

## From the archive

Remembering a key advance in our understanding of the cell cycle, and a tale of perilous Antarctic exploration.

### 50 years ago

Stephen R. Pelc, who died suddenly ... at the age of 65, was internationally famous for his pioneer work on the techniques of autoradiography and for his use of these techniques in cellular biological research. ...[A]t the Hammersmith Hospital, London, ... he began his studies on the action of ionizing radiations on photographic film which led to his development of stripping-film autoradiography ... He is most famous for work begun in the early nineteen-fifties with Dr Alma Howard. By incorporating  $^{32}\text{P}$  into dividing cells and removing all but the DNA by acid hydrolysis he was able to time the incorporation of  $^{32}\text{P}$  into nuclear DNA. He showed that DNA synthesis did not occur after prophase, as had been believed previously from staining evidence, nor did it occur continuously throughout interphase. He and Howard showed that, for each type of nucleus, there was a particular period of interphase, which he called the "S" (synthesis) stage, during which the DNA content doubled; this DNA was stable and became divided equally into the two daughter nuclei. Before and after the "S" there was a gap in his knowledge of what metabolic processes occurred in the nuclei and, understandably, he named these "G<sub>1</sub>" and "G<sub>2</sub>".

From *Nature* 23 March 1973

### 100 years ago

*The Worst Journey in the World: Antarctic, 1910–1913.* By Apsley Cherry-Garrard — This is the sixth book to give the story, or part of the story, of Capt. Scott's last expedition, and it is in some ways the most remarkable of them all. Mr. Cherry-Garrard took part in three of the worst journeys ever made in the Antarctic or anywhere else, and the iron of his sufferings has entered into his soul and imparted a ferric quality to his recollections ... If poetry be indeed definable as "emotion recollected in tranquillity," Mr. Cherry-Garrard has given us a true epic of exploration.

From *Nature* 24 March 1923



that researchers use to transfer DNA into cells. The authors created a set of plasmids that use the modified codon language of the engineered cells to encode components required for plasmid replication in bacteria. These plasmids can function only in cells with a synthetic genome and engineered tRNAs, a scenario that strikingly reduces the risk of engineered DNA being unintentionally transferred to wild bacterial populations (Fig. 1c).

It currently takes a huge effort to establish a working synthetic genome, with only a handful completed so far. Our capabilities on this front are slowly scaling up, with a full synthetic genome for a eukaryotic cell (one that contains a nucleus) expected to be finalized in the next few years<sup>6</sup>, and work towards a human synthetic genome project also under way<sup>7</sup>. As the number, size and ambition of synthetic-genome projects increases, so, too, will our ability to study and manipulate biology. The impressive feats achieved through codon repurposing in this work will be immensely valuable to bacterial biotechnology, in which viral contamination is a persistent and expensive problem.

The biggest impact of this work will probably be in providing a foundation for similar

strategies in synthetic genomes for other organisms. Increasingly, key medical products, such as vaccines and protein therapeutics, depend on the use of mammalian or human cell-culture systems that are vulnerable to viral infection, with substantial implications for cost and product safety<sup>8</sup>. Controlled, reliable manufacturing processes that are protected from problems of viral infection will be crucial for maximizing these industries' positive impact on health and well-being, while ensuring that the processes are safe, contained and retain public confidence.

**Benjamin A. Blount** is at the Biodiscovery Institute, School of Life Sciences, University of Nottingham, Nottingham NG7 2RD, UK. e-mail: benjamin.blount@nottingham.ac.uk

1. Nyerges, A. et al. *Nature* **615**, 720–727 (2023).
2. Fredens, J. et al. *Nature* **569**, 514–518 (2019).
3. Robertson, W. E. et al. *Science* **372**, 1057–1062 (2021).
4. Zürcher, J. F. et al. *Science* **378**, 516–523 (2022).
5. Mandell, D. J. et al. *Nature* **518**, 55–60 (2015).
6. Richardson, S. M. et al. *Science* **355**, 1040–1044 (2017).
7. Boeke, J. D. et al. *Science* **353**, 126–127 (2016).
8. Barone, P. W. et al. *Nature Biotechnol.* **38**, 563–572 (2020).

The author declares competing interests. See [go.nature.com/3mqzn6s](https://go.nature.com/3mqzn6s) for details.

This article was published online on 15 March 2023.

## Engineering

# Hazards help autonomous cars to drive safely

Colin Paterson & Chiara Picardi

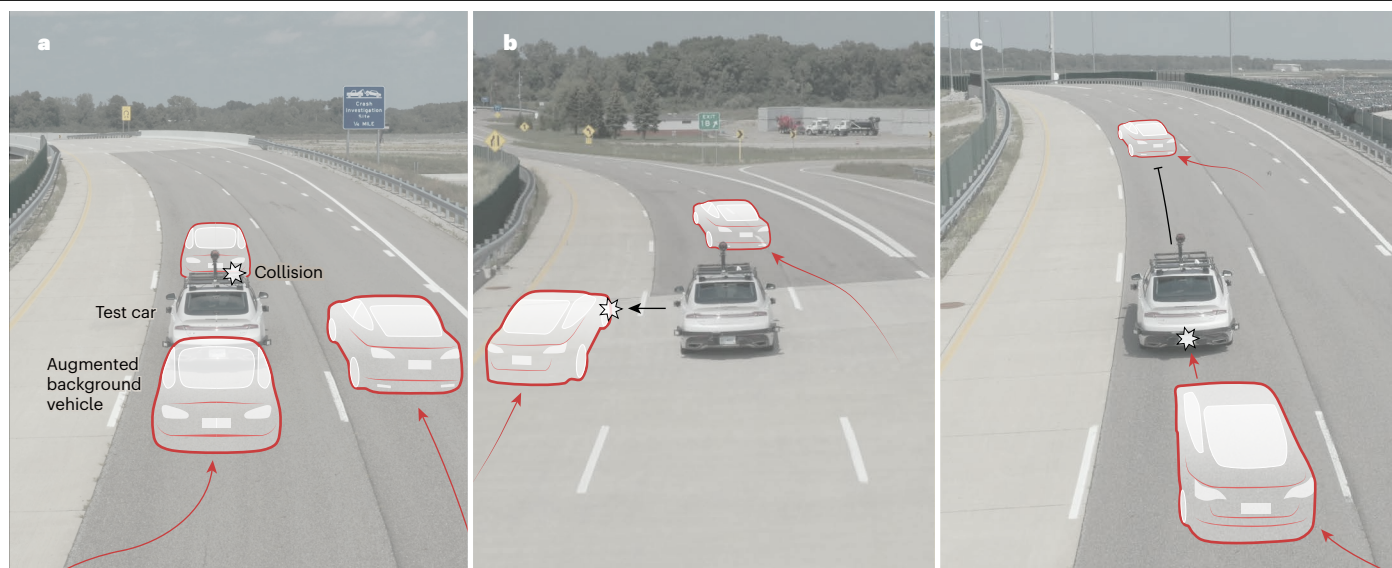
Collecting training data by focusing on dangerous scenarios offers an efficient way for artificial intelligence to improve the safety of autonomous vehicles. Augmented reality allows the approach to be tested without risking lives. **See p.620**

Every year, collisions involving road vehicles kill or seriously injure tens of thousands of people in the United Kingdom alone (see [go.nature.com/3ekkek4](https://go.nature.com/3ekkek4)). Autonomous vehicles could reduce these numbers, but their safety is yet to be guaranteed (see [go.nature.com/3ykw43v](https://go.nature.com/3ykw43v)). Identifying potentially hazardous situations and testing how an autonomous agent will react are crucial parts of the safety-assurance process. But people are not necessarily adept at recognizing situations that would be hazardous for non-human drivers because they do not register small changes in visual information that might confuse a machine. On page 620, Feng *et al.*<sup>1</sup> introduce a method that uses artificial intelligence (AI) to validate the AI of autonomous vehicles.

The challenge of thinking like a machine is not the only reason it is difficult to test

situations that pose a hazard for autonomous vehicles. The first barrier is the sheer volume of data to assess. Human drivers in the United States are estimated to crash once every 850,000 kilometres. Autonomous vehicles currently fare even worse than this: a human operator has to take control of a self-driving vehicle around once every 80,000 kilometres to avoid a crash. However, this still amounts to a lot of safe driving between collisions, which means that searching for test cases is like looking for a needle in a haystack. Simply gathering more data is, therefore, unlikely to improve road safety.

The hazard itself also complicates the task of testing for safety. Trialling autonomous vehicles by placing real people in danger is out of the question, so instead researchers must examine data from a limited set of real-world



**Figure 1 | An augmented reality test track for autonomous vehicles.**

Feng *et al.*<sup>1</sup> developed an artificial-intelligence agent for testing the safety of autonomous vehicles. The agent controls ‘background’ vehicles in a simulation to find potentially unsafe situations, thereby building a dense set of dangerous events on which to train an autonomous test vehicle. The authors then use augmented reality to implement this training approach on a real

track. **a**, A background vehicle changes lanes behind the test car while another vehicle prevents the test car from changing lanes, causing it to collide with a third vehicle in front. **b**, The test car changes lanes to avoid a background vehicle cutting in front of it, but collides with a second vehicle. **c**, The test car decelerates to avoid a background vehicle changing lanes ahead, and a second vehicle collides with it from behind. (Adapted from Fig. 2 of ref. 1.)

examples, and imagine and simulate unsafe behaviours that might never be encountered.

Such simulations will always be based on incomplete models because complexity is computationally costly. And in simplifying the model of the world to reduce complexity, researchers run the risk of removing the key factors that cause collisions. Indeed, some safety concerns might arise only when these factors are combined, so understanding which factors to remove is not easy.

Feng *et al.* approached this problem by first assuming that an AI testing agent (under the guidance of safety engineers) is better placed than a person to find situations in which an autonomous system will fail. Like its human counterpart, the AI tester is interested in finding cases that cause failures and therefore tries to create an environment in which such cases arise. The authors used reinforcement learning to train an AI testing agent to come up with a plan for controlling the behaviours of ‘background’ vehicles that can manoeuvre between lanes in a simulation. They then selected all the cases in which the choices made by these background vehicles led to a dangerous situation (Fig. 1).

Feng and colleagues’ AI agent learnt to ignore the vast set of situations in which vehicle behaviours are benign, and instead built a dense set of test data for safety-critical events. In this way, unimportant scenarios were removed from the data set, decreasing computational complexity and increasing the efficiency and effectiveness of testing.

In reality, of course, a simulation does not capture the true complexity of the real world. Small changes in environmental conditions,

such as dirt on a camera lens or a sudden change in lighting conditions, can cause safety failures. To tackle this problem, Feng *et al.* demonstrated a way of testing a real car, driving autonomously on a test track, using a form of augmented reality. As it circles the track, the car must contend with virtual drivers that have been specifically trained by the AI to display adversarial behaviours. These encounters test the response of the autonomous system to dangerous situations, in an environment that closely resembles the real world – without putting any lives at risk.

Feng and colleagues’ work is a key step

**“A human operator has to take control of a self-driving vehicle around once every 80,000 kilometres to avoid a crash.”**

towards assuring the safety of autonomous vehicles, but several challenges remain. The authors’ approach relies on motorized vehicles being the only adversarial agents in the test space. This makes defining dangerous encounters that arise from interactions with other motorists easy to define and therefore straightforward to test. However, an autonomous vehicle driving in the real world also needs to be aware of cyclists and pedestrians. And there are external factors – including system failures and environmental conditions – that could lead to safety violations beyond the AI’s control, especially when they occur

simultaneously. These factors currently preclude autonomous vehicles from being considered safe under all possible conditions.

The authors’ approach provides valuable insights into possible hazardous situations faced by an autonomous vehicle, and the safety violations that might result from these events. But it is not yet clear how this information could be used to improve controllers, vehicles and overall safety. We must accept that even a rigorously tested autonomous system will occasionally have to make difficult decisions, for which even human drivers do not have good answers (see *Nature* **562**, 469–470; 2018). Defining acceptable behaviours under such conditions remains challenging. Making sure that all stakeholders understand the perception and decision-making abilities of machines will be key to persuading the public that autonomous systems are safe<sup>2</sup>.

Despite these challenges, Feng and colleagues’ method improves the viability of testing autonomous systems in complex environments. It does not replace the need for human oversight in constructing reasoned arguments for safety, but – as a tool in the belt of the safety engineer – it could well support the validity of those arguments.

**Colin Paterson** and **Chiara Picardi** are in the Department of Computer Science, University of York, York YO10 5GH, UK.  
e-mails: colin.paterson@york.ac.uk;  
chiara.picardi@york.ac.uk

1. Feng, S. *et al.* *Nature* **615**, 620–627 (2023).

2. Burton, S. *et al.* *Artif. Intell.* **279**, 103201 (2020).

The authors declare no competing interests.