thus boosting the availability of food for krill[6,7].

Iron is almost insoluble in seawater, and the bulk of this crucial element in productive ecosystems exists in living biomass. Krill are exceptionally versatile animals, and can channel their astoundingly diverse food sources — phytoplankton such as diatoms, sea-ice algae, fluff settled on sediments, copepods and other zooplankton — into their biomass. This enables krill populations to act as a gigantic, mobile iron reservoir. Observers during the pre-whaling era described the sea surface as being coloured red by swarming krill, and reported that water spouts of feeding whales stretched from horizon to horizon[7].

Making the reasonable assumption that the former krill stock was three times the size of the whales' annual krill consumption, I estimate that such stock, spread out evenly over the whaling grounds (an area of approximately two million square kilometres), corresponds to 300 krill per square metre, which would be enough to colour the surface of the water red. That biomass of krill would hold enough iron, if released through biological recycling, to fuel a massive bloom of diatoms in the water column below. In reality, roving krill swarms would probably have been concentrated in offshore regions that favoured the accumulation of diatom blooms, which the whales would have fertilized with the iron from their faecal plumes.

Left undisturbed, diatom blooms form snow-like aggregates that sink to the deep sea three to four weeks after an initial iron fertilization, taking the iron with them[8]. A roving krill swarm once grazed down a diatom bloom that my colleagues and I were studying, leaving behind clouds of loose, slowly sinking faecal threads full of undigested food and living cells[9]. If a feeding whale had pursued this swarm, the turbulence associated with the animal's energetic swimming, lunges and filtration would have dispersed the threads and mixed their contents into the water column, rather like the way in which manure is ploughed into a field. The energy invested in such actions would have a larger return for the whale in the form of blubber amassed from subsequent feeding. This must have been an optimized, sustainable recycling ecosystem, which operated at high levels of biomass in the past — the more the merrier.

Krill started declining after the decimation of the whales, with the last large-scale surface swarms having been recorded in the early 1980s[10]. Removal of a predator is often accompanied by a rise in prey numbers, and this surprising decline in krill is consistent with a model in which whale-aided iron cycling supported the growth of krill populations. Krill biomass is now a fraction of what it once was, and the hugely productive ocean pastures dominated by diatoms, described in the 1930s[11,12], have since reverted to the classic iron-limited, high-nutrient, low-chlorophyll, microbially dominated state that is now characteristic of large areas of the ocean's surface. This degraded ecosystem dominates the former whaling grounds, presumably because the hard-working whales are almost completely absent.

A 2020 survey that found 55 blue whales in former whale-feeding grounds made the news as a sign of hope (see go.nature.com/3bffqla). However, the fact that it was newsworthy is alarming. How can a handful of whales, subsisting on the meagre food offered by the vagaries of nature in a degraded ecosystem, ever restore one of the previously hottest hotspots of animal biomass on the globe[13] if not helped by humans? We have in our power the means to mimic the iron fertilization mediated by whales to create diatom blooms, to feed the krill and thereby to feed the whales. This might restore the former pastures of plenty whose evolution the whales worked so hard to shape. The open-ocean experiments necessary to test this hypothesis[6] are waiting to be carried out.

Victor Smetacek is at the Alfred Wegener Institute Helmholtz Centre for Polar and Marine Research, 27570 Bremerhaven, Germany.
e-mail: victor.smetacek@awi.de

1. Laws, R. M. *Phil. Trans. R. Soc. B* **279**, 81–96 (1977).
2. Pauly, D. & Zeller, D. *Nature Commun.* **7**, 10244 (2016).
3. Savoca, M. S. *et al. Nature* **599**, 85–90 (2021).
4. Goldbogen, J. A. *et al. Science* **366**, 1367–1372 (2019)
5. Williams, T. M. *Science* **366**, 1316–1317 (2019).
6. Smetacek, V. in *Impacts of Global Warming on Polar Ecosystems* (ed. Duarte, C. M.) 45–81 (Fund. BBVA, 2008).
7. Nicol, S. *et al. Fish Fish.* **11**, 203–209 (2010).
8. Smetacek, V. *Protist* **169**, 791–802 (2018).
9. González, H. E. *Polar Biol.* **12**, 81–91 (1992).
10. Holm-Hansen, O. & Huntley, M. *J. Crustac. Biol.* **4** (5), 156–173 (1984).
11. Hart, T. J. *Discov. Rep.* **21**, 261–356 (1942).
12. Hardy, A. *Great Waters: A Voyage of Natural History to Study Whales, Plankton and the Waters of the Southern Ocean* (Harper & Row, 1967).
13. Bar-On, Y. M., Phillips, R. & Milo, R. *Proc. Natl Acad. Sci. USA* **115**, 6506–6511 (2018).

The author declares no competing interests.

Coronavirus

An algorithm to target COVID testing of travellers

Ziad Obermeyer

Optimizing the testing of incoming travellers for COVID-19 involves predicting those who are most likely to test positive. A machine-learning algorithm for targeted testing has been implemented at the Greek border. **See p.108**

It seems an obvious combination: machine learning and the fight against COVID-19. And yet, despite intense interest and increasing availability of large data sets, success stories of such combinations are few and far between. On page 108, Bastani *et al.*[1] describe a system that they designed and deployed at points of entry into Greece, starting in August 2020. The algorithm, which is built on a method called reinforcement learning, markedly increased the efficiency of testing for the coronavirus SARS-CoV-2, and contributed to Greece's ability to keep its borders open safely. The work also provides a clear warning about the shortcomings of the comparatively blunt policy tools that most other countries continue to use.

Testing is a problem that machine learning is well suited to solve. Imagine a border-control agent on a Greek island. A flight has just landed, and the agent's task is to identify and detain anyone who has COVID-19. The agent might want to test all arriving passengers, but the testing capacity on the island is very limited and, more generally, it is never possible to test 100% of any population 100% of the time. The alternative — shutting down the border completely, in an economy highly dependent on tourism — has its own perils. These would include not only a huge financial cost associated with the loss of jobs and income, but also the negative effects of such losses on public health[2]. So the border agent faces a difficult decision: who should be tested?

As has been noted[3], the value of a test depends on its eventual outcome. In this scenario, a negative test generates only costs: the cost of testing and a delay for the traveller. By contrast, a positive test generates tremendous benefit: prevention of all the cases of COVID-19 that a traveller infected with SARS-CoV-2 would have caused. So, in deciding who to test, the border agent's optimal strategy is clear: predict which travellers have the highest likelihood of testing positive, and test them. This strategy maximizes the value of testing, because it detects the most travellers with COVID-19 using the lowest number of tests.

MILOS BICANSKI/GETTY

Figure 1 | COVID-19 testing of travellers arriving at Eleftherios Venizelos International Airport in Athens.

If the border agent could predict which incoming passengers are most likely to test positive, tests could be allocated efficiently (Fig. 1). Conveniently, data about incoming passengers — their country and region of origin, age and sex — are available digitally, on the passenger locator form that all travellers complete 24 hours before arrival in Greece. It seems straightforward enough to use data from past tests of incoming passengers to predict which 'types' of passenger might be more likely to test positive in the future. But, as decades of research in statistics and computer science have shown[4], this strategy runs the risk of getting locked into yesterday's pandemic: given the rapidly evolving dynamics of COVID-19 spread, an algorithm must quickly adapt its predictions to stay one step ahead and still test the right passengers.

This is where the value of machine learning becomes clear. Just as an algorithm can be trained to play the game Go[5] by learning which moves lead to winning the game, Bastani and colleagues trained an algorithm to allocate scarce tests, by learning which passengers are likely to test positive.

Crucially, the algorithm balances two goals. The first, and most natural, goal is to test passenger types who are likely to test positive, by exploiting patterns learnt from previous data about the outcome of tests for SARS-CoV-2 in these different groups. The second — perhaps less intuitive, but equally important — is to explore patterns not reflected in previous data, by testing passenger types about which the algorithm knows little.

Then, at a given port of entry on a given day, the algorithm delivers targeted recommendations to border agents about which passengers to test, while respecting the budget and resource constraints imposed by supply chains, staffing, laboratory capacity and delivery logistics for biological samples. These constraints are real and binding: the authors note that, at the peak of the summer tourism season, there was capacity to test only 18.4% of incoming travellers — even after the Greek National COVID-19 Committee of Experts wisely approved group testing to drive efficiency gains in the lab.

The authors draw on the reinforcement-learning strategies that have powered advances in online commerce and marketing[6]. But using such an algorithm in the real world raises its own technical challenges. For example, the algorithm must learn discontinuously, from large batches of testing results, rather than one-by-one from individual results. And the feedback from batch results is delayed, forcing the algorithm to operate uninformed while waiting for results. Solving these challenges required substantial tweaking of the algorithms that are typically designed for easier, more data-rich online settings.

The thorniest challenges, however, are legal and political ones. To comply with the European Union's General Data Protection Regulation (GDPR), the authors deliberately limited the data available to the algorithm — and thus its accuracy — in close consultation with lawyers, epidemiologists and policymakers. The potential limit placed on the algorithm's performance by the GDPR highlights how well-intentioned laws to protect privacy can have both positive and negative consequences. In a pandemic that does not respect individuals' privacy, such regulations can ultimately hamper the ability of a government to protect the health of its citizens. The authors also adapted the algorithm with a policymaker audience in mind, choosing their optimization methods to showcase clearly the value of both algorithm goals: testing high-risk passengers and testing high-uncertainty passengers.

The results are impressive. The automated system doubled the efficiency of testing — the number of cases detected per test — allowing border agents to test and quarantine the right passengers, many of whom were asymptomatic, while letting others through to their final destination.

The success of the algorithm presented by Bastani and colleagues highlights the inadequacy of the border policies of nearly all other countries. The decisions underlying these policies — for example, whether to deny all travellers entry to the country or to force the testing or quarantine of all travellers from a given country — have two key flaws. First, these decisions are made about entire countries, rather than individuals, disregarding vast differences between people within countries. Second, they are typically made on the basis of country-level epidemiological data that, as the present study shows, have notable shortcomings.

Had border agents denied entry to all passengers from countries that had concerning metrics, they would have prevented those people with COVID-19 from entering Greece — but at the cost of crushing a key pillar of the economy. Had they simply tested people proportional to a country's reported COVID-19 metrics rather than algorithmic predictions, however, their testing efficiency would have been much lower. This is because reported COVID-19 metrics can be very different from actual disease prevalence among incoming travellers. Travellers are not randomly drawn from their countries' populations, and passively collected data on cases of COVID-19 or deaths associated with the disease reflect large reporting biases and systemic barriers to access[7].

Indeed, by efficiently testing incoming passengers, the authors' algorithm was able to anticipate spikes in SARS-CoV-2 infection rates among traveller populations almost 9 days earlier than if they had used country-level epidemiological data alone. This indicates the enormous value of intelligent, deliberate data collection — and the dangers of relying on blunt, flawed, country-level data for important decisions.

Bastani and colleagues' work will be remembered as one of the best examples of using data in the fight against COVID-19. It is a compelling story of how a group of researchers partnered with enlightened policymakers to produce a tool that has enormous social value. It highlights the best parts of both academic research and the civil service, and shows the great promise of artificial intelligence for making good decisions — which in many settings can be the difference between life and death.

Ziad Obermeyer is in the Division of Health Policy and Management, School of Public Health, University of California, Berkeley, Berkeley, California 94720, USA.
e-mail: zobermeyer@berkeley.edu

1. Bastani, H. *et al. Nature* **599**, 108–113 (2021).
2. Marmot, M. & Wilkinson, R. (eds) *Social Determinants of Health* (Oxford Univ. Press, 2005).
3. Mullainathan, S. & Obermeyer, Z. *Diagnosing Physician Error: A Machine Learning Approach to Low-Value Health Care*. National Bureau of Economic Research Working Paper 26168 (2021).
4. Thompson, W. R. *Biometrika* **25**, 285–294 (1933).
5. Silver, D. *et al. Nature* **529**, 484–489 (2016).
6. Li, L., Chu, W., Langford, J. & Schapire, R. E. in *Proc. 19th Int. Conf. World Wide Web* 661–670 (2010).
7. Wu, S. L. *et al. Nature Commun.* **11**, 4507 (2020).

The author declares no competing interests.
This article was published online on 22 September 2021.

Cryptography

Relativity could ensure security for cash machines

Gilles Brassard

Entering your personal identification number using the keypad of a cash machine is notoriously insecure. A clever application of the special theory of relativity could make identification safer. **See p.47**

When you type in your personal identification number (PIN) at a cash machine, you feel safe — provided you cover the keypad with your hand. But even the machines attached to banks are vulnerable to attack by fraudsters, some of whom go as far as to add fake machinery to legitimate machines as a way of stealing PINs (see go.nature.com/3p9r431). To prevent this type of fraud, a solution is needed that allows people to prove their identity without disclosing any secret information. On page 47, Alikhani *et al.*[1] describe an experiment that achieves this goal with unprecedented security, guaranteed by Albert Einstein's special theory of relativity.

The identification technique used by Alikhani and colleagues is an application of a concept known as a zero-knowledge proof[2]. Imagine that Alice wants to convince her friend Bob that she knows how to do something, but she wants to keep her technique secret. For example, suppose she is capable of distinguishing between two brands of cola in a glass simply by looking at them. She asks a sceptical Bob to switch around identical glasses containing the two types of cola while her back is turned. If she can still tell the drinks apart, and can repeat this feat several times, Bob will be convinced of her ability — but he won't have learnt how to do it himself. Alice has performed a zero-knowledge proof.

If Alice's trick relies on secret knowledge that is unique to her, then being able to do this task correctly functions like a cash-machine PIN for Alice's identity. Alice proves that she has information known only to her, but she doesn't share this information. So even if the cash machine is equipped with fake machinery, the person who installed it cannot later use what was learnt to impersonate Alice.

Remarkably, there is a zero-knowledge proof for any mathematical statement for which a conventional proof exists[3] — and there are highly efficient schemes for applying zero-knowledge proofs to the task of establishing someone's identity[4]. However, implementing a general zero-knowledge proof involves encoding the answers with an entirely different mathematical problem, such as factoring a large number. This means that the alleged security of most conventional zero-knowledge proofs[3,4] depends on how difficult this other mathematical problem is to solve. Unfortunately, once quantum computers are readily available, it will become possible to solve many of these other problems in a period of time that is sufficiently short to defeat the validity of the zero-knowledge proof[5].

To make a zero-knowledge proof unconditionally secure, Alice would need to prove her identity using two separate devices that cannot communicate with each other for the duration of her interaction with the bank[6]. This is similar to a detective interrogating two suspects in different rooms to determine the consistency of their joint alibi. Ideally, the two devices, provided by the bank, would require Alice to supply biometric information to activate them. They would be inserted into a pair of slots on the cash machine and perform zero-knowledge proofs on her behalf (Fig. 1). Because each of Alice's devices is kept ignorant of the questions asked of the other device, their answers will sometimes be inconsistent if they do not possess the secret they claim to have — that is, if they are fraudulent. But how can we be certain that the two devices cannot communicate with each other? This is where Einstein's special theory of relativity comes to the rescue.

Special relativity tells us that information cannot travel faster than light. Suppose Alice's two devices are one metre apart. Any signal requires more than 3.3 nanoseconds to travel between them. Therefore, if each device is required to respond within 3 ns of receiving their question, and if the questions are asked within a time window of 0.3 ns, the devices will be prevented from choosing their answers on the basis of the questions put to the other device. Such exquisite precision was thought to be technologically infeasible because it seemed to imply that an enormous amount of data would need to be communicated in this short time frame. However, a much