News & views

Artificial intelligence

Neural networks overtake humans in a racing game

J. Christian Gerdes

Driving a racing car requires a tremendous amount of skill. Now, artificial intelligence has challenged the idea that this skill is exclusive to humans — and it might even change the way automated vehicles are designed. See p.223

A modern Formula 1 race is a breathtaking display of engineering precision. Yet the popularity of the sport arguably has less to do with the performance of the cars than with the skill and daring displayed by the drivers as they push those cars to the limit. Success on the race track has been a celebrated human achievement for more than a century. Will it now become a similar triumph for artificial intelligence (AI)? On page 223, Wurman *et al.*[1] take a step in this direction by introducing Gran Turismo (GT) Sophy, a neural-network driver capable of outperforming the best human players of the video game *Gran Turismo*.

The objective in racing is easily defined: if you complete the circuit in less time than your competitors, you win. However, achieving this goal involves a complicated battle with physics, because negotiating the track requires careful use of the frictional force between the tyre and the road, and this force is limited. Using some of that friction for braking, for instance, leaves less force available for rounding a corner.

More specifically, each tyre can produce a frictional force proportional to the vertical force, or load, that connects it to the road. As the car accelerates, the load shifts to the rear tyres, leaving less frictional force for the front tyres. This can induce understeer, in which the steering wheel cannot generate more cornering force and effectively becomes a hand rest as the car ploughs out of the turn. By contrast, when the car brakes, the load shifts to the front of the car. This can lead to oversteer, meaning that the rear tyres lose traction and the car spins. Add in a complicated track topography, and the complexities of tuning load transfer with the suspension of the vehicle, and the challenges of racing become obvious.

To win the race, the driver must choose trajectories that allow the car to stay within these ever-changing friction limits as much as it physically can. Brake too early going into a turn and your car is slow, losing time. Brake too late and you won't have enough cornering force to hold your desired racing line as you near the tightest part of the turn. Brake too hard and you might induce a spin. Professional racing drivers are eerily good at finding and maintaining the limits of their car, lap after lap, for an entire race.

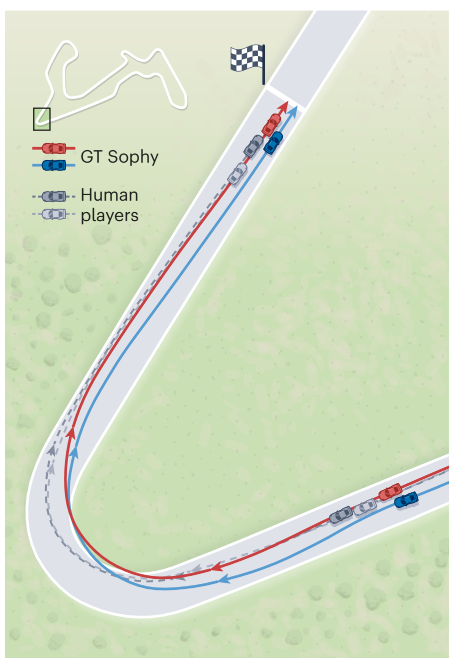

Figure 1 | Neural-network drivers outperform human players. Wurman *et al.*[1] report a neural-network algorithm — called GT Sophy — that is capable of winning against the best human players of the video game *Gran Turismo*. When two human drivers attempted to block the preferred path of two GT Sophy cars, the algorithm found two ways to overtake them. (Adapted from Fig. 3d of ref. 1.)

As complex as the handling limits of a car can be, they are well described by physics, and it therefore stands to reason that they could be calculated or learnt. Indeed, the automated Audi TTS, Shelley, was capable of generating lap times comparable to those of a champion amateur driver by using a simple model of physics[2]. By contrast, GT Sophy doesn't make explicit calculations based on physics. Instead, it learns through a neural-network model. However, given the track and vehicle motion information available to Shelley and GT Sophy, it isn't too surprising that GT Sophy can put in a fast lap with enough training data.

What really stands out is GT Sophy's performance against human drivers in a head-to-head competition. Far from using a lap-time advantage to outlast opponents, GT Sophy simply outraces them. Through the training process, GT Sophy learnt to take different lines through the corners in response to different conditions. In one case, two human drivers attempted to block the preferred path of two GT Sophy cars, yet the AI succeeded in finding two different trajectories that overcame this block and allowed the AI's cars to pass (Fig. 1).

GT Sophy also proved to be capable of executing a classic manoeuvre on a simulation of a famous straight of the Circuit de la Sarthe, the track of the car race 24 Hours of Le Mans. The move involves quickly driving out of the wake of the vehicle ahead to increase the drag on the lead car in a bid to overtake it. GT Sophy learnt this trick through training, on the basis of many examples of this exact scenario — although the same could be said for every human racing-car driver capable of this feat. Outracing human drivers so skilfully in a head-to-head competition represents a landmark achievement for AI.

The implications of Wurman and colleagues' work go well beyond video-game supremacy. As companies work to perfect fully automated vehicles that can deliver goods or passengers, there is an ongoing debate as to how much of the software should use neural networks and how much should be based on physics alone. In general, the neural network is the undisputed champion when it comes to perceiving and identifying objects in the surrounding environment[3]. However, trajectory planning has remained the province of physics and optimization. Even vehicle manufacturer Tesla, which uses neural networks as the core of autonomous driving, has revealed that its neural networks feed into an optimization-based trajectory planner (see go.nature.com/3kgkpua).

But GT Sophy's success on the track suggests that neural networks might one day have a larger role in the software of automated vehicles than they do today.

So, will the Formula 1 battles between Lewis Hamilton and Max Verstappen give way to contests between GT Sophy variants? After all, the physics of *Gran Turismo* is a close match for real racing cars. *Gran Turismo*'s director, Kazunori Yamauchi, even used the video game to find ways of tweaking his real racing car to overcome a recurring problem that he was having when taking a corner at the Nürburgring, a Grand Prix track in Germany that has the nickname The Green Hell (see go.nature.com/3tw22aa). It also helped me to familiarize myself with Laguna Seca Raceway before I started racing school.

Still, some challenges remain in moving from the console to the track. For example, GT Sophy has not yet learnt that it is sometimes better to follow the car ahead to make up time, instead of dogfighting at every corner. Of course, Wurman *et al.* report GT Sophy's rookie season, and there is no obvious reason why such a strategy could not be learnt with greater experience, too.

More challenging might be the variation that occurs with each lap. Unlike in the *Gran Turismo* races used by Wurman and co-workers, the condition of the tyres on real racing cars changes from lap to lap, and human drivers must adapt to such changes throughout the race. Would GT Sophy be able to do the same with more data? And where would such data come from? It's easy to run simulations, but no racing car in existence has completed enough laps to train GT Sophy in its current form, much less an AI that could handle tyre variability. However, there is evidence that neural networks can capture changing vehicle dynamics on different road surfaces[2], so perhaps Verstappen and Hamilton should keep one eye on their rear-view mirrors.

J. Christian Gerdes is in the Department of Mechanical Engineering, Stanford University, Stanford, California 94305, USA.
e-mail: gerdes@stanford.edu

1. Wurman, P. R. *et al. Nature* **602**, 223–228 (2022).
2. Spielberg, N. A., Brown, M., Kapania, N. R., Kegelman, J. C. & Gerdes, J. C. *Sci. Robot.* **4**, eaaw1975 (2019).
3. Fujiyoshi, H., Hirakawa, T. & Yamashita, T. *IATSS Res.* **43**, 244–252 (2019).

The author declares competing interests. See go.nature.com/3it9ddz for details.

Plant sciences

Hard graft problem solved for key global food crops

Colin Turnbull & Sean Carrington

Grafting has long been used to join tissues of different plants in horticulture and research. Methods have now been devised to extend the technique to plants called monocotyledons, which include major crops such as cereals and bananas. **See p.280**

The technique of grafting together the shoot of one plant and the roots of another is immensely beneficial in a variety of contexts. However, efforts to use this approach have long failed for certain key crops. On page 280, Reeves *et al.*[1] report success in developing a grafting method that can be used for plants called monocotyledons, or monocots.

Plant grafting has an ancient history, dating back to early civilizations. More than 2,000 years ago, *De Agri Cultura* ('On Farming'), a book written by the Roman senator Cato, details the grafting of vines and fruit trees, indicating it to be commonplace. Such grafting practices remain widespread today.

Yet one major group of plants, the monocots, have proved problematic for use in grafting. The name refers to the single leaf (a cotyledon) in the plant seed, a feature that distinguishes monocots from other flowering plant groups that have two cotyledons, and that are conventionally called dicotyledons, or dicots. Monocots abound in the global flora. They include all of the world's cereals – rice, wheat and maize (corn) – which together provide more than half the calories consumed by humans. Another key monocot is the banana, a staple food in many nations and the world's most popular fruit after the tomato.

Despite many attempts to graft monocots, minimal success meant that grafting never became mainstream. Indeed, many experts viewed monocot grafting as a near-impossible feat[2,3], often attributing failure to anatomical differences between monocots and dicots, especially monocots' lack of a specialized inner cellular layer found in dicots called vascular cambium. However, paradoxically, there are witchweeds – dicot plants of the genus *Striga* that live as devastating parasites attached to the roots of monocot crops such as maize and sorghum[4]. The parasite feeds through an interface that is, in essence, a natural graft plumbed into the host's transport (vascular) systems, proving that nature long ago accomplished a version of grafting that humans have struggled to achieve.

Reeves and colleagues present compelling evidence that monocot grafting is feasible, after all, and propose that the absence of vascular cambium is not a limiting factor. Instead, their work focuses on a feature shared by all plants: immature tissue that can be reprogrammed to make the essential connecting structures needed for a successful graft. The authors' use of fine surgical tools enabled the precise assembly of grafts of young germinating seedlings at a time in the plants' development when the root is just emerging, whereas the same method tried in older plants proved much less successful.

One key to plants' terrestrial dominance is linked to their extraordinary ability to recover from damage, whether arising from storms, herbivore grazing or simply humans mowing the grass. With a much greater regenerative capacity than that of most vertebrates, many types of plant cell are totipotent, enabling the replacement of missing parts. Sometimes, this regenerative process requires the formation of a disorganized group of cells called a callus, from which tissues and organs emerge. In grafting, when two cut pieces are placed together, a wound-repair mechanism makes connections between the pieces, resulting in a new whole plant. An essential feature of this process is connection of the plant's 'plumbing system' – the vascular highways of tissues called xylem and phloem that transport water, sugars, nutrients and other molecules throughout a plant (Fig. 1).

Reeves *et al.* report that, within days of making a monocot graft, they observe fluorescent dyes (applied to the cut surface) moving in both directions across the graft. Vascular cells develop, and the graft is sufficiently strong to be picked up by hand.

The authors find that genes in cells around the graft junction are rapidly expressed, as a prelude to visible signs of graft formation. The expression of many of these genes is a hallmark of regenerative processes. The genes encode wound-repair factors, regulatory proteins and hormones, as well as components needed to