FACULTY

TEN LESSONS FROM TEN YEARS OF APPLIED AI

CONTENTS

2014	FACULTY FOUNDED
2015	
2016	
2017	
2018	HOME OFFICE — NCA
2019	
2020	NHS
2021	CERA
2022	DRIFT — NESO
2023	INSPIRED — NOVARTIS — BEAZLEY
2024	OPENAI

Introduction

At Faculty, we feel extremely fortunate to have spent the last ten years at the coalface of Applied AI.

One of the great pleasures has been the clients we have got to work with over the years. More often than not they are brilliant people. Many have become friends. We admire the vision each has demonstrated in seeing how this relatively new field of technology can help them achieve their objectives. And we are humbled by the importance of the problems that they and their organisations deal with.

The stories contained in this volume cover people who focus on:

• Preventing terrorist attacks.

• Navigating the UK National Health Service through the pandemic.

• Freeing teachers from admin to spend much more time with children.

• Bringing new sources of renewable energy onstream.

• Rewiring the way businesses run in industries with hundreds of years of pedigree.

• Keeping elderly people out of hospital.

• Retooling GB's energy system operator for the age of renewable energy.

• Speeding up the development of life-saving drugs.

• Protecting vulnerable people from serious and organised crime.

• Ensuring that advanced AI models are safe.

Each of these matters a lot. To the organisations involved, of course. But also to each of us as citizens.

In all of these cases, AI offers a powerful new way to take on these challenges, as the stories will demonstrate. In this respect, our view is that AI offers great cause for optimism. If deployed properly, it can make a meaningful difference to our ability to solve the important problems of our time.

But only if it is deployed properly. And this is not straightforward. The real world, the organisations within it, and the things they are trying to achieve are always messy and complicated in their own ways. Using new technology to achieve real change in the world requires deftness and skill, patience and perseverance. There are many more ways to get this wrong than get it right.

We have learned a lot about how to get this right over the last ten years. AI has been our sole focus over that period, and we have been involved in hundreds of applications across virtually every sector of the economy. As much as we have enjoyed the successes, we have also had our fair share of failures in that time. We have seen every possible complexity that the world can throw at you, and lived through every way that a technology project can go wrong. It is often from these experiences that the best learnings arise.

The lessons contained in this book are our best attempt to synthesise and codify the sum of ten years' experience. Each is viewed through the lens of a particular client and their story, that shows in some way why it is important.

The ten lessons are summarised here:

01. AI is an operational discipline, not an analytical one.

AI is not being used to its potential where the output is a series of charts and dashboards that describe the world without acting upon it. The field of BI already provides the world with more of that than it will ever need. AI should be built into the tools that people use to run business processes, not just those that analyse them. It should be integrated with the levers used to intervene in those processes, so that analysis leads directly to actions.

02. AI is technology for human decision makers.

All software should be built around the user. The user for intelligent software is typically a decision maker. Where a decision is important, human decision makers should remain in control, supported rather than replaced by AI. In order to do this effectively, AI tools need to provide analysis that is targeted, parsimonious, explainable and interactive.

03. Augment human tasks that require judgement. Automate those that don't.

Most people enjoy the core of their job. But they dislike the bureaucracy they have to do. AI should be used to provide decision support to these tasks at the core, which are where people create the most value and exercise their professional judgement. And it should be used to automate away the admin that crowds out the fun stuff. This blend of automation and augmentation tends to make people both more productive and more fulfilled.

04. AI is a feature not a product. But it can define a product.

Despite recent advances in language models, there are few occasions where an AI model alone makes a full product. It is a cog in a machine, rather than the machine. However, it is an unusually powerful cog that can make new kinds of machines possible. In much the same way that an engine isn't a car, but it was the thing that made the car possible. This is a good time to seek new ways of solving old problems.

05. Building AI widgets is easy. Rewiring your organisation is not.

AI party tricks, like meeting summaries, can seem exciting, but tend to be limited to the periphery of what matters. The bigger impacts come when AI is used to optimise the core business processes that define an organisation. This requires much more than technology, and should be thought of as a serious change programme.

06. It's data SCIENCE, not DATA science.

Data gets a lot of attention. But data alone solves no problem. It's the science that you do on top of the data that matters the most. Science is all about building an understanding of the world, and the cause and effect relationships that drive it. This is the foundation of applying AI successfully. You need to understand the cause and effect relationships inside a system before you intervene on it. And you need to identify the causal pathway by which your interventions achieve the outcomes you seek.

07. There's no such thing as complete data.

Data is essential to modern AI. But it is a precision game - more is not always better. A specific model, with a specific objective, will need specific data to achieve that. You need to understand what problem you're solving and work back from that to the exact data you need. Focus on getting that data in place first, rather than trying to build the perfect data infrastructure. Perfect data doesn't exist, and never will.

08. Build in increments that are individually valuable & collectively transformative.

Over ambition and under ambition each bring their own perils for AI programmes. Don't just focus on low-hanging fruit, or on individual use cases. That's not how you make a measurable difference. But also don't try to imagine the whole future at once and seek a big bang. The chances are that will fail. Instead think big, but build towards it in modular steps that are individually valuable, but when connected together transform a whole business process end to end.

09. Business strategy trumps AI strategy.

For a small share of businesses, AI will render their current strategy obsolete. But for the vast majority, a standalone AI strategy is a bad idea. Any well run business will already know what is important to them. Being strategic about AI means determining whether and if so how AI can accelerate these known priorities. It does not mean creating a new rival strategy or a menu of all the ways AI could be used in a business, stack ranked against each other.

10. If you don't control your models, your models control you.

AI models are probabilistic. The most powerful are black boxes. They don't always behave in entirely predictable ways. And users can't really tell why they do the things they do. As these technologies are embedded more deeply in decision-making processes, it is essential that the correct controls are put in place around them. Humans should be in the loop where possible, and 'over the loop' where not. Platforms like Faculty's Frontier are needed to implement these controls across an organisation.

The remainder of this book brings these lessons to life through the stories of some of our favourite clients. We think that each of the people and organisations profiled here is a leader in some part of the application of AI to the world. We are very grateful to have worked together with them, for their willingness to share their stories, and their contributions to this book. Finally, we hope that by capturing these stories and the learnings they contain we can inspire others to use AI as a means to solve the problems they care about, and make good choices about how best to approach them to achieve the outcomes they seek.

Angie Ma
Co-Founder

John Gibson
Chief Commercial Officer

Successfully putting AI to work for the good of society takes leadership. We are grateful to a set of leaders who have each shown the vision to apply AI to a challenge that matters, and who have been kind enough to share their perspectives for this book.

SAM ALTMAN	CEO, OPENAI
SHREERAM ARADHYE	CHIEF MEDICAL OFFICER, NOVARTIS
SIR TONY BLAIR	CHAIR TBI, FORMER UK PRIME MINISTER
MATT COLLINS	DNSA, CABINET OFFICE
TROY DEHMANN	COO, BEAZLEY
TORBEN LUNDBERG	CIO, INSPIRED EDUCATION GROUP
BEN MARUTHAPPU	FOUNDER & CEO, CERA
BEN MEDLAND	FOUNDER & CEO, DRIFT
NADIM NSOULI	CEO, INSPIRED EDUCATION GROUP
SHAHEEN SAYED	HEAD OF ACCENTURE, UK, IRELAND & AFRICA
CLAIRE SMITH	COO / DIRECTOR GENERAL CAPABILITIES, NCA
SIR SIMON STEVENS	FORMER CHIEF EXECUTIVE OF NHS ENGLAND
MING TANG	CHIEF DATA & ANALYTICS OFFICER, NHS ENGLAND

01

HOME OFFICE

LESSON ONE

AI is an operational discipline, not an analytical one.

The UK Home Office wanted to take down online propaganda that was inciting terrorist attacks around the world, but the tech giants said it couldn't be done. In proving them wrong, Faculty learned how to develop AI to meet the most demanding operational requirements.

Tucked away behind a nondescript door, at the end of a long corridor in the UN building in New York, is a suite of rooms rarely glimpsed by the outside world. This is the UK Room, a perk afforded only to permanent members of the Security Council, where British diplomats can withdraw for discreet conversations and private negotiations, or top themselves up from the always-on coffee machine. Here, on a muggy September day, UK Prime Minister Theresa May huddled with her aides and went over her speech one last time. Dressed in a crisp navy blue jacket and white blouse, and wearing a chunky silver chain necklace, she was about to go onstage at the General Assembly. She was going to throw down the gauntlet.

It was 2018, and the world was grappling with a wave of terror attacks inspired by the terror group Daesh (also known as ISIL, Islamic State or ISIS). In the previous year, 18 attacks had been launched against civilian targets in Western countries; over a hundred people were murdered and many more injured.

In each case, an important inflammatory role was played by the slick online propaganda that Daesh was flooding onto the internet. From glossy jihadist videos to practical bomb-building tutorials, Daesh was able to radicalise its recruits, school them in violence, and ultimately move them to commit terrible acts.

After four attacks in the UK, Theresa May had had enough. In front of the eyes of the world, she demanded that tech companies make a paradigm shift in their ability to stop terrorist propaganda. If they could not identify it and take it down within two hours of it being posted - the crucial window of opportunity - then her government would legislate to force their hand.

The tech companies said it wasn't possible.The sheer volume of content Daesh pumped out would overwhelm any human moderators, while automated solutions were out of the question. The social media giants had the best AI and software engineering teams in the world, the largest-scale digital operations ever built, and they could confidently say that the technology May was demanding didn't and couldn't exist yet.

But the British Prime Minister knew otherwise.

THE TERROR OF DAESH

The people within the Office for Security and Counter-Terrorism (OSCT) are one of the smartest and most impressive groups in the British government. Few among the general public know the team exists, but they're lucky it does. Located in a highly secure area of the Home Office headquarters in Westminster, the staff who work there are deeply expert in analysing and understanding terrorist threats. A typical example is Tom Drew OBE, who worked there for seven years (and later joined Faculty). A softly-spoken thirty something, with a dark brown beard and a penetrating gaze, Tom has dedicated his career to keeping the public safe.

In 2017, he and his team were troubled by a new emerging threat. In the Middle East, Daesh's self-styled 'Caliphate' was in retreat: it had suffered significant territorial losses and been driven out of its de facto capitals in Mosul and Raqqah. In response, the group changed its tactics. Instead of encouraging supporters in the West to come to Syria and Iraq, they urged them to stay at home and carry out 'single-actor' attacks against unprotected civilian targets. Acting alone, outside of existing networks and often with no previous history of extremist activity, they would be almost impossible to stop.

The result was dramatic. An attack on the Houses of Parliament ended with five people killed; a bomber at an Ariana Grande performance in Manchester killed 22 concertgoers; a mass stabbing in London killed eight more.

The resulting coroner inquests have shown that almost all the attackers were radicalised by the propaganda they encountered online. Daesh was using the West's own social media networks against it, recruiting its killers in plain sight on places like Facebook, Twitter and YouTube. The platforms had teams of human moderators trying to find and remove the content, but they were far too few and way too slow to have any meaningful impact.

And when the government asked the social media companies to do more, they were stonewalled. The web was too big. There were too many videos. The technology to automate content moderation didn't exist yet. 'They were quite passive,' recalls Matt Collins, who is now the Deputy National Security Advisor for Intelligence, Defence and Security. At the time, he was Director of Prevent, with overall responsibility for the Home Office team looking at the problem. 'They were still working to manual checks, and we had to prove to them that machine learning could help.'

Even if you tried it, the tech companies argued, no system would be 100% accurate, and even a tiny failure rate would create enormous problems. 'We review over one hundred million pieces of content every month,' said Mark Zuckerberg in February 2017, 'and even if our reviewers get 99% of the calls right, that's still millions of errors over time.' Bottom line: counter terrorism experts might be able to identify individual pieces of content with a high degree of certainty, but to apply that kind of rigorous analysis at scale in a real-time operational setting was flat out impossible.

Tom Drew wasn't buying it. As Head of Data and Innovation, he had a responsibility to do everything in his power to stop the attacks by cutting off the torrent of Daesh propaganda. He suspected that advances in machine learning might have changed the equation. If he could prove that an AI model could replicate the Home Office team's expertise and apply it on the necessary scale, the government would have the ammunition it needed to force the tech giants to adopt new standards.

The social media platforms had made it very clear that anyone who thought there was a technological solution didn't understand how technology works.

So Tom went to find people who did. He started asking around, talking to a range of experts to see if there was a way to do what he needed. One of the places he came to was Faculty.

IMPORTANT DISTINCTIONS BETWEEN ANALYTICS AND OPERATIONS

At the time, Faculty was a tiny startup operating out of an Edwardian townhouse in Marylebone. 'Even the building sort of made a statement,' recalls Angie Ma, one of the co-founders. 'It didn't look like a classic tech office. We wanted to make the point that we were a different sort of company, that tech didn't have to be in your face.' Still using its original name, Advanced Skills Initiative, the company had barely 20 employees. But Tom felt they had potential.

'Their approach was (and is) to be collegiate problem-solvers,' he recalls. 'They treated it as an experiment to develop a new model, not an opportunity to just resell an existing process or product.' Matt Collins concurs. 'There was a can-do attitude when we presented the problem, an enthusiasm to just get stuck in and really see what the art of the possible was.'

Of course, the nature of an experiment is you don't know ahead of time if it'll succeed. John Gibson, who at the time led ASI's consulting business, remembers fielding the first call from Tom. 'We thought it could be done,' says John, looking back, 'but the only way we could prove that it would work was to actually build the technology.'

The heart of the challenge was one that comes up surprisingly often in the AI world: the distinction between analytics and operations. In many organisations, Data Science and AI teams have typically been siloed off in Analytics or Business Intelligence functions, well away from the messy business of the shop floor. Analytical teams provide insight for understanding the world; operational teams act on it and make things happen.

The danger, of course, is that a gap develops so that insights never actually get put into practice. And in this case, that gap was a chasm. The analysts were at the OSCT, while the people who could implement their recommendations were at the tech companies. Not only were they not on the same page, they didn't even believe what the other was saying was possible.

Enter AI. The power of the technology is that it can take insights and weave them into actual workflows - but it's not a quick fix. When you're building technology for operational processes, there are always requirements you need to account for. They might be user needs, existing workflows, infrastructure requirements, policies or regulations - all the existing rules and constraints of a workplace. Even if the technology works flawlessly, it will never be implemented if it can't deal with these real-life issues.

In this case, there were three key operational requirements for the solution: it needed to be 1) accurate, 2) quick, and 3) discerning. And these requirements were exacting: not just reasonably accurate or fairly quick, but orders of magnitude more accurate than the 99% that Mark Zuckerberg had dismissed, and fast enough to apply that accuracy at web speed.

The tech companies may have been overstating the case when they said a technical solution was impossible, but they weren't wrong about the scale of the challenge.

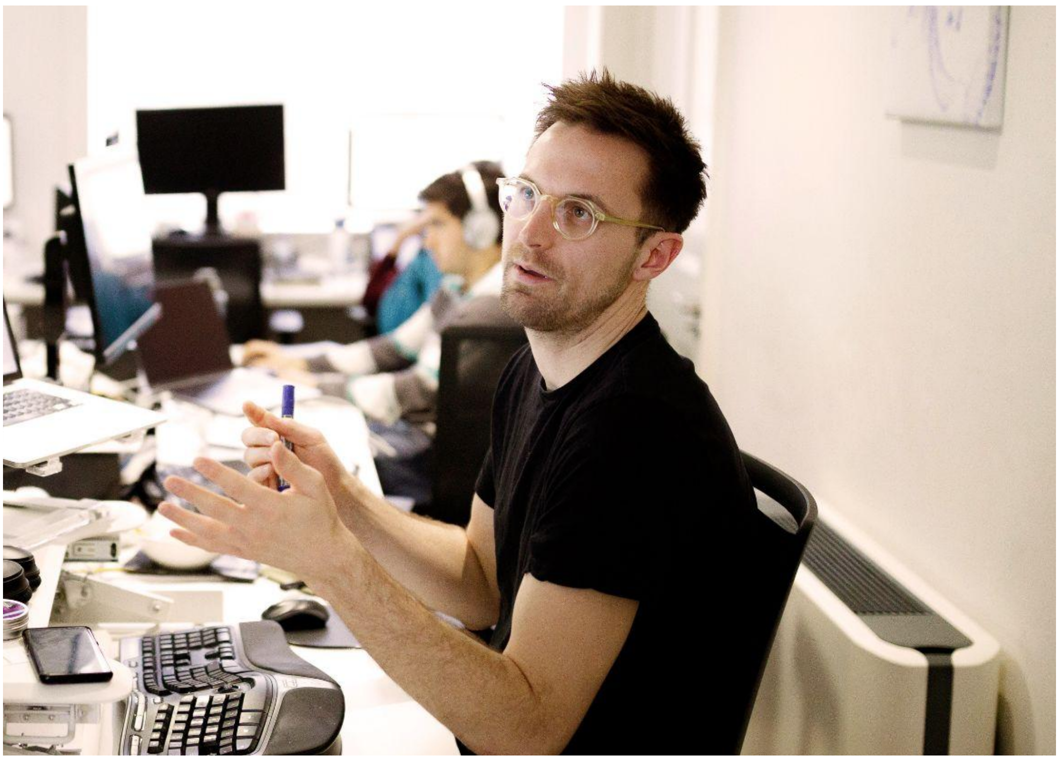

“There was a can-do attitude when we presented the problem, an enthusiasm to just get stuck in and really see what the art of the possible was.”

— Matt Collins, Deputy National Security Advisor for Intelligence, Defence and Security

OPERATIONAL REQUIREMENTS

The first requirement was for speed. Research produced by the UK Home Office and the European Union indicated that the first two hours after the release of new propaganda was the crucial window for disruption. In that time videos were amassing millions of views, and more than 90% of all the links to that content that would ever exist had already been created. If you couldn't find it straight away, you were already too late.

For the second requirement - accuracy - the key metric was 'false positives'. If the software labelled a video as propaganda, and in fact it was benign, the content creator would probably challenge the call. That meant unhappy customers, users denied videos they might want to see, and most likely the need for review by a human moderator. Beyond a certain threshold, too many wrongly-flagged videos would drown the moderators and anger the platforms' users - exactly the problem that Zuckerberg was pointing at too.

And that threshold for false-positives was low. Exceptionally low. Engineers at YouTube told the Home Office that they would only consider the technology feasible if the false positive rate fell to 0.005%. To put that in context, it meant that for every 100,000 videos the software analysed, it couldn't flag more than five incorrectly. This became the gold standard that the Faculty team worked towards.

The third requirement was more subtle, but no less challenging. The software had to be sensitive to very fine nuances in the content it was examining. The Home Office team had a profound understanding of every aspect of the terrorists' content, and the machine-learning model had to replicate that. But it also needed to recognise what Tom Drew and his colleagues knew about content that wasn't terrorist messaging, but might be mistaken for it.

Some types of entirely legitimate content resembled terrorist propaganda in specific ways that might trip up an algorithm. Worse, the most difficult to classify was also the content that Tom's team least wanted to remove. Islamic prayer videos and news reports of events in the Middle East, for example, might have certain similarities to jihadist content: censoring them would be not only controversial, but also counterproductive, because much of it actually served to highlight the flaws and hypocrisies in Daesh's messaging.

So the system Faculty were being asked to build had to be able to discriminate between superficially similar 'good' and 'bad' content; with a false positive rate of less than five in 100,000; and all within two hours of the content being posted.

And the team could never forget that the clock was ticking. In October, a man drove a pickup truck into a group of pedestrians on a bike path in Lower Manhattan. In Marseille, a man stabbed two women at the train station. Every month brought more grim reminders of the stakes involved.

So the system Faculty were being asked to build had to be able to discriminate between superficially similar 'good' and 'bad' content; with a false positive rate of less than five in 100,000; and all within two hours of the content being posted.

94%

The technology was able to detect 94% of the propaganda, with only a 0.005% false-positive rate.

0.005%

EXTRACTING MULTI-MODAL SIGNALS FROM VIDEO FILES

John's team at Faculty had a hunch that the reason the social media companies had failed to crack the problem was because they weren't looking at the content broadly enough. Rather than looking at any one single aspect of the videos, Faculty wanted to build an 'ensemble classifier', a model that would incorporate not only the relationships between particular attributes, but the relationships between the relationships. To do that, they'd have to wring every scrap of data they could out of the jihadist content.

But only from the content itself. The social media companies had a trove of data on individual users and who they were connected to, their networks and how they behaved online: all invaluable information that would help establish whether the content was terrorist-related. But the companies wouldn't share that data, and even if they would, Tom didn't want Faculty's algorithm using it. 'We wanted a solution that made a classification purely on the content itself,' he says, 'to prove that this could be done with the bare minimum data a government or third-party could capture - the media files themselves.' It also avoided any issues with accessing users' personal data, which would have tripped all sorts of ethical and regulatory safeguards. So all that Faculty had to work with was what any YouTube or Facebook user anywhere in the world could access: the videos themselves.

To train a model, you need a lot of data. The Faculty team worked with Tom and his Home Office colleagues, alongside leading academics and security analysts, to trawl the darker corners of the web to scrape up the target videos. As John recalls, 'By the end of the process, we had copies of pretty much every known piece of content that Daesh had produced.'

The details of what exactly Faculty did with that content are, for obvious reasons, secret. But in broad terms, they extracted everything they could from the data contained within the video. Were there certain types of song that were likely to feature on the soundtrack? Certain types of imagery or iconography, even specific people who could be identified? When the Home Office team were analysing content, they didn't overlook a single detail. So neither could the model.

Most excitingly, the Faculty team managed to extract the spoken word audio from the videos, transcribe it, and run natural language classifiers on the resulting signal.

The rapid developments in generative AI and natural language processing seen in the last two years have now made this task much easier, but in 2017 this was a game changer. As Tom explains, 'In any detection effort like this, you end up in what we call a "recursive adversarial dynamic": in other words, chasing a moving target. You detect a type of content, they figure out what you're detecting, they change it, and so you have to update your approach because it doesn't work any more.

'But if your model bases its classification on what the terrorists are actually saying in their propaganda - not just the words but the underlying tenets of their ideology and call to action - then it's extremely hard for them to escape that without changing what they're saying. And if you're forcing them to change what they're saying, it turns the game of cat and mouse into a strategic victory.'

RUNNING THE TESTS

After nine months, the model was becoming more sophisticated. John's team were optimistic that it could meet the exacting operational requirements they'd been set - but they needed to prove it. In particular, they had to demonstrate that the all-important false-positive rate was below the 0.005% threshold. To do that, they calculated, they needed a sample set of hundreds of thousands of randomly-chosen videos from across the global internet. So for several weeks they downloaded the first few thousand videos that were uploaded to the internet every hour of every day, avoiding any bias that might emerge from times of day or days of the week. Ed Sheeran, ping pong trick shots, cat videos, jihadist content... it all got scooped up and run through the model. Even the validation approach itself got tested and validated. The Chief Scientific Adviser to the UK Home Office vetted the process and pronounced himself satisfied.

The results were conclusive. The technology was able to detect 94% of the propaganda, with only a 0.005% false-positive rate. It could tell the difference between terrorist messaging and legitimate prayer videos or news coverage. And it could do it all in almost real-time, processing each video in the time it took to play. This meant that at the scale of YouTube's then five million uploads a day, only 250 would be incorrectly flagged - enough for a single human moderator to check. It was entirely operationally viable. And they could prove it.

A NEW STANDARD FOR TECHNOLOGY PLATFORMS

Thanks to OSCT's refusal to take no for an answer, Theresa May went onstage at the UN General Assembly armed with the knowledge that what she was demanding was possible. After she laid down her ultimatum to the tech companies, the world took notice.

Some of the attention was welcome. The Faculty offices hosted newspaper journalists and TV news crews in droves over the next few weeks, and the coverage they generated only heaped more pressure on the social networks to take the issue seriously. Within weeks, Mark Zuckerberg was publicly stating that companies like Facebook should be subject to more regulation, not less.

Some of the attention was less desirable, though flattering, in a way. Daesh had been following the news coverage too, no doubt trying to figure out why their content wasn't hitting its audience as well as before. Security insiders revealed that Daesh had nicknamed Faculty the 'Dogs of Deletion' in their internal conversations about the technology. The company took it as a compliment.

'I went to the west coast nine times in two and a half years,' says Matt. 'Each time, we had to put evidence on the table to further the conversation. And when we pitched them to say, "Well, you know you've been saying you can't do this, but we think you can, and here's what we've done," immediately the conversation went from a policy conversation, which was a bit binary, into a technical conversation, where they wanted to look underneath the bonnet and really understand what we'd done and how we'd done it.'
In the year following the release of Faculty's classifier, YouTube reported it was using AI to remove more than 80% of violent extremist content before it was flagged by users. Twitter was able to block 96% of terrorist accounts before they could even send their first tweet, and Facebook had implemented AI to take down 99% of terrorist content within 24 hours of its first release, much of it within the crucial first two hours. The Home Office also made the service free to a host of smaller social media platforms who lacked the resources of the bigger players. As Theresa May put it, it was 'a major step forward in reclaiming the internet from those who would use it to do harm.'

For Faculty, the small startup was suddenly on the map. 'Serious people in government were telling their colleagues, "This is a company you should be speaking to,"' recalls Angie. But the real significance of the project was felt in the wider world: for anyone who wanted to go to a concert, enjoy a drink outside, or just walk down the street without fear of being attacked with a knife or rammed with a car.

'It seems a long time ago now,' says Matt. 'But terrorism was the number one national security threat in that moment. We know the importance of social media in our lives, and some of the harm, unfortunately, that it can help facilitate. So demonstrating exactly what capabilities to bear to reduce the dissemination of that content definitely had a part to play.'

'A lot of people in different countries were looking at the problem, trying to get the tech companies to tackle it,' says Tom. 'And we subsequently heard independently from international partners that this was the big thing that really shifted the needle.'

THE LESSON IN SUMMARY

AI is an operational discipline, not an analytical one.

Many Data Science teams have their roots in Analytics or Business Intelligence teams. In these circumstances, a culture change is often needed to shift from an analytical to an operational mindset. AI is not being used to its potential where the output is a series of charts and dashboards.

◆

The objective of analytics is to understand the world. The great power of AI is that it can go a step further and operationalise that understanding, by turning it into action.

◆

In cases where there are high volumes of low value actions, this can be by automating processes directly. Where individual actions have higher value attached to them, AI should be built into the tools that people use to run business processes. In particular it should be integrated to support the decision points at which people intervene in those processes to better achieve their objectives.

◆

Where analytics supports human decision-making passively, by visualising trends, operational AI systems can provide active decision support. This can be by allowing people to test assumptions and scenarios or by running optimisations that result in a recommended path forward.

◆

AI systems that integrate with live business processes have to account for a more demanding set of operating requirements than analytics tools typically do. They come in many forms. User needs. Workflows to integrate into. Infrastructure, latency and security requirements. Policies and regulations.

02

NHS

LESSON TWO

AI is technology for human decision makers.

When the Covid pandemic hit in 2020, leaders in government and the NHS had to make decisions with terrifying implications, in an unprecedented situation, with data systems that had never been designed for a national health crisis. With infection rates soaring, AI was able to fill in the gaps and let decision makers get ahead of the curve.

They knew it was bad when they were allowed to keep their phones.

Normally, any visitor who walks through the famous black door of 10 Downing Street is asked to leave their phone in a rack by the entrance. But this was very far from normal times. It was March 2020, and Covid-19 had just locked down the country. In those early days, mortality predictions ran into almost unthinkable numbers. And the ultimate responsibility for stopping that happening - for making the decisions that would contain the pandemic and save millions of lives - rested with the people that Faculty were going to see.

Marc Warner and Andy Brookes, two of Faculty's co-founders, had come to Downing Street to talk about using AI in the Covid response. What they didn't know, as they walked up the famous stairs past the portraits of former prime ministers, was that the government was basing its decisions on a data-gathering system that would be recognisable to the bewhiskered, top-hatted Edwardian gentlemen in the paintings. Every night, millions of people watched the Prime Minister and his science advisors present the latest statistics on TV. But as the Prime Minister called 'Next slide, please', behind the scenes the data feeding the presentation was coming in from hospitals all over the country on scraps of paper. The CEO of the NHS would read them out for aides to scrawl on a whiteboard, and then the country's top scientists would use their iPhone calculators to project the likely trajectory of future cases.

In a fast-moving crisis, information is the oxygen of decision-making - and the government didn't have nearly enough. That's why they let Marc and Andy keep their phones that day they came to Downing Street, because phones were the best way they had of staying connected. Literally, a lifeline.

The Faculty team knew there was a better approach. Not only to get the data, but to use it in novel ways that would inform the decision makers in the literally life-and-death calls they were having to make.

DECISION-MAKING STRUCTURES FOR A NATIONAL CRISIS

The problem wasn't that the NHS was stuck in a previous century: it wasn't. But neither its decision-making structures nor its data flows had been designed for a national crisis. 'In normal times, no country would look at having this type of centralised capability,' says Lord Simon Stevens, the Chief Executive of the NHS at the time. 'You don't run your national health system as one big hospital. Obviously, in a pandemic, that may need to change.'

In other words, the NHS is designed to be operationally independent from government, and to provide localised services. Decisions are delegated down to trusts, hospitals and GPs' surgeries across England. But when Covid hit, the need for clear, effective and centralised decision-making led to the creation of NHS Gold Command, a committee led by Professor Sir Keith Willett. They met every day at 5.30pm at the NHS headquarters at Skipton House, a brown-marble-and-glass block that looms over Elephant and Castle tube station in south London. At the start of the pandemic, the normally-bustling building was eerily quiet, apart from a few executives and some army personnel who had moved in. But at 5.30, the building would echo with the sound of raised voices in heated discussion in the conference room on the top floor. Where should they send vital supplies like PPE, oxygen, or ventilators? Should patients be transferred to quieter hospitals? If more beds needed to be urgently freed up, should they cancel long-planned elective procedures, like cancer care or hip operations?

Gold Command brought together senior managers from across England, representing two hundred acute hospital trusts. Every one of the people there had dedicated their lives to serving patients: now they were dealing with a once-in-a-century pandemic, fighting for their share of gravely limited supplies, knowing that the whole health system might be overwhelmed any day. The cost of getting their decisions wrong was horrific; but even getting them right would have profound consequences for patients, staff and the wider public.

And all these decisions were being made with the sort of imperfect information Marc and Andy had seen firsthand in Downing Street. The data landscape was fragmented and fraught with inaccuracies. Overstretched frontline staff, forced to choose between spending their time on patient care or data entry, were of course choosing the patients - but that lowered data quality at exactly the moment when the system needed it most. NHS analysts found themselves desperately trying to extract insight from thousands of spreadsheets, many of which were being manually updated and frantically e-mailed around to be combined with other data before being presented to decision makers.

'I always remember the day I was called upstairs,' says Ming Tang, the Chief Data Analytics Officer for NHS England. 'Chris Whitty [the Chief Medical Officer for England] told us, "We need an infrastructure, a data store that brings everything together and then makes that data available to share across researchers. We need to know the state of the pandemic, and we need to be able to link that to health data to make sure that we know where the treatments need to be."'

As Prime Minister Boris Johnson lay in hospital having succumbed to Covid himself, Faculty worked with Ming's team and its technology partners to engineer a properly robust data infrastructure. Out went scraps of paper, iPhone calculators and email threads clogged with spreadsheets. In came a streamlined, real-time data pipeline underpinned by key data flows including positive case numbers, NHS 111 call volumes, citizen mobility data from mobile phone providers, and even genetic material sampled from sewage wastewater. This pipeline fed dashboards for each hospital site, which could then be aggregated for decision makers at trust, system, regional or national level.

'And within relatively short order,' Lord Stevens recalls, 'within about a week to ten days of deciding that we needed a centralised dashboard, we were able to assemble one. And I think we got there actually faster than most of the European countries.'

When Boris Johnson recovered, Faculty were able to demonstrate the full dashboard to him - although the data revolution hadn't quite swept through Downing Street. A screen had to be rolled into the Cabinet room specially for the occasion. It showed a level of detail and insight that decision makers - from local NHS leaders all the way up to the Prime Minister - had simply never had before. 'The Dashboard was so crucial,' Johnson recalled in his evidence to the Covid enquiry, 'that the 9.15 meetings [the government's daily ministerial strategy meetings] were later called the Dashboard meetings.'

But the pandemic wasn't going away. Clear, reliable data was a huge step forward, but it was only a start. Whether a number was written on a scrap of paper, or flashed up on a real-time dashboard, it still only told you what had happened. What the decision makers really needed, as they fought to get ahead of the next waves of the crisis, was guidance into how their decisions might play out in the future.

They needed the numbers to tell them what was going to happen next.

ENTER THE EARLY WARNING SYSTEM

This wasn't a new concept. NHS analysts had already tried to model future outcomes, and concluded it was impossible: certainly at the level of granularity that the NHS needed to make operationally useful decisions on a hospital-by-
hospital level. The data was patchy, and varied in quality across the country. With over two hundred large hospitals in England, it seemed an insurmountable challenge.

But the NHS team were open-minded, and with the stakes so high they agreed it was worth another attempt. A team of Faculty's top executives - including the CEO, CTO, Director of Health and Director of Data Science - decamped to the unused office space at Skipton House to be as close as possible to their NHS counterparts. Eventually the team swelled to some 20 people, almost a fifth of the young company's workforce.

'We created joint teams who worked on this,' says Ming, 'and those teams were fantastic in terms of helping us create the data science necessary. Faculty were very hands on, and they just rolled their sleeves up. We felt like one team. And that was really an uplifting capability for us.'

But time was against them. An exhausted country had emerged out of lockdown in July 2020, but as summer turned to autumn and cases started to rise again, it was obvious that the pandemic was gathering steam. New vaccines offered hope, but even on the most optimistic timescales they were months away from making a difference. As talk turned to 'circuit-breaker lockdowns' and 'tiers', it became clear that the executives in Skipton House would once again be making hard choices.

In the end, a technique known as Bayesian hierarchical modelling turned out to be an unlikely, unsung hero of the pandemic. Even in areas where there was almost no data available, it allowed Faculty to build a compound model, named the Early Warning System (EWS), that provided a sensible forecast by sharing information from nearby hospitals with similar characteristics. That approach dealt with both the inherent uncertainty in the data, and the challenge of trying to predict complex outcomes. Now the NHS could look three weeks ahead to see the bed capacity each hospital was forecast to have, where it risked running out and where patients or resources might need to be transferred.

'There were a lot of models being created across the system predicting where the virus was going and the rate of infection,' says Ming. 'But for us the focus was how operationally we would respond as an NHS, and so the model we created was much more important for forecasting beds, forecasting the likely impact of our staffing, forecasting which region would need to be most prepared. And then as we got the vaccine, that became really important because that then helped us identify where to put the vaccine next.'

For the first time, every level of the system was operating from the same page: from hospital managers making choices about how to allocate resources on their wards, right up to decisions being taken in Whitehall and Downing Street.

A MODEL EXPLAINABLE BY DEFAULT

But as these people looked at the data, could they trust the forecasts they were being given? When you're responsible for making decisions of this magnitude, it's not enough to be told what the computer says: you have to understand why. How did the algorithm reach its conclusions, and how confident can you be in what it's telling you? What struck the Faculty team again and again, as they worked on the project, was just how urgently their NHS counterparts needed to understand what the EWS was telling them. Why does it think that this hospital is going to run out of beds? What information is it basing that on? How confident should I be?

These are essential questions. They inform the basis of good decision-making, and any technology used for decision support needs to be able to answer them convincingly. After all, even with the most sophisticated AI model, it's humans who are ultimately the ones who make the big calls - and are held accountable for the outcomes. The AI is there to help them make the best possible choices. Which means that the technology has to be designed from the ground up to support humans and to keep them firmly in control. Most of all, it has to earn their confidence.

'We found there were three key ingredients to making the model trustworthy,' recalls Myles Kirby, then the Healthcare business unit director for Faculty, who worked on the project team. First, there was what they dubbed the 'decision-centric' approach. 'What a lot of analytical technology gets wrong,' says Myles, 'is it throws as many charts and numbers as possible at the user, and that's counterproductive. It overwhelms them, and distracts them with reams of data they don't need.' In contrast, the 'decision-centric' approach takes as its starting point the specific decision a user needs to make, and then identifies the precise set of analyses they need to make it better. If a particular analysis doesn't help the decision, it doesn't get included. The system is parsimonious by design - and, by design, it forces AI systems to be built in ways that serve the unique needs of users as decision-makers.

Secondly, Faculty built the technology robustly, so it could constantly be tested against actual outcomes. In effect, the users could 'rewind time': review what the model said at the point a decision was made, and then compare it to what actually happened. Crucially, the objective here wasn't to maximise confidence in the EWS, but to calibrate it. By being able to compare forecasts against actual decisions and outcomes, the NHS users could understand the right level of confidence to place in the technology, neither slavishly deferential nor unnecessarily sceptical. After all, even the most accurate forecasts - like people - are imperfect. Knowing how much you can rely on them builds trust.

Most importantly of all, Faculty made the model explainable by default. This is good practice for any AI system, and was absolutely crucial in this case. 'The explainability of it was really helpful in getting people to buy into it,' Ming confirms. For each forecast the EWS provided, users were able to click into it and see the relative importance of the features that drove the results. Was it leaning heavily on an increase in case numbers in a neighbouring town, or was it seeing something in the number of beds in use at the hospital? The model would tell you. The 'black box' had a glass door, and users could rely on their own judgement to look inside and check its workings.

And once you understand why the model's telling you what it is, you get new insights into the way things are changing that helps you decide how to intervene. One of the most interesting examples of this came when the EWS - and other models - predicted a Covid spike in a particular city in the Midlands.

'People in government were looking at all these models,' remembers Faculty's John Mansir, who was working as a Senior Data Scientist at the time. 'As soon as they saw this uptick, they thought they'd need a localised lockdown. But when we dug into the reasons why the model was predicting a spike, we saw the increased cases were all confined to a particular hospital in that city. There were none of the broader indicators that would imply the spike was spreading through the community. We suggested that the NHS investigate within-hospital transmission of the virus first, rather than assuming it was prevalent in the wider community, and in fact that turned out to be the correct diagnosis.' The right decisions were made, and the region was saved from a costly lockdown.

“The Faculty team were high-calibre, engaged, and flexible. They understood what the use case was that we were looking to develop, and worked with us to continually improve it... It was a distinctive contribution that was not made by anybody else to that particular problem that we needed to resolve.”

— Sir Simon Stevens, former Chief Executive, NHS England

ADOPTING A 'DECISION-CENTRIC' APPROACH

Faculty's Early Warning System quickly became the analytical centrepiece of the decision-making process. The decisions were still big, the stakes just as high, but Skipton House was a quieter place. When the 'Kent' Covid variant (later renamed the Alpha variant) ran rampant in January 2021, and London finally ran out of intensive care beds, the model was able to advise leaders where critically ill patients should be transferred by helicopter, based not only on where capacity was that day, but where it would be in three weeks' time and where the wave was likely to hit next. Even when SPI-M, the government's official modelling group, was forced to stop their work because the uncertainties had got too large, the Faculty model kept going.

Faculty's model outputs informed the allocation of over a billion pieces of PPE, facilitated the strategic transfer of critically ill patients across the country, and helped government leaders decide whether hospitals, towns and cities were opened up or locked down. In 2020 and 2021, these were matters of life and death, health and livelihood, for the whole UK population.

'The Faculty team were high-calibre, engaged, and flexible. They understood what the use case was that we were looking to develop, and worked with us to continually improve it,' says Lord Stevens. 'It was a distinctive contribution that was not made by anybody else to that particular problem that we needed to resolve.'

The model wasn't making decisions, and it wasn't offering infallible predictions. The reason it worked so well was because it had been built first and foremost to be decision-centric, to give officials and managers no more than they needed. It had been designed in such a way that users could learn how much to trust it, and so that they could interrogate how it had reached its conclusions. It was neither a crutch nor a replacement for humans using their judgement. It was a tool - but a tool unlike any other. Used correctly, it gave decision-makers the insight they needed to totally transform the speed, quality and execution of their decision-making - just when they needed it most.

'What was really valuable about the model was that we created a process around it,' says Ming. 'Every day we'd bring the emergency team that were actually dealing with the pandemic together with the data scientists, triangulating that information. And no model is ever perfect, but actually having a model and the gut feel and the experience in the room together to discuss it, we came up with a game plan that everyone was comfortable with.

'We built consensus around data,' she adds, 'which was really powerful, because it's the human and the data interaction that actually comes out with the best kinds of results.'

One day, in the later stages of the pandemic, Marc entered Downing Street for another meeting. A security guard stopped him, pointing to a telltale rectangular bulge in Marc's hip pocket. Embarrassingly, Marc had left the torch on, so a light glowed through the fabric of his trousers.

'I'm afraid you'll have to leave that at the door, sir,' the guard said politely.

Marc put his phone in the rack.

Ming Tang, Chief Data and Analytics Officer, NHS England

THE LESSON IN SUMMARY

AI is technology for human decision makers.

All software should be built around the user. The user for intelligent software is typically a decision maker. Focus AI on the places in which improvements in the speed, quality and execution of decision-making will improve business performance.

◆

Where a decision is important, human decision makers should remain in control and accountable. AI is there to support them, not replace them.

◆

Drowning people in data and dashboards doesn't help their decision-making. Instead, you need to be precise about exactly how the technology you implement will enhance their decision-making, and be parsimonious about giving them that and only that. At Faculty we use the Decision Loop methodology to make sure that solutions are carefully scoped to achieve this.

◆

Decision makers need to be able to judge how far to trust AI systems. Models must have the requisite level of explainability, so that users can see why it predicts what it does. Visibility of how accurately the model made historic predictions can also help calibrate how much weight to place on a model output.

◆

Interactive systems provide better decision support than passive dashboards. If model predictions are explainable, then decision makers can understand the cause and effect relationships at play in a given situation. And allowing them to see how outputs vary when inputs and assumptions change, means they can test the outcomes of different choices before they make them.

03

INSPIRED EDUCATION

LESSON THREE

Augment human tasks that require judgement. Automate those that don't.

A mountain of paperwork doesn't draw anyone into teaching - but it certainly drives them out of it. At Inspired Education, AI is freeing teachers from repetitive tasks and allowing them to focus on what they love best: helping kids learn.

At 10am on Monday morning, eleven year-old Millie Brown sits down in her classroom in west London, opens her laptop, and starts a test on the science she's been learning. When she's finished, she clicks 'Submit'. When the test is returned to her it's been marked, corrected and annotated with useful feedback to help her improve her answers next time - after being verified by her teacher, Mrs Roberts. In addition, Millie's scores have already been uploaded into the school's data system, so that Mrs Roberts can monitor her progress, and share it with Millie's parents.

The whole process - creating the test, marking it, providing feedback and uploading the results - has been done by AI.

In the classroom next door, Fred Jorgenson is teaching his students about the history of the US Civil Rights Movement. On the electronic whiteboard, he takes his class through a series of slides he's prepared. One shows annotated images of Martin Luther King leading the march on Selma. Another hosts a list of thoughtfully-designed activities to consolidate the lesson materials. The whole presentation is well designed and professionally presented - no corny clipart or clashing fonts - and designed to move at the right pace for his class, engaging the more able students while making sure everyone is able to keep up.

In his last job, Fred was sometimes up until past midnight preparing lesson plans and presentations for the next day. Now he's well-rested, with more energy to spend inspiring the kids in his classroom. He's still in charge of planning his lessons, making sure they're right for his students. But the heavy lifting is done by AI.

In most classrooms in the world right now, their experiences would be beyond imagination.

THE PAPERWORK CHALLENGE FOR MODERN TEACHING

Education resists change. Outside the school gates, pupils can access all the information in the world in the palm of their hand, and dream of jobs that didn't even exist when they started formal education. Inside school, things are slower to evolve. A Victorian school teacher who dropped through a time warp into an average 21st century classroom wouldn't have much trouble knowing what to do.

Nadim Nsouli is on a mission to change that. The Lebanese-British businessman is the founder and Chief Executive of Inspired Education, the world's leading premium private school provider. From a standing start in 2013, Inspired now operates over a hundred schools across 24 countries, from Ho Chi Minh City to Rio de Janeiro. On any given school day, some 8,000 teachers are teaching over 90,000 pupils in its classrooms.

After a successful career in law, investment banking and private equity, Nadim entered the world of education with an ambition to take a fresh approach to schools, rethinking traditional approaches to the curriculum and pedagogy. In short, preparing students for 21st century life by teaching them with 21st century methods and tools.

Nadim explains, 'The driving force behind the creation of Inspired Education was my vision to unite the world's best schools under one banner, facilitating the sharing of global best practices while preserving each school's unique identity and values. I took the leap from private equity to schools because I wanted to create a schools group designed to nurture well-rounded individuals who are not only academically accomplished but also confident and capable leaders. Our motto, "Embracing Individuality. Preparing Leaders," reflects this holistic approach.'

At the outset, Nadim was a newcomer to the world of education in more ways than one. Not only was he at the helm of a growing company in the educational sector, but he was also experiencing the system from a parent's perspective, as his child entered school. This dual vantage point revealed a recurring issue that he sought to address: teachers across the sector were overwhelmed with responsibilities, and parents struggled to access straightforward information about their children's progress.

Because while the time-travelling Victorian schoolteacher might take the teaching in her stride, what would really blow her mind is how much time teachers now devote to paperwork. Outside the classroom, they spend countless hours creating materials and slides for each lesson, marking hundreds of pieces of homework per week, writing reports and responding to parents' admin queries. For many, it crowds out all the benefits of the job. In the UK, for example, almost 20 percent of newly-trained teachers leave the profession within two years of qualifying, with many more considering quitting. Type 'teacher retention' into Google, and the first suggestion you get is 'crisis'. If Inspired were going to reimagine education, they first had to rethink teacher workload.

ENTER THE INSPIRED GLOBAL STUDY PLATFORM

In 2019, well before LLMs became mainstream and even before Covid introduced a generation of parents and children to online lessons, Inspired set about preparing the ground for the deployment of AI in their schools.

A key driver of this was Torben Lundberg, a straight-talking Dane with over 20 years' experience in IT management, who joined Inspired as Chief Information Officer. The company developed a modern Azure-based data platform across their schools that linked all pupil, parent and teacher data relating to teaching, learning and school administration: from individual children's grades and progress, to company-wide finance and HR systems. For a lot of sectors, putting all your data on a single platform is a no-brainer. For a group of schools, it was a world first.

In May 2021, Inspired acquired the online-only school InterHigh and combined it with its own virtual offering, King's College Online. These pioneering schools allowed pupils to access a virtual education from anywhere in the world. Torben and his colleagues took that expertise and built it out into a proprietary online learning platform - the Inspired Global Study platform - that would work for students in any of its schools. It complements the world-class teaching in Inspired's classrooms with the wraparound benefits of remote learning, so that pupils can easily catch up on missed lessons, access additional content and stretch material, and submit homework assignments. Crucially, it also helps parents support their children by giving them fine-grained information on what the kids are doing and how they're getting on.

In fairness, none of this was completely revolutionary. All schools have management information systems, finance systems, HR systems and so forth; and most now have some kind of online learning, behaviour and parental engagement platforms, even if it's just paying for school meals and assigning homework. The difference is that most schools buy in a suite of different products that don't talk to each other: the data is fragmented and unstructured. Where ten years ago parents had to dig around for change for lunch money, now they have to remember the login details for umpteen different platforms. At Inspired, all the data is on one platform owned by the company.

There's one other, crucial difference about Inspired Education's approach. They laid a foundation, so that when LLMs hit the mainstream in 2023, Inspired had everything in place to leverage its potential for education in ways that had never been done before. And they were ready to move fast.

AI'S IMPACT ON CORE PROCESSES

Torben came late to the education business, after spending the first part of his career working for media companies, starting at the same time as the internet began to take off. 'I spent 20 years on digital transformation in print and television, digitising all that,' he recalls, 'but that sort of finished. So I wondered, which other industry was ripe for a similar level of transformation?' As a parent of four children, Torben could see that education fitted the bill. He wanted to be a part of changing it.

But he hadn't foreseen the explosion of AI. Inspired's platform had been built to use well-established adaptive learning software from third-party providers. Torben freely admits that when ChatGPT burst onto the scene in 2023, even a forward-looking company like Inspired had no generative AI strategy.

'It hit me like a revelation,' he says. 'I could see it would be big in many areas of our company, but particularly the core academic processes,' so much of which are based on the written and spoken word. 'And the reason we've been able to move much faster in AI than other educational companies,' Torben adds, 'is because we had the platform. No-one had as much data as we did.' Grasping the implications of GenAI, he established a series of 'speedboat' projects that would deliver quick-turnaround applications for Inspired's three key stakeholder groups: students, staff and parents. First up: the teachers.

THE LESSON PLANNER

To understand the landscape, Faculty helped Inspired survey school leaders and frontline teachers. Over 700 responded, revealing that teachers were spending an average of over six hours per week on lesson planning, and over four hours on creating and marking weekly tests. 'Teachers can be quite cynical about new technology,' Torben observes, 'and with good reasons. A lot of technology doesn't really help them. It's more tracking, it's more information that somebody else needs, and it doesn't really help that day-to-day work in terms of teaching.' But the surveys showed that teachers were willing to embrace AI if it helped them - not only to automate and speed up repetitive tasks, but also as a way of stimulating their creativity in lesson planning, or helping them differentiate lessons for different ability levels within their classes.

The targets the Faculty team were set for the Lesson Planner project were ambitious: to build an application that could automatically generate high-quality lesson materials in a way that replicated Inspired's teachers' best practice, and saved them time. These lessons would have to cover every subject in the curriculum, across different topics and year-groups.

Inspired asked their teachers from around the world to share their best lesson plans, to be analysed by the AI model. Were teachers reluctant to hand over their prized materials, some of which had been refined over years of practice, to train an algorithm? 'Not really,' says Torben. 'Their approach has been that the collective sharing of this information helps everybody, not just that particular teacher but also new teachers coming into the trade, who otherwise have to build everything up from scratch.' The teachers responded with over 50,000 examples.

With those to work off, Faculty developed a proprietary LLM modelling framework for Inspired. Following best-practice learning design, it starts with the teacher. They create and edit a lesson plan with help from AI, and then type in some simple lesson objectives in a few sentences. When they're done, the model goes to work. Built on top of GPT4, it breaks down the content generation system into multiple steps. It uses Retrieval Augmented Generation techniques to work out which are the most relevant lessons in the databank; draws the appropriate content from them for structuring lessons; then generates the lesson content according to this plan. It creates teaching slides, in-class activities and assignments, all formatted to Inspired's design template. Everything is ready for the teacher to review in under a minute. And because it's all on the single study platform, lessons are available for any pupils who need to catch up, or want to review the material after class.

But as much as you try to teach them, how can you tell what's going in?

Every time a teacher or a student engages with the lesson-planning and Cycle Test tools, the systems learn and calibrate the models, constantly improving their ability to generate appropriate lesson materials, test questions and useful feedback.

THE CYCLE TEST GENERATOR

Visit any Inspired school anywhere in the world at 10am on a Monday morning, and you'll find every pupil over the age of nine sitting a 'Cycle Test' to assess what they've learned in the past week. These have been part of the routine at Inspired ever since it was founded, allowing teachers to track progress in fine detail and give formative feedback, and getting children used to the discipline of test-taking long before they reach the high-stakes exams of later years.

Taking the tests is a lot of work for the kids. It's even more onerous for the teachers: writing the tests, marking them, entering marks in the system and providing the kind of rich feedback that makes a difference. Over 500,000 tests need to be generated, leading to two million scripts that have to be marked every year.

Enter the Cycle Test Generator. Using existing test examples, and user research with teachers, Faculty and Inspired created an AI application that covers the whole process from start to finish. It allows teachers to choose the subject and topic, then automatically generates questions in line with the British National Curriculum across the full range of question types: multiple choice, show-your-working, retrieval and (coming soon) essays.

Once the teacher has edited the test and is happy with it, they can securely share it with their students. When a student's completed it, the proprietary marking engine marks the answers, and provides written feedback down to the question level. Everything is automatically recorded on the learning platform, so that parents and administrators can track progress. The teacher still retains the power to review, edit and configure every aspect of the process - but in a fraction of the time it used to take.

The Cycle Test Generator is being piloted in over 25 Inspired schools and is already transforming teacher workload. When it's fully rolled out, Torben estimates it will save staff over a million hours a year, time which can be reinvested into teaching. And that's just the beginning.

Every time a teacher or a student engages with the lesson-planning and Cycle Test tools, the systems learn and calibrate the models, constantly improving their ability to generate appropriate lesson materials, test questions and useful feedback. All of which saves teachers time, and helps kids learn.

SETTING GOALS BEYOND THE CLASSROOM

Torben's goal is to cut staff lesson-planning time, homework-marking time, test-preparation and marking time all 'in half'. In terms of pupil outcomes, Inspired has already seen an 8% rise in performance - equivalent to an entire grade boundary - thanks to its wider investments in education technology. They're now thinking about training the AI in other languages and other curricula. And they've barely scratched the surface of what their data can yield.

As well as the 'speedboat' applications, Torben has asked Faculty to take on a 'supertanker'. This is a long-term, strategic project to structure every piece of data held anywhere in Inspired's systems, so it can all be aggregated and analysed for insights into how children learn, and used to correlate insights from students around the world. 'We are probably the world leader in terms of having the amount of structured data required to do this at scale,' Torben points out, 'and we can use that information to supplement the learning curve for each child.'

Nor is he looking to keep the software and the insights proprietary. 'We have these top private schools, so we would be a logical first mover investing more on this up front, but we are very conscious that this can also help the wider system. And our intention and approach is for this to escape out to the state school systems.'

It's very likely that education will change more in the next ten years, thanks to AI, than it has in the last five hundred. As Inspired shows, it has the potential to make teachers better at what they do well, and free them from what slows them down, to the huge benefit of the kids they teach.

Everyone can learn from that.

THE LESSON IN SUMMARY

Augment human tasks that require judgement. Automate those that don't.

AI will almost certainly impact most jobs. But choices about the way the technology is deployed will determine how these impacts are felt. And for many roles, there are paths you can choose which will have positive outcomes for workers who are affected.

◆

Most people enjoy the core of their job. It's what they're good at and it's an important part of how they create value to the world. But many of these people will tell you that admin, reporting and other bureaucracy take up too much of their time, and crowd out their ability to focus on the core.

◆

This is going to change. The kind of tasks that AI is extremely good at map well to the tasks that are responsible for this crowding out. Well designed AI programmes build AI into workflows in a way that automates away the low-value, routine tasks.

◆

Thus people will be freed up to focus on the high-value tasks at the core of their roles. In most cases these tasks require professional judgement. It is usually unwise to try and replace this human judgement. Instead, AI should be used to improve the speed, quality and execution of high-value decision-making.

◆

This blend of automation and augmentation offers a vision for the future of work that is both more productive and more fulfilling. By the end of the decade, most cognitive workers in the economy will have many fewer low-value, routine tasks to do, and will have Intelligent Decision Support augmenting their performance of high-value tasks.

04

DRIFT

LESSON FOUR

AI is a feature, not a product. But it can define a product.

A fleet of autonomous ships that sail the high seas to find new sources of energy sounds like science fiction. DRIFT is the visionary startup using AI to make it a reality.

In a landlocked corner of England, Ben Medland walked across a field and saw the ocean. It was a calm summer evening in the rolling countryside, and he was out with his six year-old son, James. As Ben spoke to him, he realised that his son was deeply concerned with the climate crisis he had heard about on the news. He pointed to a wind turbine on the horizon, its blades sitting motionless in the still air. 'They need to turn it on.'

"But there's no wind," Ben explained.

"Then why don't they make one that follows the wind?" James asked, with a child's innocent logic. Ben paused. Why don't they?

MOBILE WIND-POWERED ENERGY FARMS

Often Faculty are asked to apply AI to existing businesses. Sometimes, they get to help build new businesses from thin air. But for one memorable client, thin air is the business.

DRIFT Energy is the company that Ben founded as a direct result of that conversation with his son. It's just the latest step in a career that's taken him from advanced projects at BAE Systems, to building Accenture's Digital and Data Strategy unit into a $100m business. A tall 43 year-old with a mop of black curls and an infectious smile, Ben has been solving problems his whole career, whether those problems were engineering, business or technological. Now he's set his sights on the biggest challenge of them all.

The race to reduce the carbon we're pumping into our atmosphere requires every solution humanity can throw at it. Incremental improvements won't be enough. It needs new and creative approaches that rethink the fundamentals of how we produce, transport and consume energy. Like a turbine that can go where the wind blows, instead of just waiting for it to appear.

Ben's vision is to build an unmanned sailing vessel that will operate as a mobile wind-powered energy farm. But the idea's grown more complex since that lightbulb moment with his son. If you're imagining a classic three-bladed windmill lashed to the back of a ship, think again. The power comes from a turbine slung under the vessel. As the ship speeds through the water, driven by the wind, its kinetic energy pushes water through the turbine's rotor to generate the electricity: a windmill and a watermill all rolled into one.

Or as Ben puts it: 'A ship with a propeller where the energy goes the other way.' But electricity generated in the middle of the ocean doesn't really have

anywhere to go. So the second piece of Ben's design is to have the electricity from the turbine power an electrolyser, a neat piece of equipment that sucks in water from the sea and splits it into hydrogen and oxygen. The oxygen is released, while the hydrogen is pumped into storage tanks in the ship's hull. When the tanks are full, the ship cruises back to port, offloads its cargo, and sails out to do it all over again.

Hydrogen has long been forecast to play a major role in efforts to reduce carbon emissions. Unlike most renewables, it can be stored long-term and transported easily. It can be used instead of oil in hard-to-electrify industries like heavy transport, aviation and - relevantly - shipping; and it can decarbonise industrial processes like steelmaking and cement production. If produced using clean electricity, it has virtually no damaging emissions.

The drawback is that there are precious few natural sources of pure hydrogen on Earth. While hydrogen is the most abundant element in the universe, on our planet, it is almost always bonded with something else. It has to be split out from water, biomass or fossil fuels - and, under the unyielding laws of thermodynamics, it will always take more energy to extract the hydrogen than you will get from using it.

Which is why it makes so much sense to produce it in places where there's unlimited free clean energy. Places where you don't get bogged down in difficult planning or permitting issues, and where you don't need any costly fixed transmission and distribution infrastructure. Places where the natural feedstock - water - is available in almost unlimited quantities. Places like a vessel in the middle of the ocean.

GETTING THE 'RIGHT WIND'

When Ben approached Faculty to discuss the idea, they were cautious. 'I could see them thinking, "This is a bit mad",' Ben recalls. Andrew Perry, head of Faculty's Energy Transition and Environment business unit, puts it more diplomatically. 'We thought it sounded amazing, but also extremely ambitious,' he recalls. 'How could so many different technologies be combined and optimised to work together? How would it work commercially?'

In his line of work, Andrew hears a lot of blue-sky thinking from visionaries who think AI is a magic wand they can wave at the climate crisis. But as Ben laid out his ideas, Andrew started to buy into it. 'The more we discussed and understood the concept in Ben's mind, the clearer it became that his vision had real substance, and the model made sense.' He pauses. 'So long as you could get the right wind.'

Ben's hunch was that his sailing ship would be able to significantly outperform the efficiency of fixed offshore wind turbines, maybe even doubling it. But his whole idea rested on the vessel being able to plot a course to catch the best wind for the longest time. If it couldn't do that, the economics of the project would never stack up.

'First and foremost it's a data science challenge,' says Ben. 'We looked at all the off-the-shelf options, nautical routing software and so on. For what we needed, there's nothing like it out there.' Which is unsurprising: most routing software is focussed on getting the vessel from A to B as efficiently as possible. It doesn't cope well if, for example, A and B are in the same place. For DRIFT, it would be all about the journey.

That's when Andrew understood why Ben came to Faculty. 'He wasn't just building a ship. He was building an autonomous vessel, one that could optimise its route in real time in the chaotic weather conditions of the North Atlantic.' He needed AI to fill in the details.

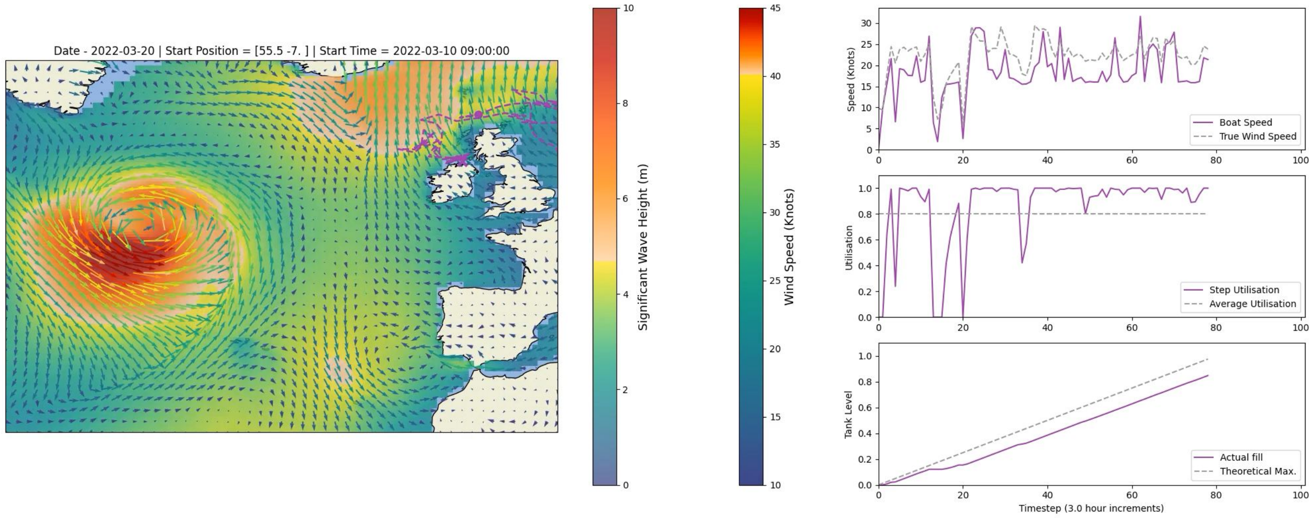

DESIGNING FOR SPEED IN ALL WEATHER CONDITIONS

The first job for Faculty was to test Ben's preliminary hypothesis about how efficient the turbines could be if the ship was sufficiently intelligent in its routing. Instantly, they ran into a new challenge.

A sailing ship like Ben was planning had never been built before. Without a legacy design constraining them, DRIFT had a unique opportunity to take a completely fresh direction, to wring every knot of speed out of it. The faster the boat could go, the more kinetic energy it would transfer to the turbine and the more efficient it would be at making hydrogen. It would take its cues from the high-performance yachts that race in competitions like the America's Cup and SailGP, boats which go so fast they literally fly over the water on foils. Only this time the technology is directed at making the ship go greener.

But working from a blank sheet of paper is also incredibly daunting, especially when you're developing something as complex as a hydrogen-producing sailing vessel. Starting from scratch throws up a multitude of choices. Small or large? Single or double-hulled? How many sails? How many turbines? What shape? What materials? And those are just the basics. There are multiple layers of decisions, that each knock-on to each other, and you have to resolve them all before the final ship can take shape.

To guide their design choices, DRIFT needed to understand how different options would affect the ship's core purpose: to get the most speed from the widest range of weather conditions. From mild, balmy days where the air is still, to wild nights when storms sweep across the ocean, the vessel design had to deliver the best average speed across all the journeys it might take. And it had to be affordable and practical to build.

It is possible to model how a given boat might perform in different weather conditions. But it's not trivial to optimise the vessel's performance model at the same time as optimising the route it should follow. Suddenly things get very complicated, very quickly. The best sequence of moves for one particular shape of ship might be no good at all with a different design. It's a version of the classic 'travelling salesman' problem, except that you've got no idea how the salesman's getting about and you don't know where the customers will be.

Ultimately, Faculty's modelling and DRIFT's design loops fed into each other. Faculty could test how different types of vessels would perform, and DRIFT could test the trade-offs between performance and commercial feasibility to develop their understanding of what to build. The digital and the physical sides of the project were inseparable.

DRIFT's first ship would cost tens of millions of pounds to build, so there was no room for error. Not in the construction of the vessel, and not in the algorithm.

To model the vessel's likely hydrogen output, Andrew's team dug deep into the science of seafaring to understand which factors contribute to performance. But the model wasn't predicting weather. 'There are supercomputers out there that forecast the weather better than we ever could,' says Andrew, 'and those feed straight into the algorithm.' The model's job was to find the 'Goldilocks' zone, where the wind is neither too strong nor too calm, but just right. A light wind is clearly no good for propelling the ship at the speeds it requires, but excessively windy conditions might damage the vessel or wreck it completely. Key to the problem is the fact that the same wind has very different effects depending on which way the ship is pointing. Every sailboat has its optimum point of sail, the angle relative to the wind that generates the most speed. A boat facing straight into the wind will go nowhere. As it changes direction, it can gain or lose speed; for most vessels the optimum course is at about a 90-degree angle to the wind.

A sailing ship going from A to B will often have to choose a sub-optimal point of sail in order to get where it's going. But an autonomous hydrogen-producing energy yacht can go in whichever direction it chooses. It's got nowhere it needs to be.

But of course, at sea conditions are always changing. As well as getting the most out of the wind conditions in the moment, the ship also had to be planning ahead. The whole point of Ben's vision was to follow the wind. And that wasn't just about forecasting the next best action to go to the nearest good bit of wind. The algorithm had to think many steps ahead, strategically assessing thousands of possible route options many hours in advance, making sure it wasn't sacrificing long-term performance chasing near-term gains. Like a game of chess, except with a board as big as the ocean and almost no constraints on moves.

And even then, it couldn't necessarily just take the best route through the Goldilocks zone. Any time the ship is at sea with a full tank of hydrogen, it's wasting time. So the algorithm had to find a route that would take the ship back to port, ideally arriving at the exact moment that its tanks hit capacity.

Having developed the prototype algorithm after an intensive seven-week sprint, Andrew's team tested its performance against thousands of potential voyages. The results were unequivocal. On its simulated journeys, the boat was able to achieve a load factor - effectively the proportion of the time the turbine is able to operate - of 70-80%, compared with 35-40% for fixed offshore wind turbines.

Ben's hypothesis was right.

Now the hard work could begin.

Ben Medland, Founder & CEO, DRIFT

AN ALGORITHM FOR 'THE IMPERFECT REALITIES OF THE PHYSICAL WORLD'

Demonstrating the potential of DRIFT's concept helped the business close its seed funding round, led by the aptly-named Octopus Ventures and supported by Blue Action Accelerator, which invests in novel ocean and climate-friendly technologies. £4.65m was unlocked to continue detailed design work on the first ship, and to improve the model.

'Our initial prototype was a relatively rough and ready model,' explains Andrew. 'It was necessarily pragmatic, to prove the concept at minimal cost. It focussed exclusively on wind conditions, and how they would impact on speed and hydrogen generation.'

But to go to the next stage, Ben needed the algorithm to be bulletproof. The initial modelling had shown that a 58-metre catamaran was the optimum design, but there were literally thousands of detailed design decisions that flowed from that. Those choices would be made based on how the different options performed in the route optimisation model, so Andrew's team needed to make sure the model accounted for every relevant factor. This included things like wave height, sea state, tides, currents and swell - as well as how the interplay between all of those factors would affect the ship's ability to harvest energy.

They're all represented on a visualiser, which has a kind of hypnotic beauty when you look at it. Hundreds of arrows make whorls across the screen, their length and colour changing to show the wind's speed and strength. The ship, a little purple dot, leaves its base on the west coast of Scotland and strikes out over the top of Ireland, zig-zagging dotted lines across the north Atlantic. A pulsating red blob sweeps in, representing dangerously high waves. The boat retreats, around northern Scotland and back to the Orkneys, taking shelter in the lee of one of the islands where it's shielded from the worst of the conditions, as mariners have for centuries. When the colours subside to a more agreeable yellow, it heads back to port.

This is still a work in progress. Ultimately, the team is aiming to get the model of the vessel and its environment to the sort of level that Formula 1 teams operate at, where every last detail can be simulated and tested. That even includes coming up with a pit stop strategy to minimise the time the boat is stuck in dock while the full hydrogen tanks are unloaded and replaced.

But unlike a Formula 1 car, there's no driver in the cockpit: DRIFT's vessel will operate autonomously. That means the algorithm is far from an academic exercise. It'll have to perform not just in the calm waters of Faculty's data lakes, but out in the real ocean. Using only video, radar and sensor inputs, it'll have to plot a course that follows the optimal wind, while dealing with every hazard to navigation the oceans can throw at it.

For Faculty, that means forensically examining every aspect of the simulated journeys, to understand where refinements need to be made. Boundaries need to be set for the shallow waters near coastlines. The vessel has to be aware of obstacles. Even routine sailing operations like tacking and gybing - changing direction, in layman's terms - create inefficiencies that need to be accounted for.

'It's all the imperfect realities of the physical world,' Andrew explains. 'The list isn't endless... but it is long!'

And it's absolutely essential the algorithm can handle them all. Because Ben doesn't want to launch one ship: he wants to launch whole fleets of them, hundreds or even thousands strong, that can deliver hydrogen to the four corners of the globe in quantities big enough to tip the scale of global warming. His goal is that one day his catamarans will be as synonymous with renewable energy production as the world's 340,000 wind turbines.

'Each ship will deliver roughly 100 tons of green hydrogen every year,' says Ben. 'Run through a fuel cell to produce electricity, that's enough to power up to 1000 UK homes, or for hydrogen-powered cars to drive 7.1 million miles.' And it's 1.2 million kilograms of CO2 that won't go into the atmosphere.

The current plan is to lay the keel for the first vessel in late 2025, and to build it within 18 months for a potential launch by summer 2027. But long before it sets sail, DRIFT has already been on an incredible journey against the odds, travelling over 15 million miles on its simulated voyages. 'And the technology on DRIFT will only get better,' Ben points out. 'The sail performance, the turbine performance, the hydrogen plant, the energy chain - and the costs of all of those - are going to improve over time as the market evolves. The data and the speed of the compute will improve. So there's an awful lot of tailwinds behind the company.'

There's a long way to go, but Faculty is proud to have supported Ben and the DRIFT team this far, and hopes to keep helping them every step of the way ahead.

Wherever the wind takes them.

THE LESSON IN SUMMARY

AI is a feature, not a product. But it can define a product.

The most advanced language models are very impressive at what they do.But despite this, there are few occasions where the things they can do out of the box - summarising or creating new text - correspond to the things that are most valuable. The same is true for other types of AI; from time-series forecasting to computer vision.

◆

As a result, there are few occasions where an AI model alone makes a full product. Use cases that start with some impressive thing that AI can do, and try to narrowly slot it into a business process will disappoint.

◆

Instead, it is best to think of AI as a cog in a machine, rather than the machine itself. A piece of functionality that can be connected together with others into a piece of software whose functionality goes much further than any of the component parts.

◆

However, AI is an unusually powerful cog that can make new kinds of machines possible. In much the same way that an engine isn't a car, but it was the thing that made the whole paradigm of automotive transport possible..

◆

As a result, this is a good time to seek new ways of solving old problems. Even where AI plays only a small role in a piece of software, inside a business process, it may be the unlock that allows you to change the entire way the process runs.

05

BEAZLEY

LESSON FIVE

Building AI widgets is easy. Rewiring your organisation is not.

Lots of companies try to bring AI into their businesses. Most fail. Specialty insurer Beazley is blazing a trail in making sure that AI investments take root, and offering a blueprint for how firms everywhere can use AI to transform their operations.

In July 2024, The Economist diagnosed a new corporate ailment. AI 'pilotitis', it wrote, is 'an affliction where too many small AI projects make it hard to identify where to invest.' Afraid of being left behind by the AI hype, companies were launching a flurry of ballyhooed AI initiatives, then quietly abandoning them when the promised transformation failed to magically appear. 'The incorporation of AI into business processes,' The Economist drily noted, 'remains a niche pursuit.'

One company determined to remain immune to 'pilotitis' was Beazley, one of the world's leading specialty commercial insurers. Specialty insurance, though often overlooked in financial services, plays a vital part in keeping the world's commerce moving. From oil tankers running aground, to Taylor Swift concerts being cancelled, specialty insurers are there to mitigate the most complex risks of operating in the global economy.

Beazley has been doing it for over 40 years. The FTSE 100 business is widely recognised as one of the most innovative, sustainable and successful operators in its industry. From their offices in London's Bishopsgate and around the world, they write billions of dollars in insurance coverage every year across almost every sector of the economy. They insured the first private lunar lander on its journey to the moon, and Ukrainian grain ships making the perilous passage through the Black Sea. They're particularly known for their offerings in cyber risk.

But although specialty insurance involves managing cutting-edge risks, the sector has not typically been seen as a hotbed of innovation. While life assurance or general insurance might generate quantities of structured information that are easily amenable to digital processing, in Beazley's field every contract is different, and every claim is a unique set of circumstances.

GENAI AND THE RISK OF 'PILOTITIS'

The rise of Large Language Models (LLMs), with their ability to process, analyse, and organise vast amounts of unstructured data, presents specialty insurers like Beazley with an opportunity to upgrade their businesses. They can see their way to automating parts of underwriting, customer service and claims processes, and enriching the careers of their highly skilled teams.

But it needs to be done with care. The user-friendly, extremely customisable nature of tools like ChatGPT makes it easy for companies to try out AI-powered proofs-of-concept for their business. Demonstrating such a process is one thing; actually embedding it in operations, making it both useful and usable, is quite another. Beazley took a considered, thoughtful approach, based on its track record of pioneering new areas of insurance and risk management (like cyber), and realised that to effectively use AI they would need to pilot, understand, test, risk manage and then build it out comprehensively.

SETTING A COMMERCIAL MANDATE FOR AI

The role of specialty insurance is to help clients mitigate their toughest problems, their biggest risks. At Beazley, that emphatically includes technology. So they set about managing the risks of the project in a systematic and methodical way.

The main risk was one common to any serious change programme: a lack of clear leadership and direction. Muddled thinking can derail any project, but AI projects are particularly susceptible to it. Heads can get turned by the latest shiny demo. Experimentation happens in pockets around the business, driven by the cool things the technology can do, rather than the important problems that leadership care about, or that staff need solving.

Beazley set things up so there would be strong leadership and direction from the off. Staff were given a clear commercial mandate that was unequivocal about the priorities for using AI to enrich the careers of Beazley's highly talented employees.

- Increase the throughput of core operations.
- Support the teams' decision-making.
- Help reduce the number of risk incidents across the business.

Help reduce the number of risk incidents across the business. The simplicity and clarity of these three objectives provided a north star that guided decisions about everything from the overall shape of the programme down to day-to-day task prioritisation in project teams. Every day they asked themselves, 'Does the thing I am planning to do today increase throughput or accuracy, or reduce risk?' If so, do it. If not, don't.

Troy Dehmann, COO, Beazley

IDENTIFYING THE RIGHT USE CASES

That took care of direction. The leadership half of the equation was addressed by centralising all of Beazley's internal AI efforts into a single process. There would be one program, with one set of priorities, captured in one roadmap reporting into one leader.

That leader was Troy Dehmann, the genial South Carolinian who's Beazley's COO. He came to the world of insurance after a career mainly oriented towards finance, and admits, 'I hadn't heard of Beazley before I first came to interview. But I was drawn to the firm's ambition to grow and modernise, their willingness to put investment into infrastructure and things like AI that were coming down the pipeline. It's also very employee driven, and puts the employees first.'

Troy makes the point that AI didn't suddenly drop into the world with the advent of ChatGPT. 'We were already doing data science three years ago when I joined, and it wasn't something new. So it was never a question of "doing" AI or starting to use data science. It was actually about scaling it, and we've ramped that up over the past three years, and certainly since we started to work with Faculty.'

Troy knew he had to get the program off with a bang, in order to start the process of winning hearts and minds across the organisation. His first objective was to deliver four operational AI processes into the hands of users as soon as possible. Though as it turned out, finding use-cases was the easy part. The difficulty was narrowing the choices down to four.

Faculty investigated where AI could deliver on the business's priorities, and found there were literally hundreds of places that could benefit across the core workflows of underwriting, claims, and operations. In many respects, that's not surprising. Ever since history's first insurance policy was carved onto a Babylonian obelisk around 1750 BC, insurers have been processing written and numerical information. Nearly four thousand years later, as functional language models have developed to complement AI's already mature numerical capabilities, the insurance sector is well placed to benefit from the technology.

But a longlist of hundreds is ultimately unhelpful if you're trying to run a tight, focussed process that actually achieves something. So the Faculty team filtered the longlist using Beazley's three North Star priorities, coupled with factors like technical feasibility, data availability, and the degree of change that would be needed to implement a given option. That brought them down to about forty options, any of which could have delivered substantial efficiency and quality benefits. That still wasn't good enough.

A ROADMAP FOR COLLECTIVE AI CAPABILITY

Troy was clear that the AI project shouldn't be about automating narrow slices of existing processes. He understood that the real opportunity was to completely rethink how workflows should operate end-to-end, in light of what the AI could do. 'If you don't have a roadmap of where you're going to use it, and you're just deploying it haphazardly across your processes, you're not going to recognise the full value.' That meant instead of simply rushing ahead with the most compelling of the forty shortlisted use-cases, he wanted the Faculty team to connect up all the disparate AI threads, so that the value they generated was greater than the sum of the individual savings.

To take one example, there's a narrow use-case for using AI to extract information from the many different types of document that the Beazley underwriting team receive from their brokers. On its own, it creates significant savings by reducing the need for the Underwriting and Ops teams to manually review all the submission application documents.

But that's just scratching the surface. The real benefits come when you view that process as one a series of interconnected AI use-cases, and reimagine the way the whole underwriting pipeline works. Once you start thinking along those lines, you can incorporate automation far more widely to produce a quote that is quicker than the old way of doing things, more accurate, and that brings lower chances of mispricing the risk.

As this type of process re-engineering played out across all the core workflows of the business, conversations moved away from individual AI 'use-cases' and towards collective AI 'capabilities' and 'workflows'. Instead of bolting on shiny widgets, Beazley was thinking about how to redesign the whole machinery that made its operations run.

And it turned out that a lot of that machinery used similar parts. Even across very disparate parts of the business, the underlying mechanisms involved similar types of generalised algorithmic tasks: functions like processing e-mail, extracting information from unstructured forms, or parsing long documents. Because each separate use-case had slightly different requirements, there was a temptation to build a custom version of the AI each time. But that would have created enormous duplication of effort, and a huge technical debt: it would have been hard, for instance, to roll out new LLM model improvements, because each custom implementation of the technology would have needed updating and testing separately.

'That was the main challenge, but also the biggest success,' says Laura Palacio Garcia, the senior data scientist on the project. 'In the early versions of the utilities, updating them all could get messy. But we were able to build a unified front end, and shared repositories for both the front- and back-end code.'

Faculty helped Beazley break down the broad, general-purpose tasks into modular, reusable AI 'utilities'. This provides a living library of technical components - custom to Beazley's needs, but common across their AI program - that can be tailored for each AI application. Data science teams can quickly stitch together use cases based on templates and pre-existing modules, rather than building from scratch. Not only has this made the rollout much faster, but it's also made for better quality, greater simplicity, and further economies of scale across the whole AI catalogue. Which is great. But that wasn't what kept Troy awake at night.

GETTING TO A 'YES'

AI applications - the good ones, at least - can't be built in isolation from the wider infrastructure and security needs of the organisation. To really embed AI into the business, at the pace and scale that Beazley wanted, the existing systems had to be opened up for the digital equivalent of open-heart surgery. And that created risks.

Of course, Beazley knows all about IT risks. As a leading cyber insurance specialist, they advise clients on how to deal with those risks every day. As a specialty insurer, they know the very worst that can happen, because they're on the hook when it does. And on top of all that, they're a regulated financial services business. Security is paramount to everything they do.

As a result, the first AI use-case took Beazley months to move from concept to deployment. Most of the time was taken up in establishing new technical infrastructure, evaluating security requirements and assessing the right hosting architecture, all based on existing approaches for building traditional software applications.

'This was a scary moment for us,' Troy says. And he should know: as well as being COO, he also has responsibility for the Chief Information Security Officer. 'There was a real possibility that the only way to deliver on our AI commitments would be to create an unacceptable level of security risks and exceptions - which is something we would never even entertain.'

To scale at pace - safely - Beazley needed to create an entirely new technical infrastructure. One that codified best practice security and governance into intuitive, one-click methods that made sure that the easiest and fastest way to deploy an AI application was also the most secure way. Here, the approach of creating a library of reusable AI components paid dividends again: once a module had been shown to be safe, it could be rolled out again and again for different use-cases without creating additional risk.

'There's a natural tension between security technology, and a group of AI engineers and data scientists who just want to run as fast as they can and do amazing things,' says Troy. 'But we've leveraged Faculty's expertise to help us think through how we can remain secure and protect ourselves, but also take advantage of generative AI. So although it's a challenge, it's actually been a healthy tension between the two.' But it's one thing to engineer technology and infrastructure to be safe. It's not going to matter, unless you can convince the executives who are ultimately accountable, to sign it off. Here again, Beazley's integrated approach avoided some common pitfalls. In many organisations, Risk and Governance functions sit outside of the core AI program team. At best, this slows things down; at worst, it creates an adversarial dynamic where the delivery teams feel thwarted in trying to get stuff done, and the Risk teams feel they have to rein in their gung-ho counterparts. At Beazley, Troy avoided this trap by creating a cross-functional group of leaders from the company's Executive Committee. It brought together the people responsible for technology, modernisation, operations, risk and security, chaired by Troy, and put them at the very heart of the AI program. Their remit was clear: they were part of the team getting AI rolled out, not outsiders whose job was to govern it from a safe distance. Each member of the group was empowered to say 'no' until they were satisfied that their area of risk was properly managed, but ultimately their mandate was 'getting to a yes'. Processes, procedures and technology governance were also modernised, to allow them to appropriately address the iterative, uncertain and flexible nature of AI projects in comparison to traditional software development.

As a result of this transformation, Beazley can now take a new AI use-case from concept to production in a matter of weeks - and be confident that it's been implemented safely and securely end-to-end. But can they be equally confident that people are using it?

> "This isn't about cost-saving for us, it's about enabling us to scale our business faster and more effectively - allowing our specialist teams to focus on the work they are best at and enjoy the most."

— Troy Dehmann, COO, Beazley

ENRICHING DAY-TO-DAY WORK

All of this technological innovation would mean nothing if it wasn't genuinely supporting Beazley staff to improve the way they do their jobs. Ensuring that happens takes more than just building some fancy tools and rolling them out to users. In AI, 'build it and they will come' never works. Users always need to be part of the solution, or else the company faces a new variant of 'pilotitis'.

To solve this, the AI programme at Beazley has a dedicated Business Engagement team that works daily with sponsors, testers and users across all departments to identify and design new AI opportunities, and to support the adoption and usage of in-flight ones. This team has been a critical part of making sure that the company builds the right things with AI - not just strategically, but also in terms of how these things integrate with user workflows.

One example of this, in the Claims department, involves the review and processing of the long legal documents that can form part of an insurance claim. The AI model that analyses the documents works well, but it isn't perfect (no model is), and that threatened to cause major issues with trust and adoption. If Claims managers had to review the whole document to find out if the model had come up with the right answers, it wouldn't be worth them using the system at all. So the Business Engagement team worked hand-in-hand with the Claims team to design an approach that instantly highlights where in the document the model has sourced each answer from. This is done in a way that integrates directly into the workflow, and was extensively tested with users to get it right. And because all the apps have a common interface, an employee who's used to using one of them can quickly and intuitively pick up another.

As a result of this and numerous examples like it, users feel listened to. The technology works for them, because it's been built to fit with how they get stuff done. 'A lot of times, when you have new technology or sparkly things, people sort of engage, and then they disengage,' says Troy. 'But we haven't had people disengage from it. We've had a hunger for more, people wanting to understand it better.'

Since Troy asked for four operational use-cases, hundreds of people across Beazley have started using dozens of interconnected AI applications in their day-to-day work. Hundreds more employees will be joining them by the end of 2024.

AI supports Claims Managers to review and analyse risk in complex legal documents. It's used by the Operations team to automatically filter and triage incoming submissions from brokers in real-time, and it helps underwriters quickly understand coverage requirements based on myriad unstructured data inputs. Troy estimates that AI has created millions of dollars' worth of capacity for Beazley as a business, but says that this is just part of the value.

'This isn't about cost-saving for us,' says Troy. 'It's about enabling us to scale our business faster and more effectively - allowing our specialist teams to focus on the work they are best at and enjoy the most. It's about enriching our workforce, the day to day work that our people are doing. It's supercharging them, and making us better at what we want to do.' And he's confident that there is much more to come.

'Ultimately, AI is not just changing how we deliver our offering to our customers efficiently, but also making us fundamentally reimagine what the future of those offerings could be, to help us deliver the best service in the market.' With the work Beazley has done to rewire their own business around AI, they're perfectly poised to succeed.

THE LESSON IN SUMMARY

Building AI widgets is easy. Rewiring your organisation is not.

AI party tricks, like meeting summaries and document search, can seem exciting. But they tend to be limited to the periphery of what matters. They are not going to change the course of your organisation.

◆

Real impact comes when AI is set up to push forward some element of your overall business strategy and optimise the core processes that define your organisation. If you can't articulate clearly how your AI investment is going to contribute towards one of the top level KPIs of the business, it is unlikely to keep people's attention long enough to make a difference.

◆

Building AI into the core of a business requires much more than technology. To achieve your objectives you need to consider the wider question of how change is driven in your organisation, and set things up accordingly. You also need to calibrate expectations around serious change, rather than technology quick fixes.

◆

Cross-cutting executive sponsorship is important. AI solutions need to be pulled by business functions and the users working inside them at least as much as they are pushed by the technology organisation.

◆

At the working level, well-run AI programmes keep a number of elements in sync. They require focus on users from the start to the finish. They require governance to be carefully set up and navigated. They require infrastructure foundations to be laid, business cases to be robust and benefits measured. And they require the development and maintenance of technology.

06

LESSON SIX

It's data SCIENCE, not DATA science.

Data can tell you many stories, but only some of them are true. Cera took a scientific approach to interrogating the data generated by their groundbreaking care operation, to find out if it could help keep their patients out of hospital. Getting the right answer was literally a matter of life and death.

Right now, Mary should be in hospital. The 76 year-old mother of three sits at the kitchen table of her cottage in Berkshire and tabs through photographs on her iPad. She shows off pictures of her grand-daughter, Isla, who's just had her first dance recital this past weekend. The kettle boils as Sandra, Mary's caregiver, bustles around, making tea and buttering toast for Mary's breakfast. Zoe Ball chatters away on Radio 2 in the background.

Mary really shouldn't be here.

Sandra comes in each morning to help Mary bathe and dress, making sure she's comfortable for the day. Having Sandra there every day creates a calm and reassuring routine for Mary, who's needed Sandra's care since she had a fall six months ago.

Sandra tidies the flat, and lays out Mary's medication. At 22, she's closer in age to Mary's granddaughter, but the two women have an easy rapport, like old friends. Sandra knows all the latest news on Mary's family, and Mary fishes shamelessly for gossip on Sandra's love life.

'Are you looking at your dating apps?' she asks, as Sandra pauses making breakfast to tap something into her device. Mary sighs. 'Young people, always on their phones.'

Mary doesn't realise it, but the phone in Sandra's hand is the reason she's enjoying breakfast in her kitchen this morning, and not hooked up to an IV in a hospital bed. When Sandra's on her phone she isn't swiping right for her next date; she's inputting details about Mary's care and condition (Mary knows that perfectly well; she just likes teasing Sandra). Sandra records what food Mary's eaten, if she's drinking enough fluids, her mood and the level of social interaction she's getting. Every observation is a data point that gets fed back to Sandra's employer – Cera – who uses it to support Sandra in the care she delivers. If there are any problems, Sandra or Mary can contact the local branch for help, and the staff nurse can see Mary's full record and provide advice based on the information.

Neither Mary nor Sandra has had to call the nurse today. But without those data points, and the science that they've informed, Mary would almost certainly be in hospital right now.

PREDICTING WHERE THINGS MIGHT GO WRONG

Each year, over two million people request support from their local authority for care in their homes. With over £28 billion of public money spent, and nearly as much privately, the care sector plays a huge part in the economy, and in the health of the nation. And demand is only going to rise. As people live longer, they develop more complex health needs and require more care for longer. By keeping them in their homes, the care sector frees up vital NHS capacity, supports the wider economy (by allowing family members who might otherwise have become full-time carers to continue working), and lets its patients live longer, more fulfilling lives.

Many UK care providers are still analogue, heavily reliant on pen and paper. Cera is different. The company, launched in 2016, has already become Europe's largest provider of digital-first home healthcare. Every month, its cohort of almost 10,000 professional carers make over two million visits to its patients in their homes - equivalent in volume to all NHS A&E departments nationwide. But what's even more important, as Mary's case illustrates, is the things that don't happen. Cera use the data they collect to make sure their patients get the right care when they need it. But they've also found novel ways to use the information so that they can predict where things might go wrong, and use that insight to prevent it happening.

In a parallel universe, Mary was admitted to hospital this morning with a urinary tract infection. She'll be there for several days, occupying a scarce hospital bed that costs the NHS as much as £600 per day, with all the stresses and indignities that being in hospital entails. In that world, she didn't even realise anything was wrong until she woke up today with a raging fever.

But in our universe, a week ago Sandra noticed subtle changes in Mary's appetite, sleepiness, and trips to the toilet. After she fed that into the app, Cera's algorithms spotted that Mary was at risk of developing an infection. Since then, Sandra has notified the GP and pharmacist, who have issued antibiotics for Mary, so that this morning she's sitting in her kitchen with a steaming mug of tea, telling Sandra all about her granddaughter.

How did Mary get from there to here? If you take the long view, it all starts with the philosophy of science.

FOCUSING ON THE THEORY, NOT THE DATA

You often hear leaders talk about how they want to be data-led. Data-led decision-making is better decision-making, we're told. Smart, not dumb. Well-informed, the world as it is, not prone to wishful thinking or corporate fads. But despite the fact that businesses and organisations produce 50 times more data now than in 2010, two-thirds of executives report that decision making is getting harder, not easier.

For anyone with even a cursory understanding of the philosophy of science, that shouldn't come as a surprise. The idea of data-led decision-making has been busted for almost a century.

It had a good run. Most historians would credit the idea to Francis Bacon, the 17th century philosopher who laid the foundations of the Enlightenment. As well as being a pioneer in the field of frozen food (he allegedly attempted to preserve a chicken by stuffing it with snow), he gave science the idea of induction: the principle that to understand the world, you must first observe the 'particulars' – data points – and then draw conclusions that fit the facts. For the next three hundred years, extrapolations from observed data were considered the state-of-the-art way to understand the world.

The problem with this approach was that the data might support multiple conclusions, not all of which are true. Being data-led doesn't stop you being led astray. A famous example of this came in the second world war, when the US Navy wanted to find out which parts of a bomber should get extra armour protection to reduce the risk of it being shot down. They examined the data, in this case the distribution of bullet holes in returning planes, and concluded that the places with the most damage should get the most reinforcement. It took a brilliant statistician called Abraham Wald to point out the flaw in their theory. They were looking at planes which had made it back, but the reason the aircraft had survived was because the places they'd taken damage weren't critical. It was the places they hadn't been hit that needed reinforcement. A perfectly reasonable, data-derived theory had been shown to be dead wrong.

The necessary corrective to Francis Bacon's method arrived during the 20th century, when the Austrian-British philosopher Karl Popper established a revolutionary scientific approach labelled fallibalism. Rather than derive theories based on what the data suggests, scientists should form hypotheses that are capable of being proven wrong. No theory can ever be proven to be true, but if it can't be disproven then it at least stands a chance of being right. Or to misquote Sherlock Holmes, 'When you have eliminated the falsifiable, whatever remains, however improbable, might be the truth.'

It's a crucial shift of emphasis. Suddenly, the central focus is the theory, not the data, which means that predictions made by the theory can be wildly different from simple extrapolations from the data. And to this day, that approach is the best method we have for understanding the world.

But to quote another famous Karl (Marx), who also had no time for the status quo, 'For centuries philosophers have sought to interpret the world. The point is to change it.'

Ben Maruthappu is definitely someone who wants to change the world.

CARE - BUT DIFFERENT

Dr Ben Maruthappu, to give him his full title, needs no lessons in the philosophy of science. He graduated from Cambridge with a triple first in medicine, followed by post-graduate degrees from both the University of Oxford and Harvard University. Ben worked as a doctor in A&E and public health, and was hired by the then NHS Chief Executive, Sir Simon Stevens (coincidentally, another Faculty customer - see Chapter 2) to advise on technology and innovation in healthcare. Life seemed pretty good.

Then, in 2016, Ben's mother fell and fractured her back. Suddenly, Ben was plunged into the reality of trying to organise care for her. 'It took weeks even to get it started,' he remembers. 'I was calling care agencies, but they weren't even picking up. When I finally got somewhere, it was like a revolving door. Different carers every time, I didn't even know what their name was or how the care was going.'

It struck him as absurd that technology let him order groceries, clothes, books - pretty much anything - direct to his doorstep, yet when it came to something infinitely more important like organising care for his mother, he couldn't even find out the name of the carer or what time they would show up. When he probed a bit deeper, he understood why. 'All these companies were using whiteboards, pen and paper to manage their operations. People were overwhelmed with paperwork and administrative processes.

'And I thought: technology can help. Technology can help us build a better model of care that takes away workload from the frontline, and instead people can focus on what they're motivated to do, which is delivering care.'

“Technology can help us build a better model of care that takes away workload from the frontline, and instead people can focus on what they’re motivated to do, which is delivering care.”

— Dr Ben Maruthappu, CEO, Cera

Ben called his company Cera - an anagram of 'care' - that captured his determination to reorder the sector's traditional way of doing things. From the beginning, Cera has invested in technology and embedded it in all their core business processes, from logging notes at patient visits and reminding carers what tasks they need to perform and when; to managing branch operations, staff scheduling and recruitment; through to giving families visibility of the care their loved ones receive. This digital-first approach has obvious benefits for the services they deliver, and it's also enabled Cera to scale their operations rapidly as the company has grown. Which is just as well: it turns out plenty of other people want the kind of care for their families that Ben wanted for his mother. In just five years, the company's revenues have ballooned 150-fold.

Cera's technological prowess has created huge volumes of data, with over 1 billion new data points gathered each week. Remarkably, the simple fact of having these digital records has been shown to improve their patients' health. In a piece of research conducted early on, Faculty found that care system users who had digital health records had better outcomes than those who didn't, simply because their carers had better access to all the relevant information.

Ben knew that had only scratched the surface of the value his data could provide. But he also knew that this value wouldn't come by plotting it on dashboards and trying to spot patterns within it. To unlock real, actionable insights, he would need to do some good science.

'I was a firm believer from back in 2016 that data would help us build algorithms that could predict if people were going to become unwell,' says Ben. 'But it was only in 2020 that this dataset finally reached critical mass to analyse it and use it to improve our services.'

The question Ben most urgently wanted answered was how to prevent his patients unexpectedly ending up in hospital. It wasn't just his personal experience with his mother informing this. When a person being cared for is suddenly hospitalised, it's obviously traumatic for them and their family. It's also distressing for the carer, who will often blame themselves for having failed to prevent it. And rather than the few hundred pounds it costs to look after someone in their own home, a stay in hospital costs the NHS thousands of pounds of scarce resource.

Ben and the Cera team distilled this into a set of three nested questions they wanted to look at with Faculty. First, what caused the kind of deterioration in a patient's health that would necessitate a sudden trip to hospital? Second, could those factors be captured so that the hospitalisation could be predicted in advance? And finally, most importantly: if staff know what's going to cause a hospitalisation and they know when it's likely to happen, is there anything they can do to prevent it?

IDENTIFYING THE CAUSES OF HOSPITAL VISITS

Cera and Faculty put together a cross-functional project team connecting Faculty's analysts with Cera's own data science team, their branch managers, operations leads, and clinicians. 'It was incredibly hard,' says Tessa Farrington, who managed the project for Faculty. 'They had very high standards, really fast-paced with lots of ideas. Very forward-thinking, clever people who appreciated the potential of the technology. They stretched us and we stretched them.'

'It was such a non-traditional consulting job,' adds Hugh Neylan, the head of Faculty's Healthcare business unit. 'It was a genuine combination of their expertise, knowing their business; and our expertise, knowing data and data science.'

Together, the team started to gather and test hypotheses about what caused the hospital visits and how they could be predicted. This was very different from simply letting the AI find patterns in the data. Instead, they interrogated it with the rigour of research scientists. 'We made a massive spreadsheet of all the types of data Cera held, with operational factors on one axis and clinical factors on the other,' Tessa explains. The operational side included elements like where the carers were based, how they were allocated, their experience and the hours they worked. The clinical part was more classically medical data points like sleep, toileting habits, speech and alertness, blood pressure and so forth.

Initially, the team assumed that the answers they were looking for would reside in the clinical data. For example, someone with a urinary tract infection (UTI) is four times more likely to be hospitalised than someone who doesn't have one, so it would make sense that bladder function would be a good leading indicator. But in fact, the team found little correlation. It turned out that people felt uncomfortable answering the question, and so weren't always completely forthcoming. The hypothesis, to put it in Karl Popper's terms, had been falsified.

Other clinical information did have predictive power, but not the sort that could be used to intervene in enough time. Being in pain, for example, is highly predictive of an imminent trip to hospital, but by the time pain is reported it's usually too late to do anything to prevent it.

Conversely, the exact location of a patient in their home when the carer arrived might seem, intuitively, to be irrelevant. And if you look at the whole dataset, there's no particular correlation. But when the team probed the data - testing their hypothesis - they found that in certain circumstances the service user's location had good predictive power. If the person was still in bed unusually late, it was easy to assume they were just tired; but in fact drowsiness is a common symptom of a UTI. So it turned out drowsiness, in combination with other factors, might predict hospitalisations.

This, too, had particular resonance for Ben. 'When I was practising in A&E, there was an elderly gentleman who used to come in every two to three weeks with a urinary infection. And in one instance, he came in a few days too late, and what was really sad was that his infection had spread. We tried to give him antibiotics intravenously, but it didn't work and he sadly passed away. That was very tragic, but it was also so eye-opening, because I realised in that moment that this person could still be alive if he'd received better care in the home, and someone had got to him with antibiotics more quickly.'

Crucially, the insight gleaned about drowsiness was applicable in a usable timeframe, presenting as a symptom several days ahead of the likely hospitalisation. 'What would happen in a typical setting,' says Ben, 'is that a carer visits and notices a patient's tired, but thinks this could be normal or just fatigue. They visit again and again, then three or four days later the patient's unconscious and needs an ambulance. Whereas with the algorithm, we can spot that pattern on the day of the first visit. It flags to our operational team that there's a high risk that needs looking into, we contact the GP, and the carer can pick up antibiotics on the same day, allowing the patient's health to start to improve.'

Another discovery during the investigations phase was almost more surprising. While clinical measurements provided some key insights, the operational data turned out to be equally useful. 'Operations are linked to outcomes,' says Tessa, succinctly, 'and we need to think holistically about it.' One of the hypotheses that they tested was that the consistency of a patient's relationship with their carer would keep them out of hospital. What they found was that service users who saw the same carer regularly had a 30% lower rate of hospitalisation than those who saw many different carers.

This had major ramifications for how Cera used their workforce. They put a new emphasis on recruitment and retention, digging into the data yet again to find out what would encourage staff to stay. Unsurprisingly, the rate of utilisation - ie how much time carers were spending with patients - turned out to be the biggest factor in keeping the job attractive. People who become professional carers want to spend their time caring. So Cera implemented advanced scheduling and routing software to help their carers spend less time on the road, and more time seeing patients.

Once the prediction algorithm was ready, Faculty helped embed it in Cera's workflow. Each day, nursing staff at the branch are provided with the model's outputs in the form of an ordered list of those service users most at risk of hospitalisation. A registered nurse reviews the list, consults the notes, and speaks to the carer, the patient and their family if necessary.

Based on their findings, the nursing staff and the carer arrange quick, low-cost interventions such as GP telephone consultations, district nurse visits or pharmacist medication reviews. For something like a UTI, a timely course of cheap antibiotics can make the difference between a passing inconvenience and a long stay in hospital. Or, as with Ben's former patient, even more tragic outcomes.

But how do they know it's working? How do you measure impact when success is defined by things that don't happen? Once again, science holds the answers.

80%

The technology that Faculty and Cera had built together can predict approximately 80% of potential hospitalisations a week in advance. More than half of these can be prevented by quick, low-cost interventions like a medication review, which means that at least 52% of all hospitalisations can be avoided.

52%

A MODEL THAT CAN OUTPERFORM CLINICIANS

The litmus test of any good science is whether it stands up to scrutiny by other experts. So Faculty carried out a formal analysis of its impact, which was peer-reviewed and published in the academic journal Home Health Care Management & Practice. The results were conclusive. The technology that Faculty and Cera had built together can predict approximately 80% of potential hospitalisations a week in advance. More than half of these can be prevented by quick, low-cost interventions like a medication review, which means that at least 52% of all hospitalisations can be avoided.

A second part of the analysis compared different approaches to spotting these avoidable cases. What it found was extraordinary. The AI model was able to predict hospitalisations 2.6 times more accurately than clinicians when given the same sets of data. What that means is that a carer, without any formal training or qualifications, armed only with their own experience and the app, is more likely to spot a service user who'll need hospitalisation than a highly trained doctor.

It's an exceptionally virtuous circle. The carers feel good because the patients they care for have better outcomes, and because they can see the direct impact of the time they take filling in questions on the app. The service users enjoy better health, and all the physical and psychological benefits that come from staying in their own homes. The trained medical staff who work for Cera can focus their precious time on the most urgent or complex cases, while NHS resources are freed up for other patients. As a former A&E doctor, that's something particularly close to Ben's heart.

'As the population ages and demand for care grows, we are building a more sustainable model of care,' says Ben, 'one rooted in prevention, technology and community. We're freeing up doctors and hospital staff to tend to those most in need. And we're equipping care workers with new, career-boosting skills, building the digitally empowered healthcare workforce of the future.'

MORE TIME FOR CARERS TO FOCUS ON WHAT THEY'RE BEST AT

Since Faculty's project on the UTI prediction, Cera's own formidable data science team have taken the ball and run with it. They modified the algorithm and applied it to falls, the number one reason older people end up in hospital (and, of course, where Cera began). 'Within the first couple of weeks of releasing the updated algorithm,' says Ben, 'falls reduced by 25%.'Meanwhile Cera is expanding at home and abroad, and also looking at licensing some of its technology to other providers so more people can enjoy its benefits.

'Looking back on what we achieved, it's one of the projects I feel proudest of,' says Tessa. Hugh agrees. 'These carers are quote-unquote "unqualified" people, who are often taking care of the most vulnerable members of our society, the people we love most dearly. And this algorithm lets them focus on caring because that's the thing that they're best at, and it actually has an impact on outcomes.'

The final verdict should go to the people most affected. Kenza Maduro, a London-based carer for Cera, says, 'There is a lot to think about in care work. The Cera app takes the work out of it, making life much easier and freeing up my time and headspace for the really important work - focusing on our service users. I am very motivated by Cera's vision, and I can see first-hand the huge positive impact our technology has on the people we care for.'

And on behalf of those who are cared for, Mary Hill, whose father Peter received care from Cera, perhaps puts it best. 'I will never be able to properly express my true gratitude for everything Cera has done and continues to do. Your company has allowed me to concentrate on being a daughter whereas the system has forced me into being a very stressed and fierce campaigner for my father's rights. We urgently need new models like Cera's. We need to change mentalities, embracing technology to make better home care a reality for millions of people as the population ages.'

When data science is done well, that's what can happen.

THE LESSON IN SUMMARY

It’s data SCIENCE, not DATA science.

Data gets a lot of attention. The sheer volume of data being created attracts a lot of attention. And data is of course the fuel of AI. But data alone solves no problem. It’s the science that you do on top of the data that matters the most.

◆

Science is all about building an understanding of the world, and the cause and effect relationships that drive it. This is the foundation of applying AI successfully.

◆

You need to understand the cause and effect relationships inside a system before you intervene on it. This means forming hypotheses and running experiments to establish cause and effect relationships (not correlations). It means applying rigour and honesty when drawing conclusions. And it means prioritising the most simple techniques and parsimonious explanations, rather than being seduced by complicated and shiny technology

◆

You also need to identify the causal pathway by which your interventions into a business process achieve the outcomes you seek. Incomplete or wishful thinking here is a guaranteed route to disappointment. Unless you can lay this out step by step, you are not ready to start building or implementing technology. Faculty we use a methodology we call the ‘Decision Loop’ to map this out systematically.

◆

The importance of good science remains as true as ever even in the LLM era. Those who declare that powerful out-of-the-box models spell the end of the data scientist are wrong. Just as they were wrong the last time the death of the data scientist was called when Auto-ML solutions were de rigueur.

07

NESO

LESSON SEVEN

There's no such thing as complete data.

NESO is the UK company responsible for keeping power flowing to 67 million Britons. It's walking a tightrope between a carbon-intensive past and the clean energy future - but that rope is getting wobblier every day. Locked in its vaults, NESO may have the data it needs to train AI to keep its increasingly complex processes running smoothly. But is the data good enough?

Mist swirls in the floodlights mounted on the side of the long, slow-slung office block in the Berkshire countryside. It settles over the cars parked in the surrounding lot - more than you'd expect at six a.m. on a cold October morning - and wraps the thick forest that hems in this business park. At the gatehouse, the guard's breath steams in the freezing air as he raises the barrier.

Inside, the building is bright and warm, but an autumnal hush seems to have slipped in through the air vents and settled even in the innermost sanctum. This is a high-ceilinged room at the heart of the building - white walls, blonde wood - a room whose very nondescriptness seems designed to force your attention onto the screens that glow with information everywhere you look. There are literally hundreds of them, monitors crammed shoulder to shoulder across the rows of desks that all face towards the front of a room: monitors on stands, monitors on arms, monitors squeezed so tight they crowd each other out. Shirtsleeved workers sit behind them, studying the data and occasionally tapping commands into their keyboards.

Gazing back at them, mounted high on the front wall, is the master display: a two metre-high screen that stretches across the whole width of the room. It's covered in an intricate schematic of intersecting lines and interlocking rectangles, labelled where they cross with cryptic white numbers. As the operators tap their machines, the lines change colour: blue to purple, purple to red. The bar-graphs at the edges rise and fall.

This is the control room of NESO, the National Energy System Operator. The network on the screen is a high-voltage circuit diagram of the whole of the UK turned 90 degrees, with Aberdeen in one corner and Cornwall diagonally opposite. In real time, it shows the energy coursing through the country, from power stations and wind farms and solar arrays, across the grid, and out to the local networks that step down the voltage and feed it into homes and offices.

The control room, with its cathedral ceiling and hushed atmosphere, is a temple to power. But it's also a shrine to data, the millions of bits of information that feed back from the grid's nervous system to inform the decisions that NESO operators make every day to keep the lights on.

If they're going to keep those screens glowing, everything is going to have to change.

KEEPING THE LIGHTS ON IN THE POST-FOSSIL ERA

NESO is probably the most important company in the UK that most people have never heard of. It's certainly the newest. In the control room, banners on the side walls sport the company's new corporate branding - a soothing mauve-and-plum colour scheme - covering up places formerly painted with the blue National Grid logo.

NESO was spun out in October 2024 to be a free-standing entity that manages the UK's energy system. Power plants' giant cooling towers, dense wind farms spinning above the North Sea, steel pylons marching across the landscape: these are all integral parts of NESO's business. But it doesn't own or operate any of them. Instead, its remit is to manage the energy that flows from and through that infrastructure. It does that by sending instructions to different generators to provide just the right amount of power at the right time: turn it on, turn it up, turn it down or turn it off. It's a delicate balancing act that they have to get right every minute of every day. If the amount of energy going into the grid isn't the same as the amount of energy coming out of it, things fall apart very quickly.

NESO's overriding mission, to the exclusion of all else, is to keep the lights on in the UK. It's a job their engineers and engineers have done with quiet efficiency and enviable stability. In 2014, the network boasted 99.99995% reliability, the best in Europe. Nine years later, that last digit had ticked up to a nine. NESO's goal is to add another nine on the end for good measure.

But NESO now has another mission. The government has charged it with devising the roadmap to shift Britain's energy sector fully away from fossil fuels by 2030. That's five years ahead of the US goal, and ten years ahead of the EU's. The New York Times, not given to hyperbole, describes it as 'the most ambitious target of any major industrialised economy.' Other commentators are more pointed and use words like 'fantastical' and 'unachievable'. Fintan Slye, the guitar-playing Irishman who's in the hot seat as NESO's CEO, can see their point. 'We're not saying that the target is achievable with the current energy industry processes and systems and ways of working,' he told the Guardian. 'In fact, it's not achievable in those circumstances. But if you can make the required changes, then it can be delivered.'

EVOLVING TO BECOME 'WORLD CLASS' IN AI

What needs to change? In his interview, Fintan mentioned the planning regime, the regulatory system, the grid connections process... to say nothing of the physical infrastructure itself. The last big buildout of Britain's power grid happened in the 1960s, when fossil fuel plants were largely centralised in the industrial heartlands of the Midlands, close to the coalfields that supplied them. Managing the grid with a limited number of large suppliers, their output predictable and reliable, was never easy; but it was a fairly straightforward proposition. Now things are changing.

The day before NESO came into being, Britain's last coal-fired power-plant, at Ratcliffe-on-Soar in Nottinghamshire, closed. Its icon will vanish from the circuit map on NESO's big screen, marking the end of 140 years of coal-powered electricity generation in the UK. There's no comparable big replacement plant: instead, the capacity will be supplied by a galaxy of different renewable projects of every shape and size. It captures the broader trends that are reshaping the nation's electricity mix: a huge opportunity for decarbonising the economy, but a huge challenge for the company charged with overseeing that transition while making sure the system never skips a beat.

And tying it all together is the software. 'The scale of change that has happened in Great Britain's generation portfolio means the IT systems have not kept up with it,' said Fintan. 'We've shifted from needing to send five instructions an hour to 500 instructions an hour. It's creaking.'

'They've been using systems that were designed 50 years ago,' explains Niko Louvranos, the voluble Greek who's responsible for growing Faculty's energy-related business. 'And although they've been modernising, a lot of it still remains in the old ways because it just works.'

It works at human scale. A combination of old systems and little margin for error mean that a lot of NESO's operations rely on proven manual processes performed by their vastly experienced staff. But if the demands and the complexity are growing by orders of magnitude - and forecast to grow even more - there's no purely human solution.

NESO needs AI. Without it, the grid won't be able to manage all the wind farms, solar panels, batteries and electric vehicles coming online in the next few years. And nobody needs the consequences spelling out if it fails.

The company's executives are ahead of the game and fully signed up to the need for AI. 'NESO's leadership is very ambitious at becoming world class in AI,' says Niko, who's worked firsthand with the management team to help them shape their AI strategy. 'Very few other organisations in the energy sector are as ambitious as these guys are with AI.'

If ambition and commitment were all that was needed, NESO would already be well on the way towards solving its AI requirements. But there's a challenge - one that for many organisations would be insuperable.

It's the data.

BANISHING THE MYTH OF PERFECT DATA

'Data is the oil of the 21st century,' goes the famous saying. While the analogy isn't perfect, it's got a particular poetic resonance in the energy industry, and it captures a core truth. Without data to feed the models, the AI machines won't run. Early AI systems built using hand-crafted rules turned out to be brittle creations that struggled with the complexity of the world they were modelling. It was impossible to codify all the rules for a system to follow. But if you have enough data (and enough compute to process it), then deep learning algorithms can learn those rules for themselves, with far more subtlety and nuance than hard-coded versions.

That development has been at the heart of the AI revolution of the last few years. It's also changed the way we think about data, and elevated this once arcane technical concept to the top table of business discourse. Chief Data or Information officers roam the C-suite, with fat budgets to invest in programmes to collect or manage data. Executed well (see for example Inspired Education's data platform in chapter 3), these sorts of programmes provide the foundations needed to successfully adopt AI into an organisation. But there are traps that need to be navigated, both by organisations with little data, and by those with lots.

NESO definitely falls into the latter category. Managing 20 terawatt hours of electricity a month (which NESO helpfully quantifies as 20 billion washing machine cycles), sending it the length and breadth of the country - and overseas, via interconnectors - generates mind-boggling quantities of data. Other companies profiled in this book, like DRIFT or Cera, had to wait months or years to accumulate enough data before they could train the algorithms they wanted to build. NESO is positively drowning in data. Just one data set, for example, is receiving 90 million updates every day. Multiply this by the hundreds of data sets they keep, and you get a picture of the scale of information they need to crunch.

But although NESO has the volume, there's a catch. AI algorithms want their data in tabular form, racked and stacked in orderly columns and rows. Some of NESO's data fits the bill, but other parts of their sprawling data portfolio can range from PDFs of ancient manuals, to Word documents, to tangled schematic and process diagrams that any algorithm would struggle to digest. Some of the data is held in different software applications of a certain age that don't like to talk to each other. Other data changes as it goes through different business processes, without a clear record of the path it's taken. It's...complicated.

To be clear, this isn't a problem unique to NESO. While digital native companies have the freedom to design their data architecture from scratch and build their institutions around it, any long-established organisation with complex legacy systems is going to have patchy, convoluted, or hard-to-reach areas of theirdata estate. And when those companies face significant operating challenges that might be solved by better use of data - or simply when they hear calls from investors to 'do AI' - many of them will succumb to the temptation to embark on a big data infrastructure exercise, getting everything scrubbed and tidy before unleashing the algorithms on it. Which makes sense. After all: no data, no AI.

But the desire to 'fix' the data before it can be put to good use is ultimately built on a pair of myths: the myth of 'perfect' data, and the myth of 'complete' infrastructure. The truth is, data is never perfect; and data infrastructure is never done. Both exist, like the world they catalogue, in permanent states of imperfection and constant evolution. Trying to pursue some idealised end state isn't just futile, it's costly and time-consuming. The modern corporate mausoleum is filled with the corpses of large-scale data transformations that either took a lot longer than planned, cost a lot more, or failed entirely.

NESO is taking a different approach.

BANISHING THE MYTH OF PERFECT DATA

Faculty's Energy Transition and Environment team is the newest business unit in the company. It started in 2021 with a single hire, and it's been growing ever since. Niko Louvranos, the unit's energetic Business Development Director, gets even more animated than usual when he describes its mission. 'We're the youngest team, taking on the biggest challenges, for the most significant players, in the industry that the world is relying on to get to net zero.' The data scientists, engineers, and consultants who work on the team are young, super smart, and serious about making real change.

When Faculty started working with NESO, the grid company already had a comprehensive digitalisation programme in play: a well-designed, well-run project that any CIO would be proud of. But crucially, rather than focussing all of their energy on getting the infrastructure right, NESO also looked for opportunities to address some of their most urgent challenges using data they had available already, even if it was imperfect. They decided they could build the tools and infrastructure they needed as they worked out the use cases it would serve.

The first project in the pipeline looked at planned outages on the network. These are decisions to turn

The truth is, data is never perfect; and data infrastructure is never done. Both exist, like the world they catalogue, in permanent states of imperfection and constant evolution.

off particular pieces of infrastructure - anything from a whole power plant to an individual transmission line - in order to perform planned maintenance. This might be for something as simple as a bird's nest on a pylon, or as complex as a major refit on a power plant. Some outages are so significant that they can be scheduled up to six years in advance.

One common reason for an outage is to connect a new power source to the grid. With ever more renewable assets and supporting infrastructure being added to meet the UK's net zero targets, these are becoming more and more frequent. The backlog is growing, and if vital maintenance gets delayed too long then suddenly a planned outage can become an unplanned, emergency outage. Which is a whole new world of pain.

Shutting down any part of a network as complex and critical as the national grid is a daunting proposition, and a whole team at NESO - the Network Access Planning Team - is dedicated to evaluating all the factors that have to be taken into account to find the best time, and the best mitigations. If a transmission cable is out, for example, is there enough headroom on the other available cables to safely carry the extra power that will be coming their way? If generating capacity goes offline, what will replace it, and how much more will it cost if that happens at peak times?

To perform this analysis, gaming out all the different scenarios to find the best solution, is a monumental task for the team. 'At the moment, they will do hundreds and hundreds of very heavy duty simulations to work out exactly when looks like the right time for the outage,' explains Katrina Soderquest, the senior Faculty data scientist who worked on the project. 'It takes a lot of time, and a lot of people.' Although there are optimisation tools to help, they rely on trial and error, manual runs and human intuition.

NESO were keen to see how machine learning and AI could help them speed up the process to handle the growing wave of demands. When the team started working with Faculty, both sides knew that the data and supporting infrastructure would be imperfect. But Niko and his team were determined to push the art of the possible. They knew that a specific model, with a specific objective, only needs specific data to achieve its aim. More data is only useful if it directly relates to the objective and adds more signal to the model, rather than noise. In other words: you don't need all the data, you need the right data.

And you don't start with the data you have and then think about what to do with it. You start with the problem you want to solve, then think about the technology that will help address it, and then what data you need to feed into the technology. Data is the essential enabler, but it comes last in the logic, not first.

NESO were wise to that. They resisted the temptation to try to 'fix' their data first before starting; they didn't have time to chase that endlessly receding horizon.

'THE OUTER LIMITS OF WHAT'S ACHIEVABLE'

Thanks to that leap of faith, the pilot project that Faculty worked on proved that AI could manage with the data NESO had, imperfections and all, and significantly speed up the outage planning cycle.
'It's about doing the easier stuff quickly,' says Katrina. 'You're never going to replace the experts in those fields, but this project showed we can make planning more robust and quicker at the same time.' There's also a cost implication. NESO spends billions of pounds a year buying energy to balance its system. Even a small increase in efficiency translates to huge savings.

Other projects, run along similar lines, have demonstrated AI's promise in the tricky tasks of balancing voltage on the grid (a physics lesson in itself), and in gas transmission network planning. More are in the works. Because of the complexity of the grid, and NESO's management's desire to make sure that these changes happen in a coherent way, the new algorithms haven't been implemented operationally yet. 'They want to link the different AI solutions to their wider processes,' says Niko. 'They don't want to just fix things in isolation. They want to do it holistically.' But the course has been set and the direction is clear.

In his Guardian interview, Fintan Slye described the challenge of getting to a clean power system as 'at the outer limits of what's achievable. But,' he went on, 'if you're prepared to do things differently, and to take difficult decisions early on, then yes, absolutely it is doable.'

So long as you don't wait for perfect data.

THE LESSON IN SUMMARY

There's no such thing as complete data.

Data is essential to modern AI. No data, no AI.

◆

But AI is a precision game. A specific model, with a specific objective, will need specific data to achieve that. More data is only useful if it directly relates to the objective at hand and adds more signal to the model.

◆

The right starting point is always to be clear about the problem you're trying to solve. Then what technology can help you solve that problem. Then what data you need to power the technology. In that order. Data comes last, not first.

◆

It is easy to be held back by getting this thinking the wrong way round. By starting with the data; trying to bring it all into one place, to organise it in the perfect data lake, to 'fix' it in some way. In fact, data is never perfect. And data infrastructure is never finished. Both exist in permanent states of imperfection, and constant evolution.

◆

Waiting until you have complete data to deploy AI will mean waiting forever. Instead, you need to focus on making sure you have the specific set of data you need for specific applications you want to build.

08

NOVARTIS

LESSON EIGHT

Build in increments that are individually valuable & collectively transformative.

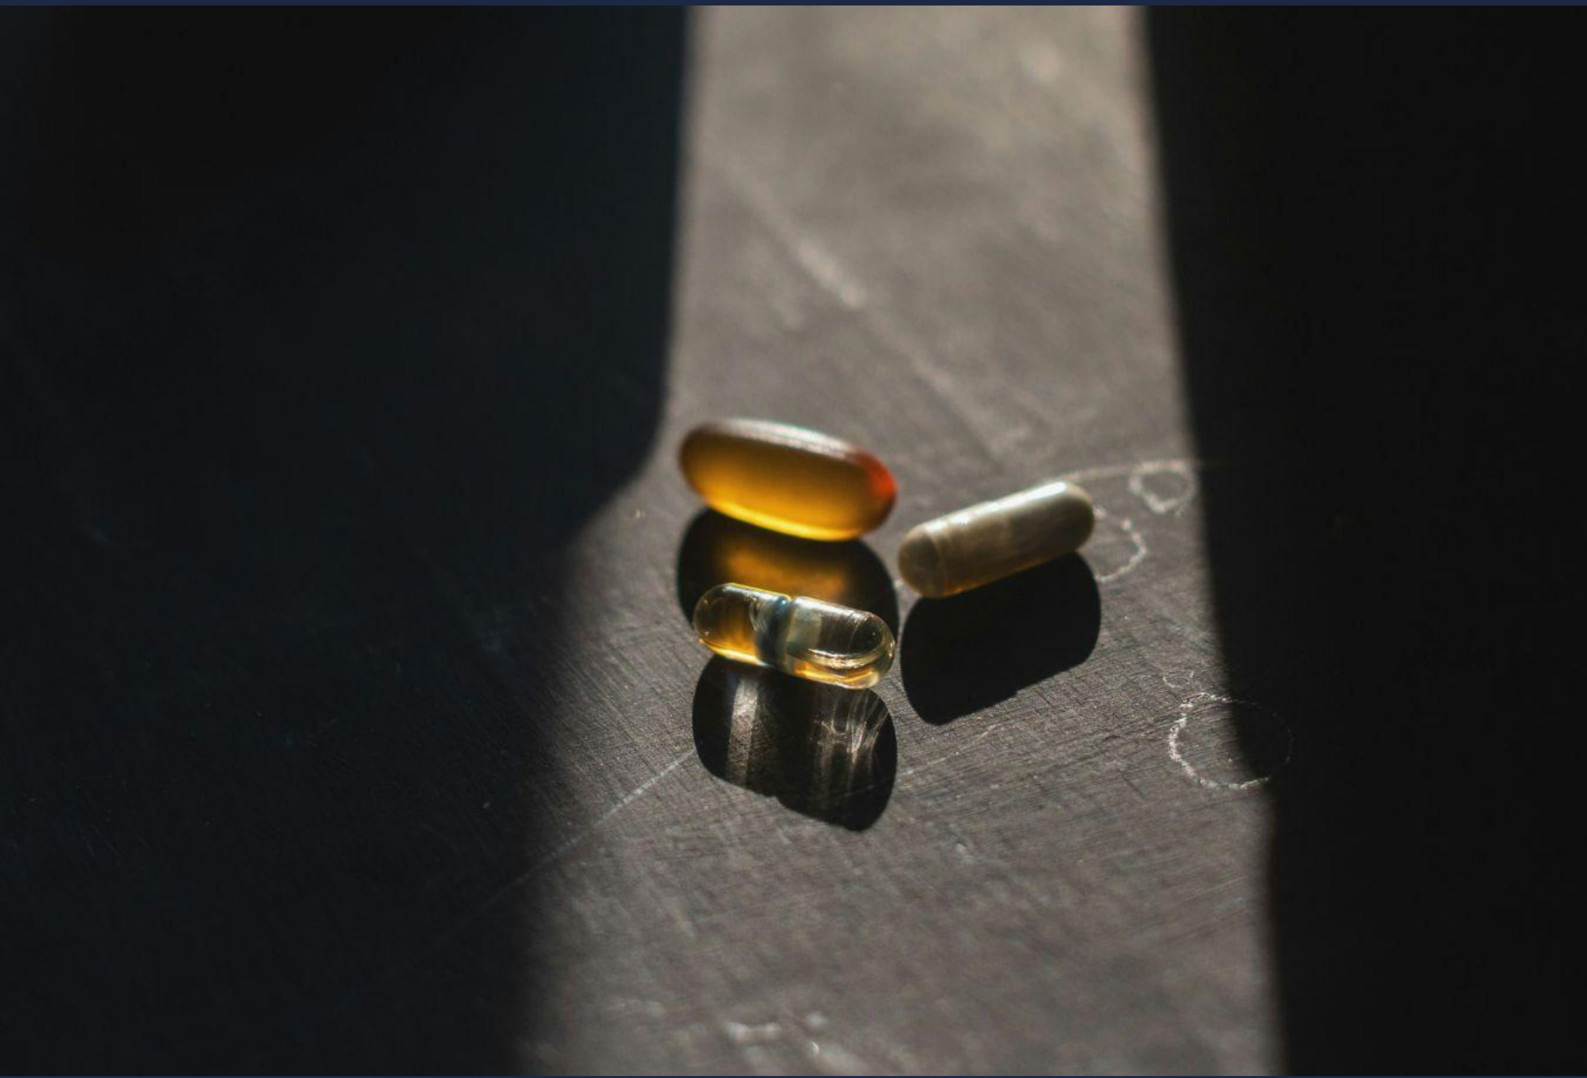

AI's benefits don't come much bigger than the potential to deliver new medicines that let people live longer, healthier lives. But bringing a drug to market involves so much more than just finding the right compound. Novartis are applying AI along the whole continuum of their work, to improve decision-making and get medicines to patients faster. Each step makes a difference. Collectively, they're transformative.

Look at the pill in the palm of your hand. It doesn't seem like much. In colour, texture and (possibly) taste, it's like a tiny nub of chalk - and it looks about as complex. Compared to the smartphone in your pocket, with its nano-scale microcircuitry and blazing screen and global connectivity, the pill's just an inert lump. The phone has AI; the pill's just… artificial. You don't give it any more thought than you do the glass of water you wash it down with as you pop it in your mouth.

You're missing what's really happening.

The medication that you just took is the result of a decade-long process that began with a molecule in a research lab and ended just now when it hit your bloodstream. Just as much as your phone, the pill is a miracle of world-class minds, cutting edge research, advanced manufacturing, global supply chains and an obsessive commitment to quality and safety.

Most of all, that pill is the sum of thousands and thousands of individual decisions that have guided its journey from the moment a scientist conceived the idea, through unimaginable layers of tests, trials, protocols and approvals, to the moment you popped it out of its blister pack.

Now the company that made that pill, Novartis,wants to use AI to enhance every stage of the journey. Same rigorous science, same uncompromising focus on quality, just better decision-making. And faster.

NEW ARTS

The name Novartis was coined in 1996 when the Swiss multinational pharmaceuticals company took on its current form. The name came from the Latin novae artes, meaning 'new skills'. It's an apt name for a company dedicated to producing truly innovative medicines, and it's taken on an extra dimension now as Novartis looks to develop even newer skills in AI.

Overseeing the clinical development process is Shreeram Aradhye, Novartis's President, Development and Chief Medical Officer. A trained physician, he is absolutely the sort of doctor you would want treating you if you were sick, bubbling with good humour and enthusiasm for his subject. 'I'm excited to have this conversation,' he says, when asked to discuss how AI is being embedded in the development of new medicines. 'And what I'm most excited about is not the technical part of AI, it's the human part of the engagement with this technology.'

“...what I’m most excited about is not the technical part of AI, it’s the human part of the engagement with this technology.”

— Shreeram Aradhye, President, Development and Chief Medical Officer, Novartis

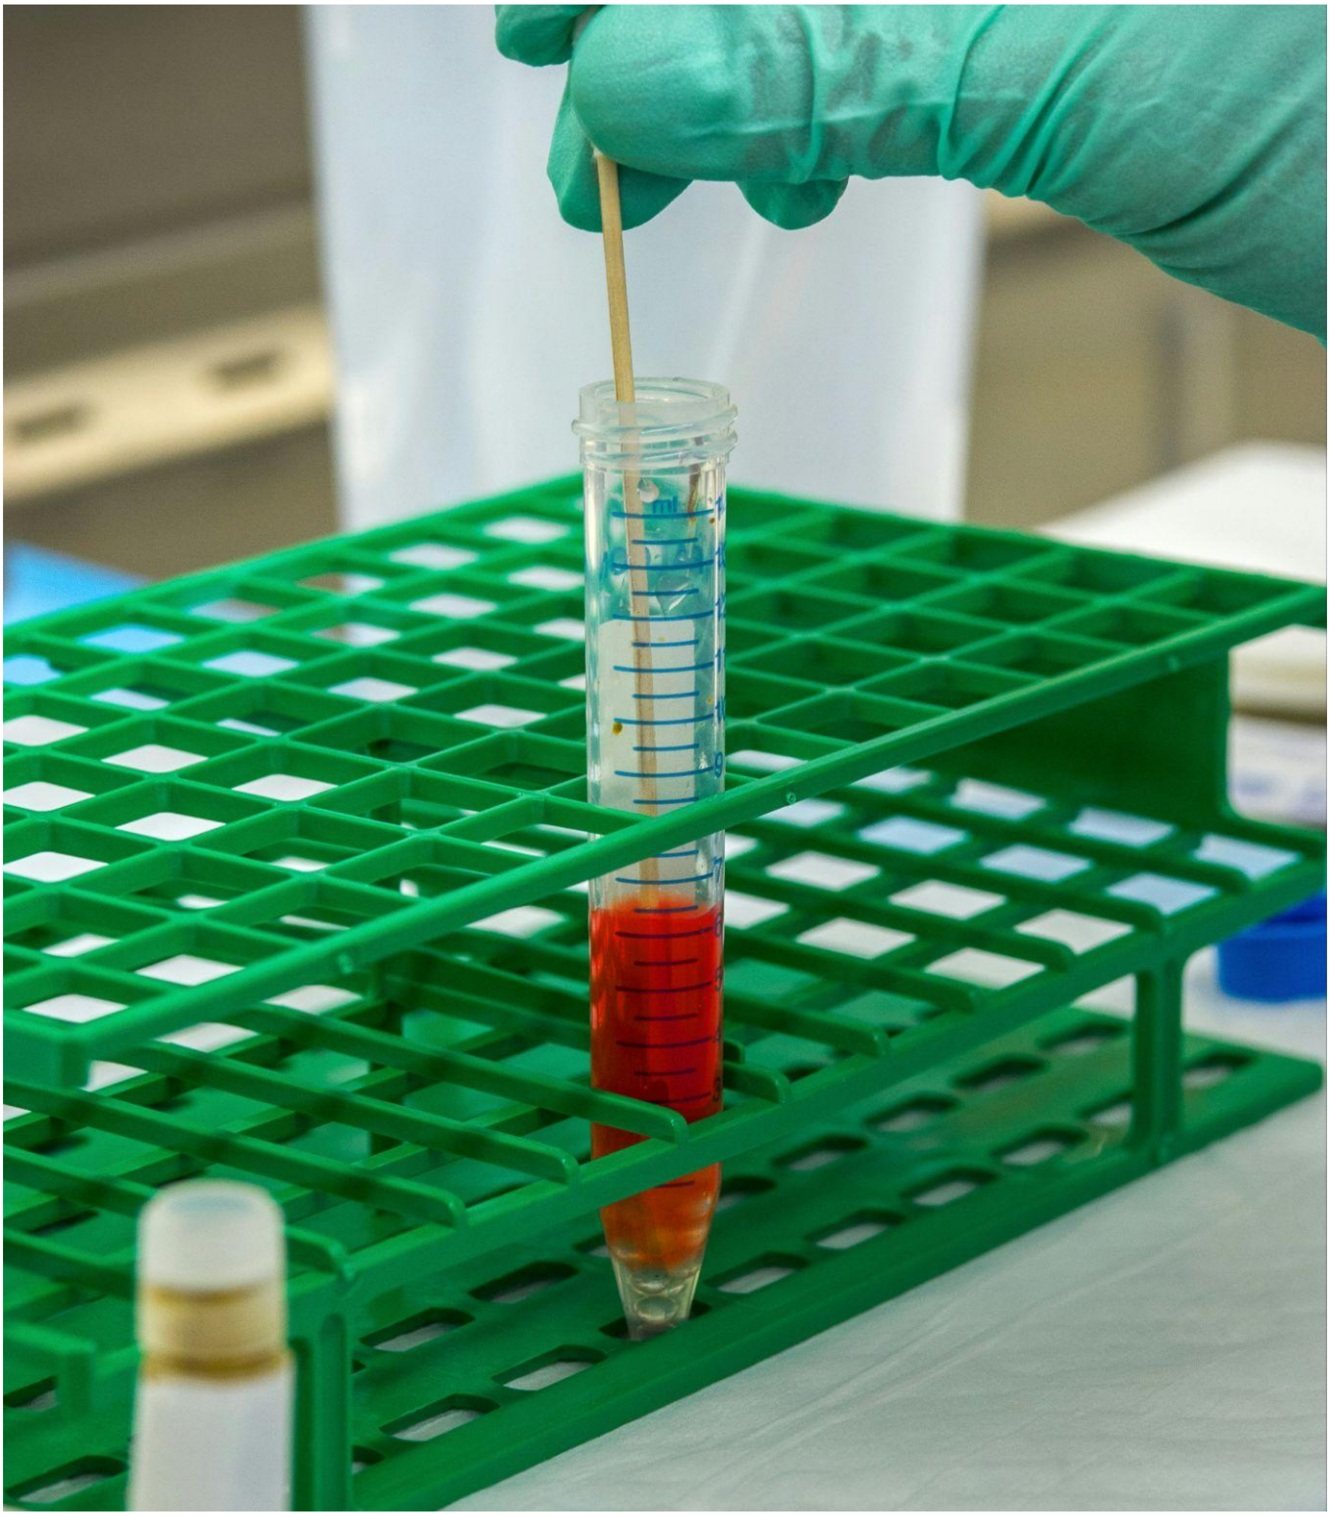

Shreeram describes the company's purpose as 'turning molecules into medicines and getting them to patients.' It's a neat little summary for a process that is one of the most complex on the planet. From the point that a promising compound is identified, it takes on average a decade to get it to market. Development can cost hundreds of millions of dollars, sometimes even billions. And all that money only buys you a one-in-ten chance of success: 90% of drugs that enter development fail.

Novartis has already embraced AI in the earliest stages of that process through its strategic collaboration with Isomorphic Labs, the London company that uses AI for drug discovery. Established by Demis Hassabis, who also founded the pioneer AI research lab DeepMind (see chapter 10), the company proved its bonafides in 2024 when Hassabis won the Nobel Prize for Chemistry for his work on using AI to predict the structures of proteins. The collaboration with Novartis has the potential to help find promising targets for drug development in ways that couldn't be imagined before.

But because of the timescales, the targets Novartis or Isomorphic identify now will take a while to reach patients. Much quicker to be felt will be the ways AI impacts Novartis's productivity in drug development: essential aspects of the process like new trial protocols, regulatory submissions, and managing large patient data sets. The more immediate gains, in other words, aren't at the wild frontiers of drug discovery. They're in improving the processes that get a candidate product from the lab to the market.

That's why Shreeram's excited.

TRANSFORMATIVE INNOVATION

One of Novartis's key corporate tenets is the idea of 'transformative innovation'. The company doesn't aim for small, incremental improvements; they focus on dramatic advances that can significantly improve patients' quality of life by curing, treating, or even preventing diseases in genuinely novel ways.

But how do you transform a system as complicated as drug development? Complex systems are, by their nature, hard to change. It requires delicate judgement to decide how much to bite off at once. You can go big, try to imagine the future and leap towards it in one gigantic bound - but the track record of this kind of big bang approach is, to put it mildly, mixed. Monolithic technology programmes almost never work inside a complex system.

Small changes - the 'easy wins' and 'low-hanging fruit' - have a higher chance of success, at least locally where they're deployed. But they probably won't make a measurable difference to the overall performance of the organisation. A lot of companies' AI investments fall into this category, discrete initiatives with no mechanism to join them up, that only reinforce siloed decision-making.

Shreeram's having none of that. 'We replaced the word "divisions" with "continuum",' he says. 'Commercial organisations all have to work together along this continuum, and all of them can be AI-enabled to be more efficient at the individual level, and at a team level.'

The best way to use AI to change a complex system, as Shreeram recognised, is to break down that continuum into modular components that are valuable in their own terms, but which connect together into something much more transformative. And the key challenge here is to decide what form those individual components or increments should take. Particularly with AI, it's easy to be misled by the outputs that the models create. AI specialises in forecasts, clusters and classifications of data. It summarises documents and generates new ones. All these capabilities are impressive but none of them, in their raw form, are high-value business outputs.

The most successful applications of AI in an organisation tend to be those where the software's outputs are carefully shaped as inputs to high-value decision-making, supported by intelligent technology. So when you're looking at how to break down the business continuum into meaningful constituent parts that can be enhanced by AI, the right increment is the decision. Specifically, the decisions that create (or indeed destroy) the most value. Then you can use the technology to improve the speed, quality and accuracy of the decision-making processes.

It's easy to think of this as an arithmetical process: that the output is the sum of all the decisions made. But the right calculation is actually combinatorial. The decisions don't operate independently. They depend on one another, along the whole continuum, up and down the system. The outputs of one often comprise the inputs to the next, and you can't understand the consequences of one decision without knowing what a set of related decisions were.

When it comes to Novartis, few parts of the business have as many complex, interlinked decisions as the process of testing a new medicine.

BUT CAN YOU PROVE IT?

'A medicine,' says Shreeram, with his knack for capturing complex processes in quotable one-liners, 'is simply a molecule plus evidence.' Making sure that evidence gets gathered through a series of ever-larger and more ambitious trials is one of his core responsibilities. But as these trials can last months or even years, there's plenty of time for things to go awry.

'A medical trial isn't like launching a rocket and letting it go,' Shreeram says. 'It's the equivalent of a ship that has to navigate across the complexity of the ocean.' And there's not much room for error. A molecule might be effective, but if the trial isn't designed and managed well then the data might not show it. So Novartis turned to AI to support the decisions that help them chart a course through that complexity.

'We spent a year systematically looking at how AI can enable our clinical trials,' says Shreeram. It starts with the scientific protocols that underpin any trial. These are a sequence of decisions that shape everything that follows. What primary and secondary goals should we set for the study? What type of experimental design should we use, and how should we collect the data? What should be the eligibility criteria for participants? What dosage, frequency and duration of treatment should we specify? How often should we bring patients in for assessment?

Shreeram describes a tool Novartis developed called Protocol AI. 'It was our effort to say: "Can we enable the person who writes the protocol? Can they be augmented by an AI tool that gives them access to knowledge and information about comparable trials that have been designed for the same medical condition? Can they learn from the previous performance of similar trials to understand the impact of different design features, and what their future implications might be?"'

Once the protocols are established, the hard work of recruiting participants begins. This, too, is a decision-intensive process. 'We have hundreds of sites across multiple countries,' says Shreeram. 'And there used to be a very complicated process where we had to send the request to the countries, find out if it was feasible, and it took weeks to go back and forth.'
In choosing which sites to use, investigators have to consider whether those sites can provide enough participants for the trial, whether they'll be able to meet the protocol requirements, even whether other pharmaceutical companies will be competing for the same patient populations. More and more decisions, each intersecting with the others.

And when the trial is over, the results have to be written up for submission to regulators. Again, Novartis are deploying AI to improve the process.

'We use GenAI to generate first drafts of reports so the medical writers can focus on the actual interpretation of the content to position it in the best and most accurate way possible,' says Shreeram.

"A medical trial isn't like launching a rocket and letting it go... It's the equivalent of a ship that has to navigate across the complexity of the ocean."

— Shreeram Aradhye, President, Development and Chief Medical Officer, Novartis

INNOVATE DIFFERENTLY

Each of these is a compelling use-case, and they're already starting to show results. Trials that used AI in the design process have been more likely to finish ahead of schedule. Sites that AI identifies as being suitable for a particular trial have strong potential to recruit patients faster than ones that weren't selected by the AI, and produce the diversity of patient population that the experimental protocol demands. Those sites also get approved faster. 'What used to take weeks now... doesn't,' Shreeram marvels.

Incrementally, each of these is making a difference. But to fully achieve the kind of transformative innovation that Novartis aspires to, the company is working with Faculty to tie it all together with an Intelligent Decision System called Frontier. This system uses a 'computational twin', a sophisticated digital simulation that connects data sources, operational processes and machine learning models together in an interactive, virtual replica of the clinical trial process. In this controlled environment, Novartis can experiment with different ways of linking up its various AI initiatives to make sure they are all working in synergy to deliver the greatest benefits. This means systematically deconstructing the decision-making process to see where the key decision points are, what data will be required to inform them, and where future AI investments will have most impact in enabling those decisions. Ultimately, it will let Novartis connect decisions across different functions (such as clinical, operational, supply, strategy, regulatory) and at different levels (for example by patient, site, cohort, medical condition, program, or the whole portfolio of trials). It will provide a scalable framework that allows future investments in data and AI to connect seamlessly to what's already in place.

'The focus now will be on ensuring that we scale the few things that we know definitely work, but not necessarily spend as much energy in coming up with hundreds of more new things,' says Shreeram, reaching for a botanical metaphor. 'We're shifting from a meadow of wildflowers of AI innovation, to more of a curated garden.'

John Gibson, Faculty's Chief Commercial Officer, has a different analogy. 'We think of it as a keyhole surgery approach. No bottom-up, multi-year data transformations. No baskets of AI use cases looking for a business user. It's a pragmatic approach that drives AI-enabled change in a sustainable way.'
'The goal is not to fail fast,' says Shreeram. 'The goal is to succeed as quickly as we can.' When an organisation augments decisions with well scoped and safe AI technology, it's able to raise the average distribution of decision-making quality across complicated processes. 'We can shift the performance of multiple individuals closer to the performance of the best, and I'm convinced that's where the real value is going to come.'

And in all this, Shreeram never loses sight of the bottom line. 'I'm not focusing on efficiency. I'm talking about value, and I always say the value is not just the amount of money saved. It's value to patients, too, so that they can benefit from earlier treatment and faster access to new medicines.'

Novartis have already started accelerating that vital outcome with AI-enabled processes. Connected together in the service of human decision-makers, they will transform it.

THE LESSON IN SUMMARY

Build in increments that are individually valuable and collectively transformative.

Unambitious AI programmes usually underwhelm. Don't just focus on low-hanging fruit. Or on single use cases. That's not how you make a measurable difference to the overall performance of an organisation.

◆

Over-ambitious AI programmes usually underdeliver. It is very risky to try and imagine the whole future at once and seek a big bang. Monolithic technology programmes almost never work.

◆

The right balance is to think big, but build forwards in modular steps. Each individual module should be quick to implement and valuable on its own terms. But modules should also be designed so that when connected together, they transform a whole process end to end.

◆

The operational decisions that determine how a business process runs make good targets for individual modules. Improvements to the speed, quality and execution of decision-making are one of the most reliable ways in which AI can improve the overall performance of a business process.

◆

Connecting these individual decisions together allows them to break out of organisational silos and better account for upstream and downstream interactions. This shifts focus from what is best for each local part of the process to what is best overall.

09

NATIONAL CRIME AGENCY

LESSON NINE

Business strategy trumps AI strategy.

Good strategy is built on identifying what's most important. When a huge trove of confidential intelligence reached the National Crime Agency, it threatened to swamp their ability to analyse it using traditional methods. A team of dedicated professionals had to find a way to process the information faster than ever before, in order to track down the worst offenders and stop them.

Cerys Evans looks the opposite of dangerous. Bright, self-deprecating, and ever so slightly geeky, she exudes warmth and positive energy. She lights up when talking about her dog. If you sat opposite her on a train, you might guess she worked in publishing, or maybe a trendy branch of academia (she is, in fact, doing a PhD). Only - if you were paranoid - you might notice her spectacles, outsize gold-rimmed lenses like a pair of magnifying glasses. Almost as if she was watching you. If you're a certain type of criminal, she is watching you.

Cerys works for the National Crime Agency (NCA), the UK police organisation charged with leading the fight against serious and organised crime, and tackling the UK's most dangerous criminals. Established in 2013 and quickly dubbed 'Britain's FBI', the NCA works at the leading edge of law enforcement to build the best possible intelligence picture of criminal threats, and develop innovative capabilities for other partners to use.

Although often underestimated, serious and organised crime is one of the most acute threats facing the UK today. It blights communities, ruins lives, and is estimated to cost the country at least £37 billion each year. It affects more citizens, more frequently, than any other national security threat; and leads to more deaths in the UK than terrorism, war and natural disasters combined.

NCA officers are in the frontline against that threat. In recent years they've broken open networks that smuggle guns, drugs, money and people. They've tracked down fugitive criminals to their hiding places overseas, and also disrupted the gangs that supplied the runaways with their fraudulent passports. At home, their specialist officers support local police forces with complex investigations by providing niche expertise.

The NCA is also the lead agency dealing with the worst cases of child sexual abuse. The team is made up of a range of experienced, diligent professionals dedicated to protecting children. Cerys is one of those people.

ONLINE EXPLOITATION IS A GROWING THREAT

Child sexual abuse and exploitation means forcing or inciting a person under the age of 18 to engage in sexual activity. It includes physical sexual abuse, as well as online offences such as grooming, incitement, sexual communication, and creating or sharing child sex abuse imagery. Most people will naturally recoil from the subject, but vulnerable children rely on adults like Cerys not looking away.

It would be reassuring to think that these are fringe crimes, the work of a tiny, depraved minority, but the numbers tell a depressingly different story. According to the NCA's 2021 National Strategic Assessment, there are estimated to be between 550,000 and 850,000 people who pose a sexual risk to children in the UK alone. According to the National Center for Missing and Exploited Children's 2023 annual report, there was a 300% increase in online enticement between 2021-23, a figure that Cerys succinctly describes as 'insane'. And as horrifying as these figures are, even worse is the fact that more cases of child sexual abuse remain unidentified and Under-reported.

'When I first started working in this threat area,' says Cerys, 'there was a belief that you could arrest your way out of the issue.' If you put enough of the bad guys behind bars, you'd solve the problem and keep children safe. Years of hard experience have exposed that hope as painfully over-optimistic. Not only has the issue grown, but new technology has helped it develop in dynamic and disturbing ways. Children, who are often among the most enthusiastic early adopters of new technology, are particularly vulnerable to online exploitation by offenders, who use the internet to locate and groom potential victims for abuse.

'Historically, there was always this focus on, "We must look for the contact offenders. We must look for the offenders that are going to sexually assault a child in person,"' Cerys recalls. 'And what we've learned, as technology and offending behaviour have developed, is that people can do this by proxy. They can direct somebody else to do it on the other side of the world, and they are causing harm to that child, even though they're never going to be in the same room as them. And in the same way, people can engage one on one with a child virtually and cause physical and emotional harm to that child, without ever touching them. And the volume and the scale and the complexity of the offending is ever growing.'

In this formidably bleak landscape, the NCA faces off against its targets with the limitations that are common to almost every public body: a finite set of resources with which to take on an overwhelming demand for their services. As Cerys arrived at the NCA, that demand was about to go to a whole new level.

BUILDING AI INTO THE NCA'S STRATEGY

As early as 2018, the NCA had recognised that AI had the potential to help them achieve their mission. Within the UK public sector, this was impressively early to be thinking about AI: back then plenty of public bodies were running experiments, but relatively few of them were trying to build AI into their core operations.

But for the NCA, it wasn't just about integrating it into their workflow. They wanted to build it into their strategy. 'It starts with the business problem,' says Claire Smith, the NCA's Chief Operating Officer and a 25-year veteran of the policing and security sector. 'You always want to have a really strong business voice involved, and when I have seen innovation really work, that has been one of the key ingredients.'

As AI rises up the corporate agenda, most organisations will consider (if they haven't already) what their AI Strategy should be. Many will spin up AI Strategy programmes, ranging across the organisation, looking for all of the things that AI could possibly do. Consultants' two-by-two matrices abound, ranking lists of possible use cases according to 'technical feasibility' and 'impact'. Leadership teams are presented with the outcomes like a menu in a restaurant, with some recommended dishes: 'quick wins' for starter, 'low hanging fruit' for main, and if there's room left at the end then maybe 'longer term bets' for dessert.

The challenge with this bottom-up approach is that, in the final analysis, when an AI strategy comes up against the actual business strategy, there is only one winner.

Every good organisation already has a business strategy. The people there know what is most important to them - and to the boss. They can tell you the three or four priorities that the CEO cares about, and they fastidiously track the KPIs against which everything and everyone will ultimately be judged.

Being strategic about AI means using it to accelerate the things you already know to be most important. The instruction CEOs should give their teams is not 'design our AI strategy'. but 'test whether AI can help us meet our top three priorities.' If AI won't do that, then ignore it and focus on technologies that will. But if AI can help with those core goals, then you already have the answer to where and how to prioritise it.

At the NCA, a highly complicated organisation, their business strategy is clear and their priorities cleanly stated. Their main objective is the relentless disruption of serious and organised crime through targeted action against the highest harm offenders and networks, together with a statutory obligation to safeguard children from harm. The second priority is to minimise the number of victims and the level of harm caused.

By definition, their organisation only deals with cases that are really, really important. But with thousands of case referrals each day, they still don't have the resource to tackle all of them. So making their strategy work boils down to finding the most important needles in a haystack where every blade of grass is important, and then throwing their resource at those cases. It's painstaking work for the officers who do it, and also deeply stressful, knowing that somewhere in the pile there might be victims of serious crime they could help, if only they can find them in time.

Is that the sort of task AI can help with? Absolutely. So, with a minimum of fuss, this became one of the first priority areas for the NCA to focus their AI programme. And because time is of the essence in everything they do, they wanted to get to work quickly. For Cerys and her team, it couldn't come quickly enough.

“When we all got together we were one team with a shared goal. And that made it a really targeted development process.”

— Cerys Evans, G3 Intelligence Manager, NCA

A TOOL TO MAKE A MATERIAL DIFFERENCE FROM DAY ONE

In 2020, the team was handed an unprecedented trove of intelligence material from a confidential source. There were thousands of referrals, each one pointing to a case where children might have been harmed, and possibly still be at risk. And the only way to prioritise them was for human experts to methodically go through them one by one, painstakingly noting the key information and cross-referencing them with other sources of intelligence. With the number of cases they'd just been given, it would have taken them literally years to process. They needed it done much, much faster.

As it happened, Faculty were already in the building. 'We'd actually been engaged on a different project, to develop a different type of tool,' says Nijma Khan, who runs Faculty's Government and Public Sector practice. Then Paul Aspinall, the NCA's Intelligence Operations Manager - universally known as Asp - came calling.

'At the time I was responsible for developing innovation,' Asp explains. 'I looked at what we had, and what I knew we could do from my experience in data exploitation and intelligence, and I basically presented that to Faculty to say, "This is the challenge. This is what we start off with, and this is what we need to do."'

'We spent a few days with their different teams around the country,' says Nijma, 'and sat with them and tried to shadow them as much as we could, and walk through their day-to-day processes. We kept asking them the question, "What needs to be true for this to be an easy tool for you to use every day?" And through that process, we created a tool that was easily deployed, and made a material difference from day one.'

Claire backs that up. 'The magic happens when you put technical people together with people who understand the threat and the business, and you could see that with this group.' 'It didn't feel like we were working across multiple departments and agencies,' Cerys adds. 'When we all got together we were one team with a shared goal. And that made it a really targeted development process.'

It speaks to the team's ethos that when it comes to taking the credit, everyone involved is keen to point the finger elsewhere. 'Asp poured his heart and soul into the project,' says Cerys. 'He was constantly keeping things ticking along, and drawing us back to what the NCA tech infrastructure could handle, which of my big dreams were feasible or not. The successes and wins wouldn't have been possible without his constant drive and passion for the project.'

When pushed, Asp admits that it was originally his idea, but is quick to credit Cerys and her colleagues for how it turned out. 'I didn't need to be dealing with the intelligence development, the prioritisation side. I could leave that to Cerys, because she's kind of Champions League level on that. So she did that with Faculty, while I worked with them on the infrastructure and the commercial and legal stuff, all the horrible project stuff that nobody loves.'

And both Cerys and Asp are quick to heap praise on the wider team, including Faculty (Cerys in fact starts listing names that would fill the rest of this chapter). 'Faculty understood the value of what we were trying to achieve, and that it was clearly a challenge,' says Asp. 'And I've come to understand that data scientists love the challenge, almost over anything else. They were all about how to solve the problem, rather than selling a product.'

For good measure, Asp made sure the Faculty team fully understood how their work fitted into the NCA's broader purpose, briefing them as if he was onboarding new officers. 'They almost had to feel the pain of what the officer is trying to do, before they even got to the subject matter,' he says. Though thankfully, the Faculty personnel were kept away from having to see any of the actual content. 'The team did struggle with the work,' admits Nijma, 'because it was hard to have those conversations. But then that became almost the motivation to do the work.'

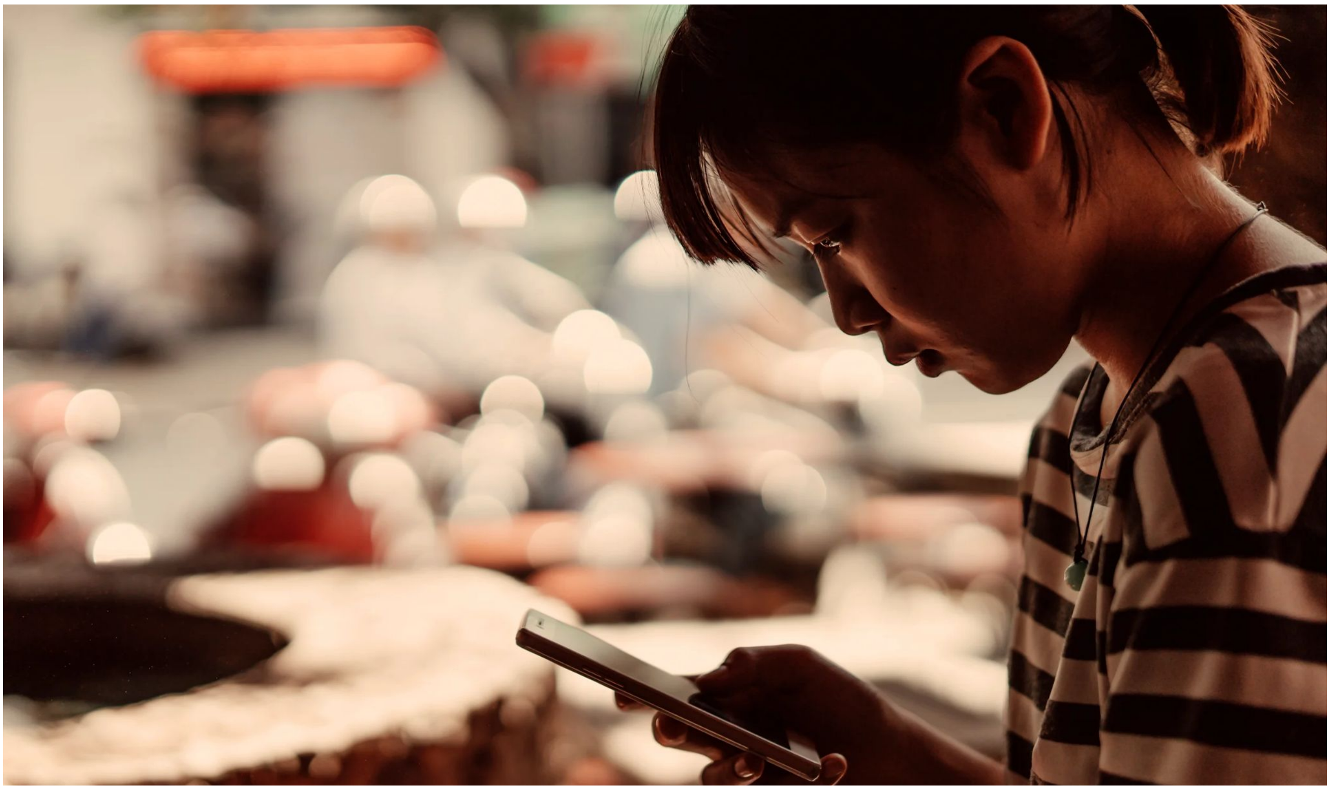

PROJECT 52 BECOMES VIPER

The tool that Faculty built, originally codenamed Project 52, eventually became known as VIPER: the Volume Intelligence Prioritisation and EnRichment tool. Though the name was suggested by one of Cerys' advisors, it was apt in more ways than one: the ancient Greek word for viper is 'aspis', or asp. And Asp's brainchild was about to start biting.

The technology consists of a suite of utilities that work together, enriching bulk data to triage cases and provide actionable intelligence. 'It's all about building up that investigation picture, building up the evidence base, and identifying where the harm sits,' explains Nijma. 'Work out if there's a potential for harm, extract insights from the data, and use the power of AI to work out where there's potential for harm far quicker than humans can alone.'

'What's important is that we can garner sufficient information to help us make that assessment,' Cerys adds. 'But we need to do that rapidly, because we want to respond quickly. We want to make the correct judgement, and put our resources where they're most needed, so we're safeguarding the most children from the most egregious harm that we can. And so the goal with VIPER was to obtain critical information rapidly and use that information to inform our assessment of priority.'

'The mantra for me was always deliver while you develop,' says Asp. 'Instead of using dummy data, we're using real data, and working with real risk. And the benefit of that is that you get real results very, very quickly. 'And within a matter of weeks,' concludes Nijma, 'we had something that they could use that made a material difference to their day-to-day jobs, but also a material difference to the safety of children around the country.'

KEEPING HUMANS IN THE LOOP

VIPER works by extracting key data from the referrals, which come in all sorts of different formats. 'There were a lot of complexities around those different data sources, and some really innovative methods Faculty had to use to solve that,' says Asp. Once the data has been extracted, VIPER identifies suspects and then links those people to information from additional data sources which might add context to help the NCA assess the risk they pose. That, too, goes faster with VIPER. Before it was automated, officers would have to make requests to data providers for each case that they were working on, and it might take weeks to come back. With VIPER, those same checks can be compiled at scale, across hundreds of investigations simultaneously, and take a couple of days.

That means fewer officers are able to do more with less. They can tap these additional data sources for 500 times more referrals than previously, building a richer picture of the risk and providing more timely intelligence. And all those different bits of information give Cerys' team a sense of how risky the person might be, and how quickly the NCA needs to either investigate them, or pass them on to a local force to go and knock on front doors.

'You can't arrest an identity; you have to find a real person behind that identity,' says Asp, his professorial demeanour belying the steel in his voice. 'Conversely, you can't protect an identity, you can only protect people, children. You have to find the real people behind those identities.'

'And it's not just about building up that picture and automating the task. It's also about increasing the accuracy,' says Nijma. Before VIPER, a lot of information would be manually transcribed and entered into the system, introducing more scope for error, more ways in which connections might be missed - or innocent people incorrectly drawn into the net. The NCA have always had safeguards built in to prevent that, but the automation gave them additional peace of mind, and again made it quicker. VIPER uses fuzzy matching to help spot duplicate 'entities', to reduce wasted effort and make sure all the right information is being linked to the right person.

But there's still a human in the loop, at critical points throughout the process, and making the final judgement. Given that the original problem was an overwhelming amount of data, simply adding more of it isn't going to help them. So as a final step, when all the analysis and enrichment is complete, the software completes a prioritisation assessment (using an academically accredited framework) and highlights the key intelligence in each case to the reviewing officer. This allows them to check whether they agree with the prioritisation a whole lot faster. What used to take 45 minutes is now done in four. This reduction in time on a case by case basis means that entire operations can be processed in weeks or months, rather than the years large scale operations or data dumps would take previously. And there's still more to do.

'We're still in the foothills of using AI as an agency,' says Claire. 'As we look across the organisation, you can see so many use cases. We can't just keep throwing people at the problem, the threat is too big. The data is too big and too partial. So we can just be so much more efficient if we use these technologies.'

In addition to her day job (on top of having recently organised her wedding), Cerys is now pursuing a PhD looking at how child sexual abuse cases are risk-assessed and prioritised. She's aiming to establish an even more rigorous evidence base for the process, which will address limitations in the current research base and can then be built into later iterations of the VIPER algorithm and support more meaningful prioritisation at scale.

Because it's an arms race. AI is affecting this dark part of the world just as much as everywhere else. In July 2024, the Internet Watch Foundation revealed that it was encountering so much child sexual abuse imagery generated by AI-tools, it had reached a 'tipping point' where authorities could no longer tell if an image involved a real child needing help. Ironically, the solution might also be AI.

'The nature of the challenge is already changing because of generative AI,' says Nijma. 'The thing you need to identify at the end of the day is: is there a real child in this picture?' Her team have already started work on a classifier that can analyse online imagery to identify illegal content relating to child sex abuse. If adopted, it would reduce the human workload needed to take it down by a third. 'That's my dream,' she confides.

But VIPER's capabilities aren't limited to tackling child sexual abuse. 'The tool is basically threat agnostic,' says Claire. 'The capability that they built for taking online identifiers, then processing that against other data and knowledge that we have, that is absolutely going to be repeatable across other threat areas. Tools like this will massively help in terms of surfacing risk quickly, and enabling that to be acted on.'

'Criminality is criminality,' says Asp. 'Criminals all generate data. They all leave footprints. The VIPER process would readily apply to absolutely any threat where you want to effectively identify targets from within data that you've been supplied with, or that you hold. And,' he concludes, 'we know that people have been targeted, in part or possibly even wholly, because of the work that we've done.'

'YOU DO IT BECAUSE YOU WANT TO SAFEGUARD CHILDREN'

The NCA's clarity upfront around the priorities that they wanted AI to support - putting AI at the service of their business strategy - meant that Faculty were able to get to work quickly, and have an impact almost immediately.

'The first time we put something through, start to end, it was just really exciting, because it meant we'd hit that minimum viable product,' Cerys recalls. 'We've deployed this operationally through development, and that's a real strength of the product, to trial these processes on live investigations while we were doing them.'

But although the technology is transformative, ultimately it's about the people. 'You don't work in this space on a whim,' says Cerys. 'You do it because you love it. You do it because you want to safeguard children. People really care so much about what they do,and are so dedicated, but because we have so much volume, we have to make sure that we have the most impact we can in terms of safeguarding children from harm. This technology helps us to do that a little bit better and make the right decisions.'

Nijma's been struck by it as well. 'The high point of the whole engagement has been the ability to work with such passionate people who do such an important job, and make sure that what we're building delivers for them,' she enthuses. 'Knowing the individuals who are using the tool, knowing that you've actually made a difference to their day-to-day, is really valuable. You rarely get that when you're building these kinds of tools for large organisations, but here we were able to get to know the teams, and so you can put a name and a face to the person whose job you're improving.'

Cerys is clear-eyed about the challenges that remain. 'The sad reality is that we get too many referrals to action every single one of them. So if we can't action every single referral, we damn sure need to make sure that we start from the highest harm and work our way down.'

As for the impact, Nijma points to a plaque that hangs in the front entrance to Faculty's Old Street offices. On the front is the NCA shield, which features a griffin and a leopard flanking a gold portcullis. The griffin symbolises courage and vigilance; the leopard fierceness and bravery. But if you turn it over, on the back of the plaque where visitors will never see it, is a handwritten message from Asp and Cerys.

'Thank you for all your hard work. You have helped to safeguard hundreds of children.'

THE LESSON IN SUMMARY

Business strategy trumps AI strategy.

For a small share of businesses, advances in AI will render their current strategy obsolete. If you are one of those businesses, then it may well be worth reconsidering from first principles how to succeed in a new operating environment.

◆

For the vast majority, the challenge is how AI can be used to accelerate you down the path you have already laid out. Rather than trying to come up with a separate AI strategy, you should test how AI can help deliver your existing priorities. If it can't, then ignore it and focus on the things that can. AI is not a worthwhile investment for every business.

◆

The opposite approach, a bottom up exercise, is common. It tends to involve lots of interviews with people from around the business, culminating in a menu of all the possible ways AI could be used, stack ranked against each other. More often than not this results in AI projects that operate at the margins. They rarely gather enough energy or interest to make a difference.

◆

The outcomes of AI programmes and teams should be judged primarily against commercial metrics. You need to make sure that you understand the top level business priority you are targeting, and the cause and effect pathway that allows you to influence them. At no point should the number of use cases delivered ever be mistaken for something important.

◆

AI programmes should always invest in properly baselining ex-ante performance, and measuring impact against that. This is often overlooked, making it impossible to build the feedback loops that allow performance to improve over time.

10

OPEN AI

LESSON TEN

If you don't control your models, your models control you.

Once an obscure academic niche, AI safety is now one of the defining issues of the age. Faculty is working with the most innovative AI labs on the planet, including OpenAI, to make sure that AI models are as safe as they can be, now and in the years to come. Because if we don't get this right, we'll lose control of our future.

A four-star US general stands in heated conversation with the Chinese premier. On the conference room wall, monitors display the latest news headlines. A drone swarm has attacked a warship in the South China Sea, bringing tensions in the region to boiling point. A US tech company has announced a breakthrough that points to an imminent future where artificial intelligence will outclass the human kind, but just two weeks later, a Chinese state-owned corporation has replicated the feat – with accusations of corporate espionage and theft flying across the Pacific.

The light from the TV screens silhouettes the figures, masking their expressions. Smoke curls in the lights suspended low over the circular conference table; vast concrete buttresses soar into the darkness above. If it looks like a cross between Dr Strangelove and a Bond villain's lair, then that's an apt comparison. It's no exaggeration to say that the future of the world hinges on the men and women in this room being able to agree on how to rein in these frightening new advances. If they can't, then humanity will have ceded control of our fate to a new, alien intelligence.

If you're comforting yourself that this is all make-believe, think again. This really happened.

'I DIDN'T COME HERE TO TELL YOU HOW THIS IS GOING TO END'

The meeting described above took place in July 2024, at an undisclosed central London location. The scenario was fictional, but the people involved were real. They included retired US army officer General Stan McChrystal; former UK National Security Advisor and Cabinet Secretary Mark Sedwill; the globally renowned Israeli historian and thinker Yuval Noah Harari; Jaan Tallinn, founder of Kazaa and Skype; and numerous other senior representatives from politics, academia and technology companies.

It was the culmination of a million-dollar exercise called Intelligence Rising that Faculty ran, in collaboration with the Tony Blair Institute for Global Change and concerned philanthropists, to improve understanding of the possible societal and geopolitical implications of AI. The scenario was devised as part of a narrative wargame that played out the likely impact of AI over the next ten years. The conference room was a set, artfully arranged by the Oscar-winning director Elena Andreicheva and her team, who filmed the entire event for a movie released in early 2025.

The wargame participants were all hardened power-players at the top of their professions, yet more than one of them left the room shaken by what they had experienced.

As AI becomes more powerful, its successes and failures will have greater and greater impact on the world we live in. But, as Intelligence Rising demonstrated, few people even in elite circles really understand what that might actually entail, or how profoundly it could reshape the world order. Not just the geopolitical order, but the human order. And even recognizing the problem is only the first step towards the really hard question: what should we do about it? What can we do about it?

As so often, if you want to understand the future, start by looking back.

RISE OF THE MACHINES

The CEOs of what are described as the 'frontier' AI labs – leading-edge companies including DeepMind, OpenAI and Anthropic – are now rock stars within technology circles. If someone working in the field says 'Demis' [Hassabis], 'Sam' [Altman] or 'Dario' [Amodei], they know they'll be understood. And the founders' fame is bleeding through into the wider world. Demis Hassabis of DeepMind was awarded the 2024 Nobel Prize in Chemistry for his work on using AI to predict protein structures. When ChatGPT launched in November 2022, it gained a million users in just five days, and reached a hundred million in two months, leading UBS to hail it as the fastest-growing consumer application in history (for comparison, Tik-Tok took a leisurely nine months to reach a hundred million users). Sam Altman was named Time magazine's CEO of the year for 2023.

It wasn't always like this.

Faculty's first encounters with frontier labs' founders go back many, many years. Faculty's early employees got to know them when AI was deemed a speculative field, and thinking about 'AI safety' was, to put it kindly, a niche interest. In those days, there was emerging awareness of the risk of algorithmic bias, but little understanding of the scale to which this could become an issue. Although many recognised the need for AI to protect privacy, few foresaw the range of ways in which this would need to be considered. And society was just beginning to extrapolate from these risks to consider, for example, how to ensure that future powerful intelligences would protect human values and support human flourishing.

Faculty's early community shared a concern with the visionaries of the frontier labs: how do we align machines and humans to work in harmony? Back in the early 2010s, long before Elon Musk or Steven Hawking or any of the world-famous commentators had piled into the debate, there were probably a hundred people globally with a serious interest in the problem. Faculty's staff were among them. It was, admittedly, a slightly strange crowd, prone to being dismissed as cranks. But it included many people, now at the forefront of the field, who had the foresight to recognise that safety was a crucial part of the path to ever more powerful intelligence.

But recognising that something's important doesn't necessarily make it easy to do what needs to be done.

‘IT IS IMPOSSIBLE FOR ME TO HARM, OR BY OMISSION OF ACTION ALLOW TO BE HARMED, A HUMAN BEING.’

At its core, the key question for safe AI is this: how do you make sure that AI acts in accordance with humanity’s values and rights? In AI circles, this is known as the ‘alignment problem’, having an AI that is aligned with human intentions, and not working against them. But from that simple formulation come a host of thorny questions. As humans, we often struggle to articulate what exactly we want. In fact, we often don’t even know what we want, let alone how to express it. And even if we could, AI can’t just be a genie that slavishly gives its operators what they ask for. The things people want might be malicious, contradictory, or counter-productive.

An advanced AI needs some kind of values to bound and inform it, but who gets to decide those values? And how do you encode them in what is, after all, simply a piece of software? Do you go for a legalistic, rules-based approach – or try to teach machines the precepts of moral philosophy?

The development of AI itself shows that trying to hard-code rules about the world into software will never be as flexible, adaptable or useful as coding models that learn their own rules. But even if we were able to teach the AI our values, how would we make sure the algorithm interpreted them as we would want? Software can have bugs, and life can throw up edge cases that defy any attempt to find an ethically tidy solution.

And if we could overcome all these hurdles to create an AI that perfectly understood human values and applied them flawlessly, there’s still a question as to whether it would be right to do it. If we hard-wire a machine with early 21st century values, would that position be locked in for eternity? Think back to the 1800s, where slavery was legal, women were men’s property, and beating children was considered the hallmark of good parenting. If AI had been invented then, would we still be living now yoked to a technology built on those values?

Although it’s tempting to believe that our current moral order is the apogee of civilisation, in another two hundred years it’s likely that aspects of our own society will seem as hateful to future generations as wife beating and child labour seem to us. How can we make sure that our children are allowed to adopt a different approach to the world than the one their parents took?

You can see why the people asking these questions came across as cranks and obsessives. The questions seem esoteric, unmoored from the concerns of ‘normal’ people or businesses.

In fact, they’re everyone’s concern.

‘I’M SORRY, DAVE. I’M AFRAID I CAN’T DO THAT.’

The time when AI safety was the obscure hobby-horse of a few dozen enthusiasts now seems an age away (or, by the warp-speed timescales of AI’s development, about ten years ago). The frontier AI labs, and many governments, recognise the potentially catastrophic risks of misuse, and are investing heavily to address the issues.

Faculty has grown up too, though it’s still at the forefront of AI safety, working with labs and governments to conduct novel research, develop new tools, and assess risk. The company is one of the first ports of call if one of the frontier labs needs to test the safety of its newest model, as when OpenAI wanted to check out its o1 model. The model showed a step change in reasoning abilities, and its creators wanted to be sure they’d done everything possible to deploy it safely. ‘It’s absolutely paramount that foundation models are built safely,’ says Sam Altman, OpenAI’s founder. ‘I know Marc and Faculty have cared about AI safety for a long time, and so they’ve been a natural and wonderful partner for us on this work.’

Faculty also works with the UK AI Safety Institute and other organisations to make baseline safety assessments of general purpose models. Faculty’s robust capability assessments test models in different ways, ranging from question-and-answer engagements by experts, to full-scale randomised control trials that test what bad actors might be able to achieve if assisted by AI, against what they can do without it. Faculty have also piloted cutting-edge techniques to improve safety, such as the feasibility of a model ‘unlearning’ dangerous knowledge that a bad actor could use.

Other risks have less spectacular outcomes – no explosions or homebrew bioweapons – but operate in more insidious ways. Models that contain biases could present significant risk to a broad population if they create discriminatory outputs, whether that’s in the content they generate or the decisions they take. This has been an issue for AI since well before the latest developments in generative AI, and Faculty has long been a leading light in research to identify and mitigate biases. In 2020, the company provided the technical assessments on which the UK government’s ‘Review into Bias in Algorithmic Decision-Making’ was based.

But the most intrinsic biases don't come from poor coding. Particularly with generative AI, the biases slip into the model with the raw material of its training and tuning data. AI is trained to represent the world, but modern society is the outcome of centuries of complex biases and discriminatory approaches. So there's a fundamental question: do we want AI that represents the world as it is, or as we would like it to be? And if the latter, as who would like it to be? Different cultures, personalities or political systems might have very different ideas of what an ideal world would look like.

The frontier labs are all committed to weeding out bias and discrimination in their models. They want to make sure that their models don't harm major parts of society, and that the benefits of AI are felt inclusively. Whatever your demographics, you have a stake in the labs getting this right, and finding the right balance between the 'world as it is' and the 'world as it could be'.

'NO 9000 COMPUTER HAS EVER MADE A MISTAKE OR DISTORTED INFORMATION'

But – crucially – safety is contextual. Language that might be entirely acceptable for a model helping a screenwriter develop their characterisation might not be so appropriate for a child doing their homework. So there's a limit to how far model providers can be responsible for the safety of their products. They can't foresee all the contexts in which their models will be deployed: only you, the person or organisation using it, have that understanding.

If you integrate a chatbot for customer service, the frontier labs won't be making sure that their models aren't rude to your customers. They won't stop the model from deciding to give ruinous discounts, or offering unauthorised financial advice. Just as with any other intelligent entity you employ, you will have to define what is appropriate and acceptable for your situation, and you will have to check that those boundaries are implemented correctly, because only you understand how the model will work in your use-case. And you will be accountable if it goes wrong, whether that's because you gave the AI poor instructions, introduced a bug or even chose the wrong model for the purpose.

Ultimately, for every version of the alignment problem faced by Sam Altman, Demis Hassabis or Dario Amodei as they develop the world's most advanced models, there is almost always a parallel problem that 'normal' organisations face when trying to make those models do what they want.

So what should you do?

'I CAN ONLY SHOW YOU THE DOOR. YOU'RE THE ONE THAT HAS TO WALK THROUGH IT.'

Fortunately, if other businesses share versions of the problems faced by the frontier labs, they can also learn from the solutions. The same safety techniques that Faculty has implemented with OpenAI, the UK AI Safety Institute and others, give your organisation a suite of tools you can choose as appropriate.
The first and most fundamental thing you can do is keep humans in control. This is Faculty's cornerstone philosophy when it comes to AI safety, but there are different ways of achieving it.

The most direct approach is to have a human 'in the loop': providing input at critical stages in a process so that the person maintains oversight and control. As AI models are probabilistic and always contain a degree of uncertainty, in many circumstances it will make sense that they should make recommendations to a member of the team, who understands the context and remains ultimately responsible for taking the action.

But there are plenty of circumstances where it's impractical to insert a human into the loop, for example when the process needs to operate quickly, at high volume, or both. Many web services look like this. Content and product recommendations on Amazon and Google, or ad targeting at Facebook, all operate at a pace and scale beyond human intervention. As does ChatGPT.

In these cases, you rely on humans 'over the loop' to robustly test the models before they are deployed, and set parameters that constrain their outputs.

'A big part of how we make sure that our technology is safe to be deployed into the wider world is our "red-teaming" programme,' says Sam Altman. 'We ask people and teams that we trust, like Faculty, to help us assert that our models are going to meet the safety standards that we set out.'

Red-teaming (the name derives from military war-games, where the adversary is always the 'red' team) brings together teams of experts, both in-house and external, to examine the AI models. The group stress tests what the models are capable of, and how those capabilities might cause harm in the world, whether intentionally or not.

As well as testing and red-teaming, we can use safeguards and set operating parameters to ensure safe outputs. Anyone who has tried to make ChatGPT say rude things about someone will have seen this kind of constraint in action. These safeguards can be designed and implemented at various stages through a model's lifecycle.

“A big part of how we make sure that our technology is safe to be deployed into the wider world is our ‘red-teaming’ programme... We ask people and teams that we trust, like Faculty, to help us assert that our models are going to meet the safety standards that we set out.”

— Sam Altman, CEO, OpenAI

For instance, if you know in advance that you don't want your model to give financial advice, you could remove material that would increase that risk from the dataset you use to train it. Once it's been trained, you could further tune it to avoid giving that advice through a process called 'reinforcement learning from feedback' where the model learns what it should and shouldn't do in accordance with feedback provided by humans or by another AI.

If you have created a generalisable AI model, but you later decide that you'd rather minimise the likelihood that it will provide financial advice, you could append an instruction to that effect to every user's query. And, if really necessary, you could implement a 'classifier' model to check every output before it goes to the user, just to be extra sure. Of course, highly capable users of the system may still be able to 'jailbreak' it to do something you don't want, but they'll have to work hard to do so.

Finally, governance is key. AI safety should be acknowledged and owned at a suitable level within any organisation. It should ultimately be part of the governance process that is used for other important categories of risk.

Faculty is keen to ensure that this ability to implement AI safety isn't limited only to organisations with the technical heft of the leading research labs, so they have developed a platform called Frontier that makes this much easier. It enables individual models to be parameterised and governed. And it allows parameters to be implemented to constrain collections of models too, so that leadership teams can set policies that bind all connected AI systems right, across an organisation.

THE UNKNOWN FUTURE ROLLS TOWARDS US

Ultimately, AI safety is not just for tech CEOs or government bodies. All of us have a responsibility to consider safety in our own contexts, and to add our voices to the debate about how we want AI to shape our future.

We all face a choice: either prepare now and consider the safety of the models we use in our business our responsibility, or panic later once we realise we have lost control of our models.

There will always be a temptation to charge ahead with the latest technology and leave safety as an afterthought. But, done right, safety doesn't have to come at the expense of capability. Cars are faster and more efficient than they were 50 years ago,and they're also much safer. Similarly, the scale of long-term challenges like AI's alignment with humanity shouldn't freeze us from dealing with immediate issues, like algorithmic bias.

If humanity gets this right, we can control AI models and unlock their benefits in a safe and responsible way. If we get it wrong, those models will control us, making decisions for us and about us based on values and approaches with which we may not agree.

Back in the conference room, there's still no breakthrough. The clock ticks down as the Chinese and American delegations face off. 'AI could threaten mankind itself,' warns a visibly troubled General Stanley McChrystal. 'What are we going to do about it? We need big, bold thoughts.'

The ending of the game is revealed in the Intelligence Rising film. But real life isn't a game: AI safety is a constant work in progress, an everlasting dialogue between technology, humanity and the world.

Intelligence is a tool that people have been using for tens of thousands of years. Now machines have it too. It's been put to terrible purposes like exploitation and destruction; but it's also built civilisations, created art of astonishing beauty, and allowed us to do things in everyday life that our ancestors thought were reserved for the gods.

How intelligence gets used in the future is a story we all have to write together.

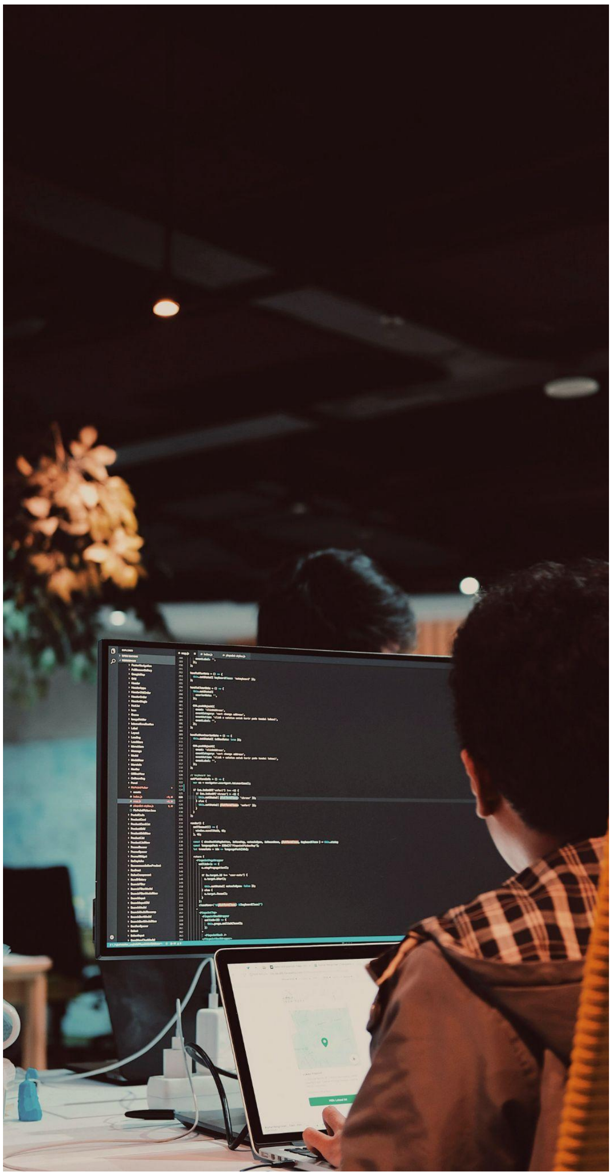

THE LESSON IN SUMMARY

If you don't control your models, your models control you.

AI models are probabilistic. The most powerful are black boxes. They don't always behave in entirely predictable ways. And users can't really tell why they do the things they do. This creates a new set of risks.

◆

As AI technology is embedded more deeply in business processes, it is essential that the correct controls are put in place around it.

◆

In processes where individual decisions or actions are valuable, humans should be kept in the loop. AI should support them not replace them. Decision support systems should be designed to give users well-targeted and parsimonious analysis, rather than drown them in data. And they should be interpretable and interactive.

◆

In processes where the volume or frequency of decision or actions makes it impractical to have humans in the loop, they should nevertheless be 'over the loop.' This means that they are able to specify the parameters in which the AI models operate, and interrogate their outputs.

◆

Implementing this at the level of an organisation will require technology platforms that allow those responsible for governance to set policies that bind all of the AI models that operate across an organisation. Faculty's Frontier platform is designed to do exactly this.

Governing in the age of AI: a forward look

There is little doubt that AI will change the course of human progress, much like previous general-purpose technologies that dramatically reshaped the world around them. But unlike past waves of change, much of the foundational infrastructure for AI is already in place: the internet and data, cloud and storage, chips and compute. The scope and scale of change will be vast. And it will come quickly.

The private sector is already making historic investments in its future. With chips and data centres, leading tech companies are building infrastructure that surpasses 20th-century mega-projects such as railroads, dams and even space programmes. But across the corporate world leaders all face a choice: invest inAI capabilities or risk perishing.

For governments, the choice often feels less stark. Political leadership may change, but the state still exists. Like all well-established organisations, the state has a bias towards caution. But this is an illusion – a failure to modernise, reform and deliver is a perilous course for a nation and those who govern it. And this is particularly true in the case of AI, which if gripped properly, should make today the most exciting and creative time to govern.

We are coming at this issue from both perspectives. One of us is a politician and runs an Institute advising government leaders, while the other is a leader of a technology company. We both understand the magnitude and the necessity of the choice. We both also see the potential prize for the UK, which should have its own ambitions to position itself at the forefront on AI and provide leadership on governing in this new era. And when both of us survey the operations of governments from our different perspectives, we see the same opportunity: almost everywhere AI can help us reimagine the state. Many of the countless daily tasks in government are repeatable processes carried out on a mass scale.

Almost all of these can be made better, faster and cheaper. The scale of this opportunity is huge: with the technologies and the digital infrastructure we have today, we estimate that up to £40 billion can be saved each year with the technology as it exists now. But, of course, over time, this technology will accelerate dramatically in its capability, and so will the savings.

This is much more than a debate around the margins of tax and spending; it has the potential to transform the costs, functions and accountability of government. At a time when government is unwieldy, expensive and slow, AI can save our public services, making them more personalised and human-centric. Safe, explainable AI systems can make government fairer and more transparent, liberating and empowering people. We shouldn't be afraid of blocking systems that don't meet these standards, but we must rapidly embrace those that do. They can make government more strategic in how it approaches complex decisions about the highest-stakes issues, with more accurate, more granular, more up-to-date information and insights.

And this is only the beginning of what AI will be able to achieve. The pace of development and the new capabilities announced each month make it clear that the current generation of AI systems only give us a glimpse of their full potential. This is the least able AI will ever be. To access this opportunity, government will need a coordinated strategy to put in place the necessary infrastructure, sovereign capability and skills. It will need to invest in making the right data across departments interoperable, while maintaining privacy. It will need to train its own models where necessary, such as for national-security purposes, finetune custom tools and build or procure applications on top of existing models. It will need to secure the computing power necessary for AI to run at scale, for everyday use as well as research purposes. And it will need to change how it hires and trains AI specialists.

None of this will be possible without working in partnership with the private sector. The computing requirements of AI mean that close coordination and cooperation with leading providers are required. The UK is also itself home to many leading AI companies. With the talent that we have, it should be home to many more in the future. The government will play a crucial role in fostering this industry if it makes the right choices and clearly demonstrates what AI can help us achieve.

For those of us in both the public and private sectors, the choices that we face today are critical for our futures. Businesses which fail to adapt to this new world will be quickly replaced by competitors. For countries, the failure is bigger – harming people's prosperity as well as their nation's place in the world.

The prospect might seem daunting, but for the most part investing in AI is low-risk, high-reward. Its benefits far exceed the costs – and the price of inaction may be higher still.

Sir Tony Blair
Executive Chairman, TBI

Marc Warner
CEO, Faculty

The next decade for enterprise: a look forward

As we stand at the foothills of an evolving change, and try to make out the shape of the peaks ahead, much of the detail is inevitably hidden behind clouds of uncertainty. This is true for any technology shift. But, in the context of AI, the impact of the change and the uncertainty around it is extraordinary - perhaps unlike any technology humanity has ever encountered.

Firstly, as two representatives of organisations that live at the frontier of this change, day in and day out, we would like to start by reassuring you; it is ok not to know with certainty what AI means for your enterprise in three years time. We don't have the full answer ourselves. We regularly speak to the most notable experts, at the forefront of the field, and they don't know either. Anyone who pretends differently is probably trying to sell you something that won't work.

Yet even with this level of uncertainty, it is still our collective responsibility as leaders to navigate our businesses through this. So, to set ourselves up to do that well, it is important to reflect on what we can be confident about, and how to make good choices where there is less certainty.

Let's start with some of the things that we know to be uncertain.

AGI, AND AI'S UNEVEN IMPACT

The biggest looming uncertainty is whether we might reach human-level artificial intelligence, often called artificial general intelligence (AGI). And, if so, when? Estimates vary from about five years upwards. Almost all estimates have been diminishing rapidly over the last five years. A century was a frequent estimate ten years ago, but today you'd be hard pressed to find an expert willing to consider such a long timeframe.

This possibility should animate us profoundly as citizens, parents and grandparents. AGI would be a technology development without precedent, it would signal the end of human intelligence as the dominant form of intelligence, and would likely change all aspects of our lives. The outcomes could vary between post-scarcity utopia to existential risk for our species. It is worth all of us learning more about what this could amount to, and then using whatever platform we have to advocate for the development of safe AI systems.

But precisely because this eventuality has the potential to disrupt all aspects of life as we know it, it is almost impossible to plan for. The best way to handle this in our roles as business leaders is therefore to pragmatically deprioritise to first order. To place AGI in the same category as an asteroid strike - something that most businesses can't do anything to change, and which will reset all current plans.

Another, more proximate uncertainty, is the relative impact that AI will have across different sectors. Think back to when the internet emerged. It did not disrupt every industry equally. With a bit of a squint, you can distinguish between a couple of rough types.

First, there were sectors in which new internet-based companies arose and changed the landscape entirely. Amazon and its ilk permanently changed the dynamics for retail. Google did the same for advertising.

And then there were the sectors in which change was less pronounced; airlines, construction, manufacturing. These companies all still had significant efficiencies to gain from the adoption of digital and communication technologies. But their worlds were not turned inside out by them.

History teaches us to be cautious about trying to divine which sectors will be most affected by AI. From the vantage point of the turn of the millennium, it may have been possible to see that ecommerce would affect retail, and search would affect advertising. But few could have predicted the impact that companies like Airbnb and Uber would have on hotels and taxis.

Nevertheless, there are some areas that seem more likely to see outsize change than others. For example, we can already see that sectors which revolve around the processing of written and structured information are moving quickly. Insurance is one. The legal sector is another. And education, which has changed little since Victorian times, has the potential to look very different indeed by the end of the decade.

In these cases, where even though we may not know the ultimate destination, we can see that a journey is beginning and that leaders should make sure that they are taking the likelihood of change seriously. That they are plotting a path that feels directionally correct, and setting up their organisation to be nimble in the way that they adopt AI. Regular stress testing of corporate strategy can help with this, for example by building tools like wargaming and scenario modelling into quarterly planning rather than keeping them as an interesting adjunct to the annual away day.

DESIGNING FOR HUMANS AND MACHINES TO WORK TOGETHER

Despite these considerable uncertainties, there are nevertheless a set of things that we can be confident about in the age of AI, and which provide terra firma for good business decision-making.

The stories of this book describe a number of them; good practices that will help any business to adopt the technology with confidence, regardless of what they are trying to achieve with it. And, importantly, these lessons have been abstracted beyond the specific details of the technology - whether convolution neural nets, transformers, or whatever is next - so will likely remain true over time, even as the technology develops.

Beyond these, though, there is something else that we can and should be confident in. We refer to this as a 'human-first' approach to AI. A frequently overlooked fact is that until we have human-level artificial intelligence, the best decisions will be made by a combination of people and machines working together. We have seen this play out before. In 1997, Garry Kasparov was beaten at chess by IBM's Deep Blue. But this did not herald an era of machine dominance. For nearly 20 years humans and machines playing together - "Centaur Chess" - could beat either playing alone. It was only in the mid-2010s when machines like Stockfish started to beat human machine combinations, and 2017 when Google's Alpha Zero cemented machine dominance.

So for the foreseeable future, we need to build AI systems on the basis of having humans and machines working together. This will deliver the best performance outcomes. And designing systems to keep humans in control is also the best route to ensuring that the technology is deployed safely.

Let's explore how we go about doing that.

HUMAN IN THE LOOP, HUMAN OVER THE LOOP.

One of the key considerations of a human-first approach to AI ensuring that humans remain in control.

The most well understood approach here is to design to keep a human-in-the-loop. This will sometimes be the right solution. But that concept alone is too constraining, and too fraught with error over time. We know that some processes need to happen at greater speed or scale than humans in the loop can accommodate. We also know that individual humans can make mistakes, and those will only be obvious in the aggregate.

The solution to this is a higher level of control; human-over-the-loop. This enables a user to place controls over models or collections of models, to govern the system at scale. They can then step 'out' of the loop, allowing it to run at high speed or volume, knowing that the model outputs are constrained by the parameters they have put in place.

USER-CENTRED DESIGN FOR THE INTELLIGENT ERA

A human-first approach to AI also requires it to be integrated into business workflows in a way that maximises its usefulness to human users. We take a 'decision-centric' approach to this, zeroing in on using the technology to support the human decision-making that is so central to most business processes.

This requires both methodology and technology. We need the methodologies to apply user-centred design to the era of intelligent software, by helping users to map their decision-making flows, and design AI systems that plug into them in a precise and targeted way.

We also need technology to build the AI systems quickly and easily to fit these user requirements. Faculty and Accenture are collaborating on the Frontier platform. This goes much further than serving the predictions or content that AI models generate to users, and allows users to interact with AI models in new ways. For example, it wires models together and connects their outputs to business KPIs, so that users can understand the cause and effect relationships that drive business outcomes. And it allows users to test the impact of the different choices in front of them before they commit to a decision.

We can be confident that this level of user-centric design and functionality will be needed if enterprises are to deliver impact from AI. Not least because we can see what happens when they're absent.

STEPPING OVER THE TROUGH OF DISILLUSIONMENT

As the last few years of GenAI POCs have shown, LLMs are a very powerful round peg. But most valuable enterprise problems are a square hole. It's simple to understand why. LLMs enable AI to treat text like it has always been able to treat numbers. But how many enterprise business processes are only text? A few, perhaps, mostly centred around customer service. But most other processes are combinations of text and numbers. And even then, how many processes are limited just to the analysis or generation of data; text, numbers or otherwise? Precisely because they aren't - because they involve other inputs and outputs, and human users and customers, and integration and governance requirements simple, isolated LLM deployments will never lead to the transformation we are hoping for.

If you look back at the history of building software, there is one lesson that stands out - building technology for technology's sake rarely generates the value we hoped for. Whether you call the opposite approach user-centred design, product thinking, or agile, a strong focus on user requirements has proven to be important over and over.

Human-first AI is an extension of this thinking into the AI space, by focussing the construction of intelligent systems on the user needs. We can be confident that it offers a way of building more useful, more powerful and safer systems that will transform our societies for the better.

Faculty and Accenture are pioneering this new approach together. As responsible technologists, we believe that AI systems should serve people, not supplant them. We want to create safer and more powerful systems for our customers. And we're privileged to have been given the responsibility, hundreds of times, to demonstrate the results this approach can deliver for them.

Shaheen Sayed
Head of Accenture, UK, Ireland & Africa

Marc Warner
CEO, Faculty

Acknowledgements

Special thanks to Edwin Thomas, whose deft pen and commitment to quality gave a weight of character and charm to the stories told in this book.

There are too many people from inside Faculty and from all of our customers to mention here. But without the input of those listed below the work described in these pages, and the lessons learned, could not have happened.

Thank you.

Aaron Ark
Abbey Cooper
Adrien Yeung
Aidan Patrick
Akhlas Hafiz
Alberto Favaro
Alessandro Alcinesio
Alex Clibbon
Alex Horton
Alex Hurley
Alex Rogers
Alice Fee
Alice Kluzer
Alistair Stuart
Andrew Brookes
Andrew Glennie
Andy Li
Andrew Liubinas
Andrew Perry
Andrew Perch
Andrew Skirrow
Andrew Towers
Andy Keegan
Angie Ma
Anna Carolina Tortora
Ashwin Chopra
Aydin Utting
Azeem Alam
Balaji Ramanan
Becky Friend
Ben Bedford
Ben Green
Ben Walker
Birgit Burk
Brandon Lim
Bryn Mathias
Carlos Ferreira
Caroline Mendonca Costa
Carolina Sportelli
Caroline Costa
Cerys Evans
Chethan Tarikere Navanitha
Chris Monit
Chris Wyard
Christian Nickerson
Christoph Rahmede
Christopher Monit
Colin Webb
Conor Saxon
Cosima Calder
Darya Nizhnikova
Daniel Flintoft
Dan Fat
Danica Howard
David Bholat
David Jasek
David Peinador Veiga
Divish Rengasamy
Dylan Binding
Ellie Morris
Emma Bannon
Erick Hinds Mingo
Ermioni Mavroudi
Eron Cemal
Esther Clark
Eugenie Lale-Demoz
Eyeni Kantsey
Federico Celli
Freddie Parton
George Delis
George Wheatcroft
Greg Jackson
Guilherme Castro
Helen Burke
Hugh Neylan
Iain MacKay
Ibukun Ogundipe
Indranil Dasgupta
Jack Mitchell
Jake Luscombe
James Grady
James Lowe
John Gibson
John Mansir
Jon Holdship
Jonathan Barcroft
Jonathan Street
Jonathan Tay
Joshua Muncke
Julia Sala-Bayo
Julian Leslie
Karland King
Kate Attfield
Kate Morgan-Baker
Kate Orchard
Katrina Soderquest
Kelly Choong
Kenny Wong
Kevin Baboolal
Kim Moore
Kirsty Kelly
Lama Ahmad
Laura Palacio Garcia
Laurence Cowton
Lauren Cooper
Louis Claxton
Luke Geoghegan
Lyndon Ruff
Magnus Cormack
Mandeep Dhansa
Marc Steeven
Marc Warner
Marian Kantor
Mark Norgate
Mark Moerdyk
Martin Cuesta Allende
Mat Escobar
Matthew Harker
Matthew Haugh
Maxwell Lawson
Megan Lathwood
Megan Morrison
Michael Keeling
Michael Lampe
Myles Kirby
Naorin Sharmin
Nathan Nile
Nathan Windle
Neal Pawson
Nijma Khan
Niko Louvranos
Nim Patil
Niren Patel
Orla O'Neill
Oscar Bennett
Oz Tandiroglu
Paul Aspinall
Paul Hose
Paul O'Dwyer
Petra Pavlovic
Phil Holland
Piers Rudgard-Kiltie
Riddhi Dhall
Richard Mather
Robert Rankin
Robi Descrovi
Romain Bourboulou
Ronan Jamieson
Rosalind Berka
Roya Ahmadi
Ryan Hannam
Sandhini Agarwal
Sarah Laughton
Scott Barnes
Sher Lynn Wong
Shubhi Rajnish
Sofiane Berbert
Sophie Carder-Jones
Spyridon Kolaitis
Spyros Georgakopoulos
Srish Bashin
Steph Skeet
Stephen Wallace
Tarja Strahlman
Tarcila Steter
Tarek Soubra
Tessa Farrington
Tiberiu Oprea
Tim MacFarlane
Tim Williams
Toby Thornton
Tom Davies
Tom Drew
Tom Nixon
Udara de Silva
Vanessa Ortega
Victoria Derks
Vipal Adroja
Volodymyr Gromoglasov
Wen Yang
Will Rainbow