key to making driverless cars widespread. "We are going to be slower to get into our first city," he says. "But once we get to one city, we can just scale everywhere."

For all the talk, there's a long way to go. While Cruise's robotaxis are driving paying customers around San Francisco, Wayve—the most advanced of the new crop—has yet to test its cars without a safety driver. Waabi doesn't even use real cars.

Still, these new AV2.0 firms have recent history on their side: end-to-end learning rewrote the rules of what's possible in computer vision and natural-language processing. So their confidence is not misplaced. "If everybody goes in the same direction and it's the wrong direction, we're not going to solve this problem," says Urtasun. "We need a diversity of approaches, because we haven't seen the solution yet."