Imperial College London
Enterprise

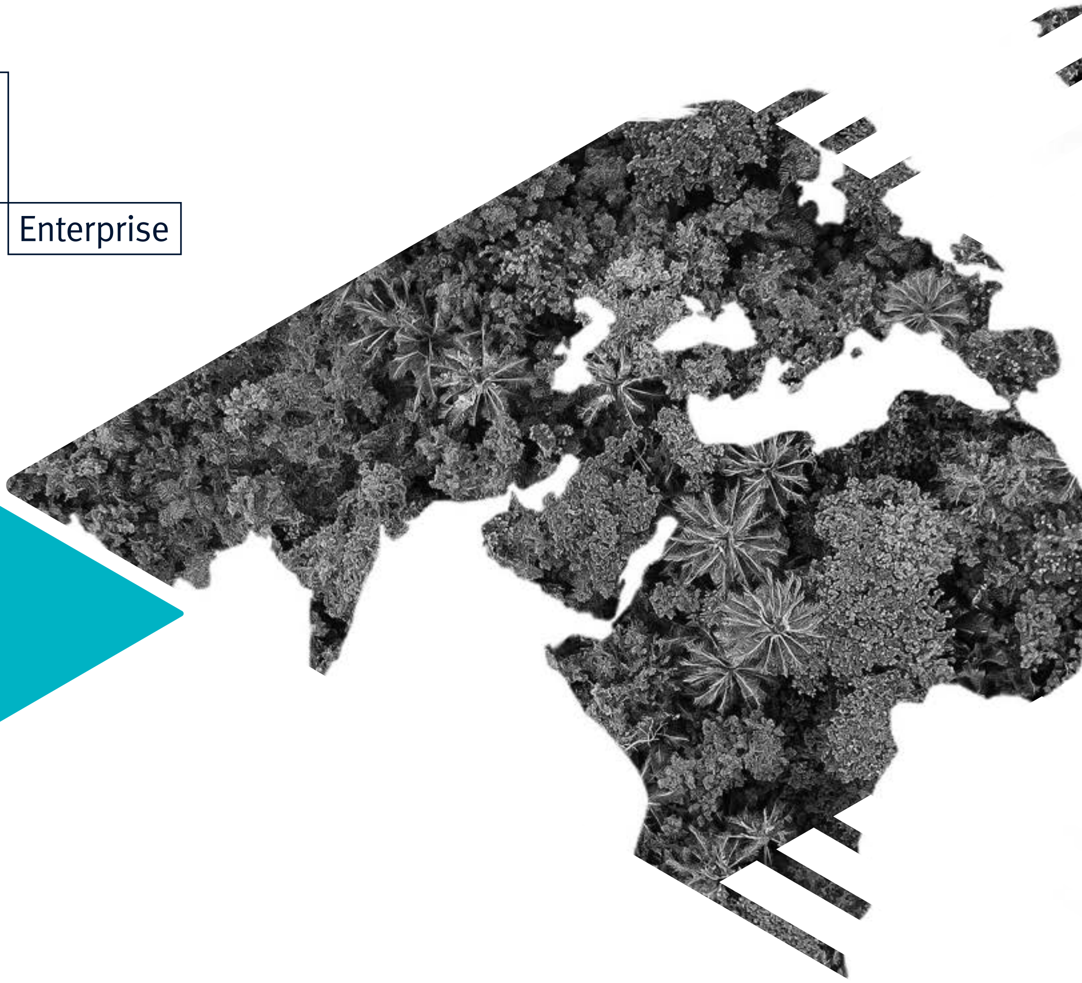

Using AI to mitigate risks from natural hazards

Phil Ormrod

As the effects of climate change accelerate, risk managers in government and business are finding it increasingly challenging to understand and mitigate the human and economic risks of natural hazards such as wildfires and cyclones. New developments in artificial intelligence promise solutions.

Imperial Business Partners

As the effects of climate change accelerate, risk managers in government and private enterprise are faced with increasing challenges as they seek to mitigate an uncertain future.

In May this year, the European Forest Fire Information Service reported that 2022 was the second-worst wildfire season on record, surpassed only by 2017. Given the current situation in Greece, this benchmark will be surpassed before long. And over four days in early June, the global average air temperature record was broken three times; previously, this had only happened twice over the preceding six years.

In other words, the pace of change is rapid and accelerating. And with vast human, ecological and financial consequences – last year's wildfire season was estimated to have cost €2 billion – there is a huge imperative to model and understand the hazards that might result. But doing so with enough accuracy and efficiency is difficult.

Global climate models, for instance, can perform hundreds of billions or even trillions of calculations to model large-scale weather patterns at half-hour intervals over a hundred years – but can't realistically be run at a high enough resolution to satisfy a risk manager seeking to understand the variation in storm surge probability along a particular coastline.

And with a run time of weeks or months depending on the hardware available, the process is difficult to iterate effectively. Iteration – running the same predictive model a number of times – is vital to the accuracy and reliability of its results.

It's also challenging to keep track of unfolding events in real time. A huge range of resources are available to track wildfires as they unfold – from satellite arrays to social media. But cleaning, weighting and integrating noisy and heterogeneous data sets is a real challenge under such intense time pressure.

At Imperial, experts across a range of disciplines have found that AI tools can help them address these challenges: to improve long-term modelling for events like hurricanes and tropical cyclones, hedge risk in vulnerable

Cover photo by Marcus Kauffman on Unsplash investments like forestry, and create systems that can model natural hazards as they develop – and the effects of potential responses, too.

Machine learning offers unprecedented power to understand complex relationships between the enormous numbers of variables in current climate models. Natural language processing enables crowd-sourced information to direct satellite surveillance; and convolutional neural networks can process the resulting information with striking accuracy. Digital twins allow responders to simultaneously trial a range of measures against an unfolding emergency.

And as the frontiers of this work continue to develop, researchers continue to improve the explainability of AI-based predictions, as well as ensuring that they reflect real conditions grounded in physics as well as statistical probabilities.

AI tools help crisis managers combat wildfires in real time

As the effects of climate change gather pace, wildfires are growing in frequency and intensity around the world: in Greece and the wider Mediterranean; in California and Canada; even Siberia and the Russian Arctic.

Responders can draw on vast streams of data to track them as they evolve; but filtering and assessing this information fast enough to make effective decisions is a huge challenge.

Dr Rosella Arcucci, a member of Imperial's **Department of Earth Science and Engineering** and **Data Science Institute**, leads the Data Assimilation and Machine Learning (Data Learning) research group. As a mathematician and computer scientist, one of her key projects is to implement the operational use of AI tools in managing just this kind of natural hazard. As she puts it: "The key word is usability: even when you're doing fundamental research, you always keep an eye on the operational priorities. For example, in healthcare, speed can be important but not crucial – it's accuracy that's key. With wildfires, that's not the case: you have to be really fast. So we work on efficiency. What's slowing us down?"

The first answer is knowing where to look. And the solution is perhaps a surprising one: natural language processing, a branch of AI more often associated with chatbots. Dr Arcucci explains: "We've built a nowcasting tool: a social media scraper, with filters for language, bots and so on, that picks up on wildfire-related keywords and uses sentiment analysis, then sorts and geolocates the relevant posts. Using this, we can aim the available satellites. That saves us something like three hours, which in this situation is a lot."

Dr Rosella Arcucci, who leads the Data Assimilation and Machine Learning (Data Learning) research group.

In other words, Dr Arcucci and her team are able to crowd-source from a vast but extremely noisy pool of data – the endlessly diverse public conversation on social media – and use a proven AI tool to quickly and simply focus on what's relevant.

Wildfire surrogate model

Having located the fire, responders need a tool that can respond to its development. This is the Wildfire Surrogate Model: a digital twin that can quickly be assembled by integrating the targeted satellite imagery into a pre-trained neural network, constrained by mathematical representations of wildfire physics.

A digital twin is a virtual representation of a real-world entity or process, which can forecast how its counterpart will develop and respond to interventions. A neural network is a type of artificial intelligence system inspired by the structure of the human brain, which processes information through layers of nodes, like virtual neurons. As the network learns, it builds connections between these nodes, whose weight depends on the significance of the relationship they represent. In deep learning, neural networks are divided into multiple layers. With tens or even hundreds of layers of nodes, the connections between them can number in the hundreds of millions, allowing the network to process huge amounts of data very quickly, and to recognise highly complex statistical relationships within it.

Using this technology, the Wildfire Surrogate Model can simulate the movement of the fire front and its response to different barrier placements, allowing responders to trial and choose between strategies. Crucially, it can do this thousands of times in a few minutes. Dr Sibo Cheng, a Research Associate in Imperial's Data Science Institute and Department of Computing, has taken a leading role in its development. "Because the model is based on a neural network, learning directly from data rather than just replicating the physics, we can calibrate the model and be ready for operation 10,000 times faster," he explained.

Of course, neural networks have already been deployed in many fields, not least in building the large language models that have dominated the headlines since the beginning of 2023. What's distinctive about the Wildfire Surrogate Model is that the neural network is integrated into a wider design that integrates live data inputs, physics modelling and interventions by operators, creating a supremely adaptable AI partner that can support decision-making in an uncertain, high-stakes and rapidly evolving scenario.

Machine learning + data assimilation = data learning

It isn't enough to simply hoover up whatever data is available. In the complex, evolving situations implied by natural hazards, no single source is authoritative. So the model also integrates data science tools that assess inputs: weighting individual satellites by age and location, and whether their view is impeded by cloud cover, for example. And it continues to assess and modify these inputs as the situation unfolds. As Dr Arcucci explains: "One of the fundamental attributes of our brains is that they're adaptive: they change their model of the world based on new information. For example, we're about to cross a road we've identified as safe, when a car suddenly appears – so we don't step out. Our model needs to assess and integrate new developments with the same speed."

That includes retrospectively qualifying its own predictions as the situation unfolds to support or contradict them. Just as certain deployments of large language models like GPT-4 can learn from their interactions with users, so the Wildfire Surrogate Model evolves every time it's deployed.

Dr Arcucci sees this dynamic approach as fundamental to properly deploying neural networks more broadly. She describes it as 'data learning': machine learning supported by constant data assimilation.

"The advantage is that neural networks themselves facilitate this approach," she says. "In the past, the temptation was to simplify data in order to assimilate it – which always entails losing information. Now, you don't have to do that: the network can sort the data itself."

From the next hour to the next century

Developed in partnership with UCLA and the Technical University of Crete, the Wildfire Surrogate Model has already been running for nearly two years of tests; Dr Cheng hopes that it will be ready for operational use within five years. Looking further ahead, the same team are in the early stages of developing a global wildfire forecasting model with the UK Met Office. But longer timescales bring their own challenges – an issue that's even more clearly illustrated by current challenges in climate modelling.

AI tools make climate models usable in real-life decision-making

Climate risk modelling is an excellent example of a field in which the greater processing speed and data volume enabled by AI tools can make research insight usable in an operational context.

"It's hard to find teams who can bring together state-of- the-art science and financial engineering, which is what we really need to plan for the risks implied by unfolding climate conditions."

As Associate Professor of Actuarial Finance at Imperial, **Dr Enrico Biffis** works on everything from sustainable forestry investment to innovative financial products in developing economies; as well as advising the World Bank and IMF; and he approaches all of this with a focus on real-world applications, where the risks studied place lives and large sums of money at stake.

"Climate risk modelling is a hugely challenging, multi-disciplinary space," he explains. "It's hard to find teams who can bring together state-of-the-art science and financial engineering, which is what we really need to plan for the risks implied by unfolding climate conditions."

Dr Enrico Biffis, Associate Professor of Actuarial Finance at Imperial

Photo: Thomas Angus

This integration of skills is precisely what he's pursuing with colleagues at Imperial's **Grantham Institute – Climate Change and the Environment**, as well as Singapore Green Finance Centre, as they work on improving the resolution of climate models to offer high-detail predictions that risk managers can use to identify a level of variation in future hazards across specific regions.

Climate model resolution presents an ongoing challenge

The first challenge is to calculate risk in enough detail to support real-life planning. **Professor Ralf Toumi**, co-director of the Grantham Institute for Climate Change and the Environment, outlines the computational challenge in stark terms: 'The challenge in climate modelling is to build projections with sufficient resolution to be useful, without demanding a level of computing power and time that makes the process impractical. Statistical methods, though much less intensive, don't offer an acceptable substitute for dynamical modelling by themselves; but nor can we take a process that already involves performing billions of calculations over several weeks, and expand it by orders of magnitude.'

To understand how AI-assisted processes can help escape this trap, it's worth first looking at the modelling process in more detail.

CMIP6 – the Coupled Model Intercomparison Project (Phase 6) – brings together sets of results (or 'runs') from around 100 different global climate models, produced to support the sixth International Panel on Climate Change (IPCC) assessment report in 2021. The resulting scenarios describe the possible evolution of the climate for the rest of the century, and underpin current research.

Most of the CMIP6 data comes from General Circulation Models (GCMs). These work by dividing the atmosphere, ocean and Earth's surface into a 3-dimensional grid, and mathematically calculating the actions of various forces and process within each unit given a set of starting points, as well as how it then affects those around it; and then repeating this at regular intervals or 'time steps'. The process is called dynamical modelling.

Factors modelled include surface pressure, wind velocity, water vapour, albedo and cloud cover – and, crucially, the effects of carbon emissions and other human-driven factors. The outputs offer a probability range for climate conditions in each grid unit at each time step: precipitation levels, wind speeds, average temperature and so on.

Researchers tend to use a set of 'pathways' to compare the potential results of different emissions scenarios. The RCP8.5 pathway implies 'business-as-usual' emissions that will lead to the Earth absorbing of 8.5 Watts of energy per square metre by 2100; there are intermediate scenarios at 6 and 4.5 Watts; and RCP2.6 is the best case, which still implies an energy increase of

2.6 Watts per square metre. Each model will output its calculations on many factors – including warming – according to the conditions of the chosen pathway.

The problem is, the resolution of the resulting data is too coarse to inform much real-world decision making. At their most detailed, the CMIP6 generation of models break the Earth's surface into grid squares 100 km wide; an actuary pricing insurance against natural hazards might need to be 10 or 100 times more precise.

This is because the level of geospatial detail a model can provide is limited by the amount of computational power available. Working on the basis of half-hour intervals (a common value) for a period of 100 years implies more than 1.5 million calculation steps per grid square, or 75 billion for the entire globe at a resolution of 100km x 100km.

Given the further complexity of the calculations entailed in each individual step, it's not surprising that climate research is a famously computing-intensive field. A single run – say, modelling the global effects of the RCP8.5 ('business as usual') emissions scenario until 2100 – can take weeks to complete, even on the specialised computing clusters available to leading institutions like Imperial.

Professor Ralf Toumi, co-director of the Grantham Institute – Climate Change and the Environment

Photo:Ralf Toumi

Enlarging this demand by several orders of magnitude simply isn't practical. Instead, Dr Biffis and his collaborators have developed machine-learning techniques to downscale (or 'zoom in') CMIP6 data to a usable resolution – that is, to work with squares as small as 1km by 1km – without exponentially increasing the computational power required.

Machine learning enables precise and efficient downscaling

The key insight echoes that of the Wildfire Surrogate Model above: that it's possible to design a machine-learning model to harness the efficiency of a statistical approach – simply seeking patterns in data and then extrapolating them – without sacrificing the greater accuracy and explainability of dynamical modelling, which uses mathematics to replicate and 'rerun' the actual physical processes that generate the data in the first place.

Concentrating on the Western Pacific, the team trained a series of machine-learning tools to understand patterns in the variation between real historical climate data and 'simulated observations' from CMIP6 models – the results given when the models were asked to retrospectively predict historical climate conditions.

These two sets of data – one real and observed, the other generated by mathematically replicating physical processes – both cover the same time period, and so they can be compared, using a process called quantile mapping. This involves comparing the distribution patterns within real and modelled datasets, and mapping the former onto the latter.

Trained to recognise these patterns of difference in minute detail, Dr Biffis' machine-learning tools were then used to bias-correct the future scenarios projected by the climate models, based on their deviations from the historical data. In other words, the predictions were recalibrated based on their likely pattern of error – and because the real observations exist at a much higher resolution, the projected data could also be produced with much greater detail – in units measuring 1km×1km.

The results were finely-detailed sets of geospatial data, offering nuanced outlooks on climate variables across Indonesia, Malaysia, Thailand and the Philippines. And further analysis provided secondary insights into heightened tail risks: the chance of extreme events, a vital consideration for insurers, as well as broader changes in average.

It's important to note, though, that the machine learning/statistical approach doesn't replace dynamical modelling. Any missing physics in the coarser resolution global model can not be simply fixed by this statistical downscaling. However, the basis of the process is still the statistical replication of physical processes within the CMIP6 models.

Instead, what the researchers have done is harness the interpretative power of AI to extrapolate from this information with enough confidence to make real-world decisions – raising the precision of the data without requiring prohibitive computing power.

Computer vision can identify patterns in noisy and complex data

At the other end of the process, AI can also be harnessed to improve the data providing the starting points for such models, to make their resulting predictions more reliable.

Let's return to Professor Ralf Toumi, whose own research offers a case in point. A physicist focussed on understanding tropical cyclones, he's led the development of the Imperial College Storm Model (IRIS) : a parametric statistical model (that is, one whose results are constrained by set parameters – in this case, the relevant laws of physics) that predicts the probability of cyclone landfall across the planet.

It does this by generating millions of synthetic storms, in order to calculate the probability of a given magnitude or landfall location over periods as long as a thousand years. This provides the basis for a particular cyclone to be described as, say, a 1-in-200 year event (in other words, one with a 0.5% probability of happening in any given year) – an important threshold for insurance capital allocation (the process by which insurers determine how much capital to hold in relation to different risk exposures, in order to meet their obligations in the case of a claim).

IRIS is built to be constrained by the actual thermodynamics of the atmosphere – particularly in the maximum intensity of the storms predicted. It can use predicted data from GCMS as the basis for this, but is substantially more accurate than them within its own specialism.

That accuracy is partly underpinned by a key insight about usability, as Professor Toumi explains: "In a research context, people tend to be interested in modelling the whole life-cycle of a storm. That can entail a lot of margin for error. But what's actually relevant for risk managers, in insurance and so on, is where it makes landfall, and what its intensity will be when it gets there. If you pick up the storm at its peak, and only model its decay, you can be much more precise."

So where does AI come into this? In highly specific, targeted applications, its role is to help the model understand and apply key relationships and variables. For example, one of the first steps when estimating the intensity of a tropical cyclone and forecasting its track is locating its centre. This is as relevant to a model like IRIS as it is to understanding a real storm as it unfolds: providing accurate initial positions and characteristics generates more precise simulations. Satellite cloud imagery can be difficult to interpret, however.

Professor Toumi has trained a series of convolutional neural networks (CNNs) to identify storm centres from the long-wave infra-red information in sequences of satellite photographs. CNNs (a sub-type of neural network) are often used in image recognition tasks, because rather than working on a pixel-by-pixel basis, they break an image into components and classify these – rather in the same way many cognitive scientists claim our own brains interpret incoming sensory information. The result is a process that's better at identifying features relevant to classification, and doesn't get swamped with the sheer volume of data in an image.

However, he's keen to emphasise that the IRIS project only deploys AI in this kind of specific, limited way. "It's important that the model itself remains transparent and accessible, and that we can explain its results in direct reference to specific physical phenomena – rather than just statistically." This is in marked difference to the 'black-box' proprietary models used by many insurers – and arguably more useful in the context of more specialised risk management strategies, such as parametric insurance. These policies, triggering a fixed payout at a threshold event, rely on the ability to reliably model the probability of such specific thresholds being crossed.

Explainability, accountability, efficiency

Explainability is a key focus of AI research right now. The projects explored here demonstrate that, as Dr Arcucci puts it, "the word means different things in different contexts." She goes on: "For us, explainability means replicability." In other words, if the Wildfire Surrogate Model is consistent in its predictions – and across the high volume of iterations it's capable of running in a short period of time – it can be relied upon by responders in the heat of the moment.

That said, Dr Sibo Cheng is currently working on causal explainability within the model's neural network. "We're using tools to determine the most important neuron in the gradient descent," he explains. "It's like backpropagation: by following the route of the data through the network as it's processed, we can identify the most important node. And as each node is associated with a particular phenomenon or process, we can then identify that process as salient.'

In the case of climate risk, these approaches are still grounded in robust and continuously evolving dynamical models of the physics. What the AI tools do in each case is make the model more usable: in determining better data to feed it in the first place, or in pulling out enough detail to make reliable real-world judgements. Dr Enrico Biffis has also used a similar approach to determine which climate model to deploy in the first place: using the trained machine learning tools to assess the statistical basis of different models and their appropriateness to a given use case. When it takes days to run a model, this pre-judgement offers valuable efficiencies.

The result of applying these techniques is that planners and decision-makers can make judgments with more confidence in the face of an increasingly uncertain future, in which exponential change becomes the norm. Faced by ever more complex situations, the challenge is to harness AI's ability to find the patterns and insights we need to keep moving forwards. Across Imperial's community of researchers, there's plenty of evidence that we can meet it.

Contact Imperial Business Partners

If you would like to know more about Imperial Business Partners and our futures consultancy, Imperial Tech Foresight, we'd love to begin a conversation with you.

ibp@imperial.ac.uk | imperial.ac.uk/ibp

Imperial Business Partners