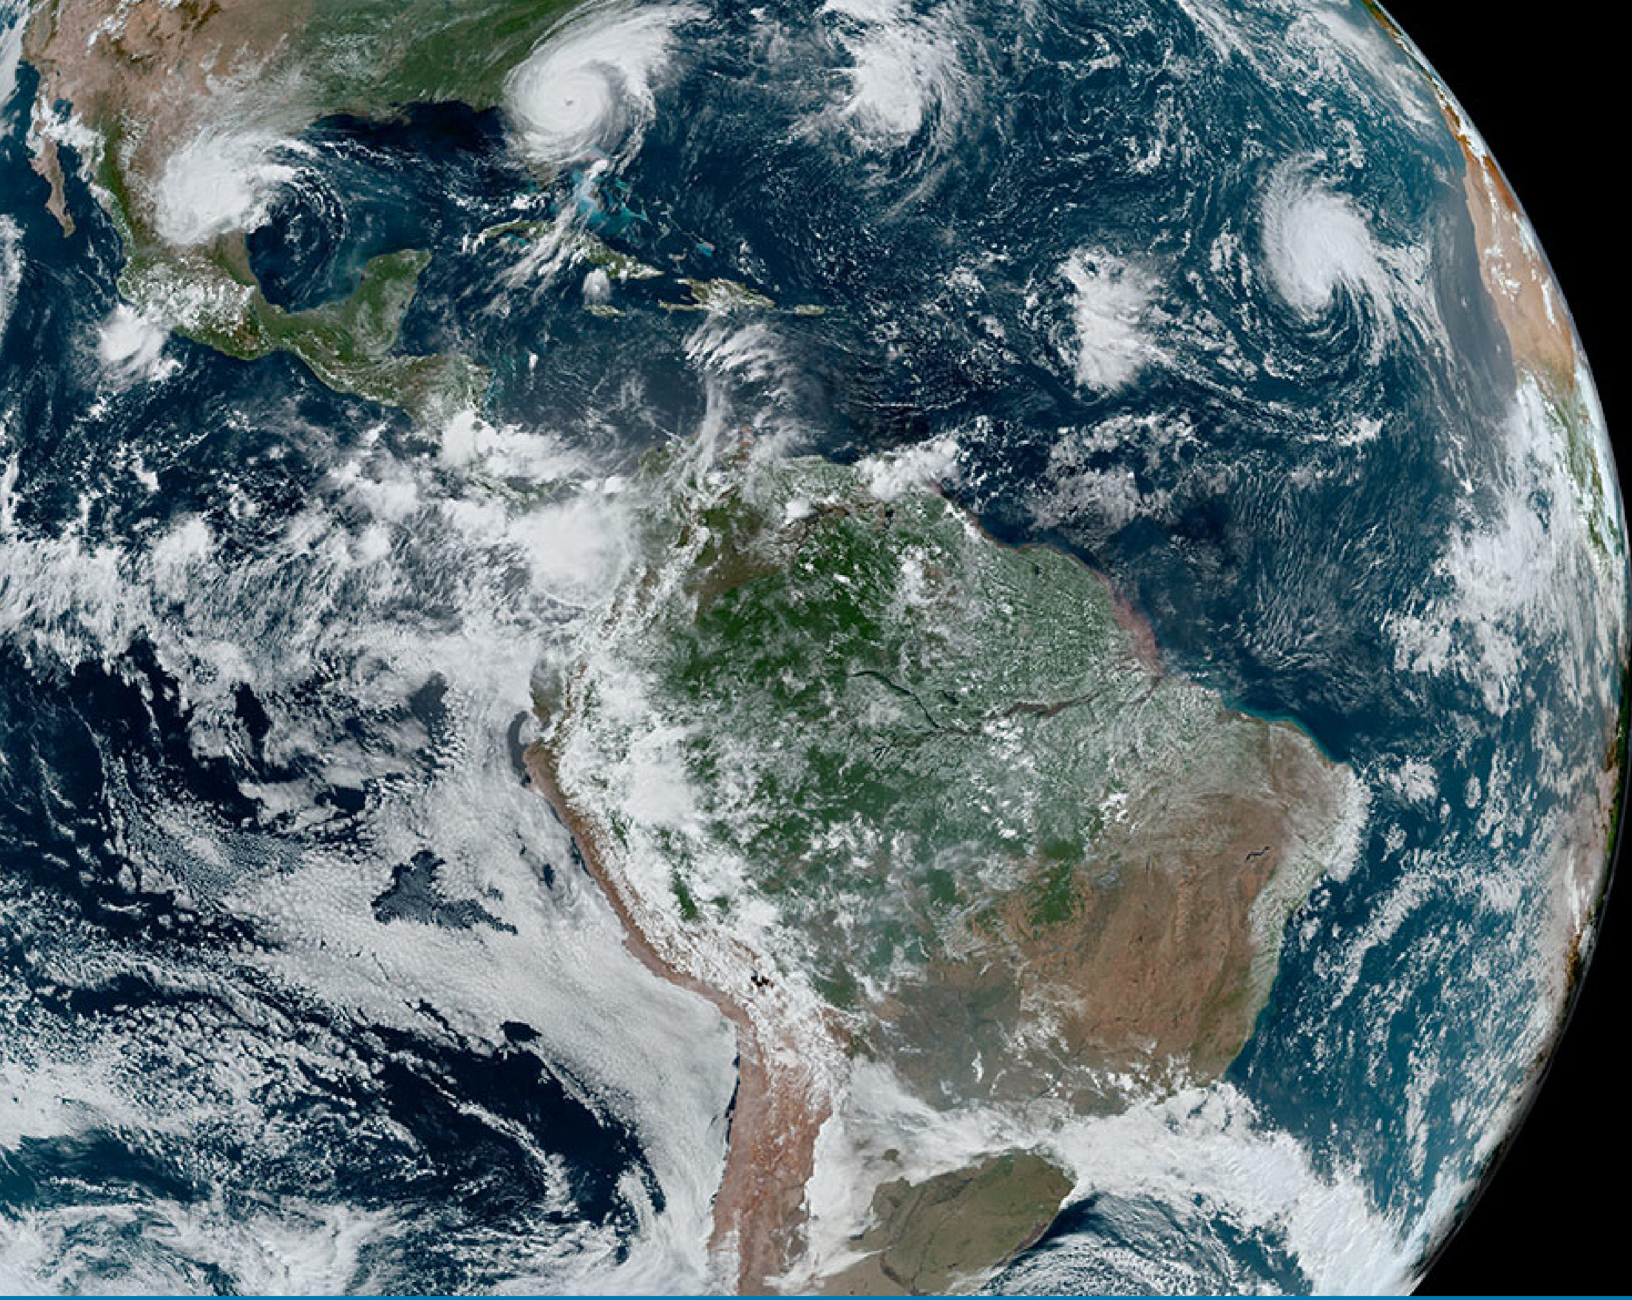

Reimagining the Future of Simulating Atmospheric Perils

A New Framework for Modeling Weather and Climate Extremes

Verisk

Contact Information

If you have any questions regarding this document, contact:

AIR Worldwide Corporation
Lafayette City Center
2 Avenue de Lafayette
Boston, MA 02111
USA

Tel: (617) 267-6645
Fax: (617) 267-8284

For more information on our Boston headquarters and additional offices in North America, Europe, and Asia, visit **air-worldwide.com/About-AIR/Offices/**

Table of Contents

Foreword

Bill Churney, President, AIR Worldwide

Climate change is a definitive challenge of our times, with strongly nonlinear, cascading impacts and many surprises in store. The Paris Agreement, a pledge signed by almost all nations to take action to limit global warming to less than 2˚C compared to pre-industrial levels, identifies insurance as a key piece of the resilience solution. The industry is in a unique position to play a leadership role in assessing, communicating, and reducing climate change risk and developing new insurance products to meet the needs of the evolving market. At AIR and across the wider Verisk organization, we continue to make significant investments in our modeling capabilities and our technology infrastructure to be effective, future-ready partners in this crucial endeavor.

To that end, AIR has worked with partners in the larger modeling community to create a new framework for modeling atmospheric perils under a future climate. Between June 2020 and April 2021, we published a series of articles that explain AIR's new framework, which is currently under development. This series has been compiled in this white paper and modified slightly for readability and currency. In it we discuss the evolution of and motivations for our catastrophe modeling techniques from our pioneering beginnings in 1987 through to the present day and beyond, as well as their practical applications.

Traditional catastrophe modeling—a discipline built on making sense of historical data—has had to evolve to create views of extreme event risk for today's climate and must continue to evolve to create views of this risk for the climate of tomorrow. The business benefits are substantial, starting with a view of risk for regional perils better rooted in science and a new quantification for diversification benefits, leading to better allocation of risk capital. In short, for those in the insurance industry, you can perform all the key functions you do today, but with greater confidence in how you will manage the uncertainties due to a changing climate."

While we continue to adapt today's models to offer pragmatic solutions to questions like "how will hurricane tracks change?" and "will storms move faster or slower?" our aim is to address the physics of atmospheric phenomena, especially at fine scales that are of interest to catastrophe modelers, in a more comprehensive way.

Part I: A Reckoning and a New Approach to Modeling Risk

Jayanta Guin, Ph.D., Chief Research Officer, AIR Worldwide

Many parallels have been drawn between the unfolding crises of COVID-19 and climate change. Of the lessons that can be learned from both, perhaps the most important is that science matters. The countries that have weathered the pandemic best—Singapore, Hong Kong, and South Korea, for example—did so because they understood the threat early on, having battled the SARS and MERS epidemics of 2003 and 2012, respectively. The scientific investments these countries made in the aftermath of those outbreaks better prepared them to deal with COVID-19. Had the rest of the world recognized the inevitability of a global pandemic and invested tens of billions of dollars at the right time, tens of *trillions* of dollars in economic damage might have been saved, not to mention the enormous societal damage that COVID-19 has inflicted.

For some years, we've been at a tipping point where scientific investments in mitigating the impacts of climate change can still make a material difference. We know that the cost will be orders of magnitude higher the longer we delay meaningful action. I'm sure that I am not alone in hoping that our experience with COVID-19 will inspire the world to finally reckon with the implications of inaction on climate change and spur the necessary investments and fundamental policy changes at the global, sovereign, and societal levels.

Here at AIR, we also feel a tipping point and are taking bold steps to fundamentally change the way we quantify insurance and other financial risks due to climate change. This part will describe our motivation and approach.

A Brief History of Cat Modeling and Climate Change

In the mid-1980s, AIR introduced a fundamental change in the way extreme event risk was quantified: the first stochastic hurricane model for the insurance industry, which initially failed to garner much attention. Hurricane activity had been below normal for the decade prior and the model's suggestion that the industry could see losses several times larger than had ever been seen before was met with skepticism. Hurricane Hugo, which in 1989 produced the largest hurricane loss to date, opened some minds, but it wasn't until **Hurricane Andrew** in 1992 that the industry fully embraced the new technology.

At the time, few outside academia and government were concerned about global warming. Early generation cat models produced robust results by relying on decades of historical observation data (augmented by scientific expertise), assuming a stationary climate. The first-generation models, which were a novelty at that time, were used to create the very early sensitivity studies of changes in the frequency and severity of hurricanes, but it's fair to say that it was not done in earnest. The first real concerns about climate change impacts on hurricanes were expressed after the record-breaking 2005 Atlantic hurricane season when the National Hurricane Center had to use Greek letters for six storm names. (The next time Greek letters were used to name storms was during the **historic 2020 Atlantic hurricane season**, and we had to reach even deeper into the alphabet to name eight storms.) The year 2005 was also when Hurricane Katrina devastated New Orleans, and Hurricane Rita not only achieved the lowest central pressure on record in the Gulf but also the largest radius of maximum winds. In response, catastrophe modelers introduced the first "climate conditioned" catalogs a year later.

Still, in the mid-2000s climate change largely remained an afterthought to the insurance and cat modeling industries. The 10 years with no single Florida hurricane landfall that followed—often referred to as the hurricane "drought"—helped push the issue to the background yet again (despite the possibility that climate change may have been responsible for the drought). For many years, the interest emanating from the insurance industry was sporadic and mostly reactive to individual events.

Today, large-loss weather events are almost guaranteed to produce headlines attributing them to climate change. And thanks to the relatively new science of event attribution, there is growing justification in doing so. On such occasions we can see with our own eyes the effects of ever-increasing greenhouse gases. Memories are short, however, and the effects remain largely invisible—except, perhaps, for those areas of the coast that are experiencing more and more frequent sunny day flooding. For most of us, climate change is a long, drawn-out catastrophe unfolding just beyond our line of sight.

This presents a fundamental challenge to maintaining focus on the issue. But all of us, as stakeholders, must overcome the challenge and recognize that we are at another inflection point or "Hurricane Andrew moment," although a far more momentous one. Whether climate change is truly an "existential" threat may still be debated, but its costs, which we are already beginning to experience and which will only increase over time, can no longer be ignored. It is no longer acceptable to think about climate change only after an extreme event occurs. And we cannot allow market pricing cycles, like the soft market that, in part, resulted from the 10-year hurricane drought, to lower our guard or commitment.

What Needs to Be Done?

We must be guided by the science in our modeling of climate change risk, not rhetoric or headlines. We must recognize the inherent uncertainties in projections of future climate states and navigate them using rigorous analytics. We must thoughtfully explore the known knowns and the known unknowns, and we must at least imagine and speculate about the unknown unknowns.

It is time for us, as model developers, to go beyond a piecemeal approach to addressing climate change, beyond incremental updates or extensions to existing models. As model developers, we must take the best-of-breed science coming out of academia and leading research institutions and translate it into fit-for-*business*-purpose climate change models and analytics. In a very real sense, we should be thinking of climate change as a new peril, and the models must be capable of answering different kinds of questions.

What Are the Questions that Only a New Breed of Climate Models Can Answer?

Since the introduction of the first cat model, we've been trying to answer questions such as, "What is the probability of a Category 4 hurricane making landfall in Texas?" Today, the relevant questions are much more complex, for example: "What is the probability that a Category 4 hurricane will make landfall in Texas, stall over Houston, and drop more than 50 inches of rain?"

In fact, AIR's existing physical-statistical-hybrid approach can get us quite close to an answer. What our existing approach cannot answer with confidence is, "What is the probability of that same scenario happening over Mobile, Alabama?" Similarly, we might speculate, "We saw Category 5 Hurricane Dorian stall over the Bahamas; what is the probability of that happening over southeast Florida?"

The only way we can begin to answer these questions is by modeling the physics that give rise to such occurrences. We must better understand how small-scale features in the atmosphere can have disproportionate impacts on large-scale planetary features (and vice versa) and how these non-linearities and teleconnections between scales and across distances drive weather extremes. Although the quality and detail of **reanalysis data sets** for the last 40 years has improved by leaps and bounds (the latest holds more than 500TB of data), they can only tell us what has happened historically. They cannot answer the question of where and how frequently a break in a planetary wave will set up a large stationary high, such as the one that caused Hurricane Sandy to make its notorious (and anomalous) westward turn into northern New Jersey; they cannot tell us the frequency of anomalous jet-stream behavior that allows the **polar vortex** to split and sag southward, bringing frigid temperatures to North America for long stretches; or omega blocks that can bring tropical temperatures to Europe and Greenland, as one did in the summer of 2019. Even if we spatially perturb the reanalysis data to create new scenarios, we do so with many unknowns, leaving us with lower levels of confidence in (probably biased) results that no longer consistently and coherently represent the dynamical nature of the atmosphere. And we are unlikely to produce simulated events that present surprises for us—surprises that we know a changing climate will bring.

Why Is the Problem So Difficult to Solve?

At its simplest, the answer to this question is clear: The atmosphere is chaotic. Weather, which refers to short-term atmospheric conditions experienced at a location over the course of hours or days, is highly variable. The reliability of weather forecasts falls off after about a week, if that long. Climate, on the other hand, refers to the statistics of weather over decades, often over 30 or 40 years. While the climate tends to change quite slowly, we do experience shorter-term fluctuations; El Niño Southern Oscillation (ENSO) and the North Atlantic Oscillation (NAO) are familiar examples of features in our ocean-atmosphere system that drive variability. When discussing "climate variability," we're describing natural (i.e., not man-made) processes that affect the atmosphere. When we introduce anthropogenic climate change (caused by greenhouse gases), it becomes quite challenging to distinguish its signal from climate's natural variability, particularly when it comes to short-term weather phenomena such as individual storms.

The task for the catastrophe model developer is to build large catalogs representing ensembles of future climate states at different timescales. In theory, we might create such an ensemble using global circulation models (GCMs) which have become quite powerful. But while some very recent GCMs are capable of running at high enough resolution to explicitly simulate small-scale features such as hurricanes, the computational cost of running them at that resolution long enough to generate tens of thousands of years of simulated hurricane activity remains impractical. Furthermore, no current GCM attempts to capture all the smaller-scale atmospheric processes and their interactions that give rise to the full range of extreme weather events. Those processes that are not resolved are parameterized, which can introduce additional bias. Careful analyses of GCM output in the recent climate reveal their limitations in accurately representing the statistics of some of the larger-scale dynamics in the atmosphere, which we know often contribute to the conditions leading to the occurrence of an extreme event.

But this leads to another question: If we successfully build new models suitable for different timescales, how do we validate them? We know that history is unlikely to be representative of the future, so what does validation of a future climate state really mean? In fact, validation will necessarily take on a new meaning. The key is to break the problem into smaller constituent parts and develop an approach whereby we can validate the recent past and then employ those features of future model projections in which we have more confidence. For example, the impact of climate change on sea level rise, **Arctic amplification**, and temperature and precipitation patterns are better understood than the impacts of climate change on tornado activity. But if we can better understand the atmospheric conditions that drive tornado activity, we will have more confidence in what the models tell us about where and how frequently those conditions might arise in the future.

At AIR, we believe we've found a solution, in which transparency and staying true to the science are key; we cannot allow ourselves to read more into the data than is there or overextend the state of science. We should not be afraid of getting something wrong, but we should be prepared to incorporate new knowledge as it becomes available and update our view of the risk accordingly.

So What Is AIR's Solution?

We see an opportunity to blend our traditional hybrid, physical, and statistical approaches with a new set of tools that come from the world of artificial intelligence—specifically, machine learning. Our approach, which represents the efforts not only of AIR scientists but of partnerships with research institutions at MIT in the U.S., Magdeburg University in Germany, and the University of Utrecht in The Netherlands, combines a novel approach of de-biasing large-scale features in a computationally fast GCM, with analysis of fine-scale features from historical data to learn the "rules" of atmospheric behavior that produce weather extremes. Million-year catalogs are suddenly possible—*global* catalogs that capture all types of dependencies, from global teleconnections to local correlations across all weather-related perils and across all regions.

The result will be a new framework for climate risk modeling, one that is deeply rooted in high quality data and deep domain knowledge of weather and climate physics. The framework will allow us to answer not only today's new climate questions, but also tomorrow's. The business benefits are substantial, starting with a view of risk for regional perils better rooted in science and a new quantification for diversification benefits, leading to better allocation of risk capital. In short, for those in the insurance industry, you can perform all the key functions you do today, but with greater confidence in how you will manage the uncertainties due to a changing climate.

AIR undertook this important project starting in 2018, but there is still much to do on this multi-year journey—a journey that requires patience, tenacity, and commitment. This paper's topics range from the fundamentals of climate and climate variability, to what climate models can and cannot do, to the latest thinking on the contributions of machine learning techniques, to what insights we might gain into the potential impacts of climate change on the locations, frequencies, and intensities of extreme events around the globe.

Also critical is *your* engagement and support to drive innovation in climate risk analytics, which the industry will need for the future. The modeling approach represents a long-term, sustainable strategy for managing climate change risk for the coming decades. In the meantime, we are executing on shorter-term strategies that include increasing our efforts on how we evaluate the impact of climate change on each peril region, development of climate sensitivity event-ensembles, and developing capability in our products so that you can do what-if scenarios and establish your own view of climate risk. Now that this part has provided the motivation for our investment, the next parts will provide the details.

Part II: The Importance of Planetary-Scale Motions for Modeling Local Weather Extremes

Peter Sousounis, Ph.D., VP and Director of Climate Change Research, AIR Worldwide

In Part I, we discussed the importance of being able to model not only the physics of a weather system itself, but the physics of the larger, planetary-scale circulations that operate on longer time scales—the climate dynamics. For example, to better model the risk of a Hurricane Harvey-like event occurring in Miami rather than Houston, it is important to also be able to model the physics in a place such as, say, Manhattan.

The distinction between weather and climate has been stated in the following ways: Climate is what you expect, weather is what you get; weather is your mood, climate is your personality; climate lasts and weather doesn't. They all convey the same message—that climate is an amalgamation of weather over a long period of time. But climate is more than that. Understanding what climate really is from an atmospheric dynamics perspective is critical to improving catastrophe models. That is what this part is about, how climate is inherently *the cause of* weather and weather variability, and the real driver behind weather extremes—not the *end result* of weather. At the same time, it is important to understand that climate and weather are inextricably connected.

Climate's Large- and Small-Scale Atmospheric Motions and the Weather They Generate

The climate that any given region on Earth experiences is determined largely by its latitude. But latitude means more than just how far away from the equator a location is and how directly the sun's rays reach the Earth's surface. Latitude also influences whether prevailing surface winds blow from the east (easterlies) or from the west (westerlies) and how strong they are. At *midlatitudes*, which are typically defined by the range of 30° to 60° north or south of the equator, prevailing winds are westerly (see F). These westerlies are the result of southerly winds in the Northern Hemisphere associated with the **Ferrel cell**—the middle cell of three pole-to-equator cells and responsible for equalizing heat imbalances—being deflected eastward by the Earth's rotation. In other words, in the Ferrel cell, air flows poleward and eastward near the surface and equatorward and westward at higher altitudes; this movement is the reverse of the airflow in the **Hadley cell**, another of the pole-to-equator cells.

At upper levels, these westerlies are guided by the polar jet stream. The jet stream and the six to eight large-scale waves along it owe their existence to the pole-to-equator temperature difference, or *gradient*, as well as to the Earth's rotation. When the temperature gradient is large, the jet stream is faster and straighter; when the gradient is smaller, the jet stream slows and begins to meander, becoming wavier. And not only does the air move *through* the wavy pattern, the waves *themselves* move.

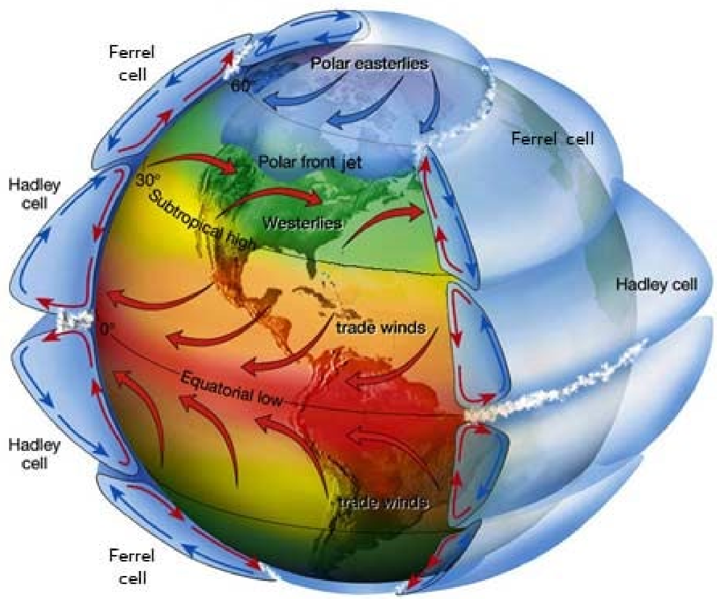

Figure 1. Locations and extents of meridional cells and polar front (jet) mentioned in text. Adapted from futurelearn.com**.**

Rossby Waves

These waves are known as Rossby waves (see box) after the meteorologist Carl Gustaf Rossby who discovered them in the 1930s:

- Very long waves the size of continents can move slowly eastward or westward at mid-latitudes—or remain stationary; these long waves can therefore influence weather patterns that last an entire season or longer
- Shorter waves—with wavelengths the size of low-pressure systems such as winter storms—do move from *west to east*
- In the tropics, waves of all wavelengths move *east to west*

The movement of these waves alone accounts for much of the weather variability for a given region. Storms approach; clouds, wind, and precipitation develop; skies clear. It's a typical winter weather pattern in the midlatitudes.

But the waves don't move in isolation; they interact, non-linearly at times and through positive feedback mechanisms. *Non-linearity* means that as waves phase together (e.g., the ridges or troughs of the waves align) the resulting amplitude can be greater than the sum of the individual wave amplitudes. Furthermore, what happens at the Earth's surface can greatly influence the motions above and vice-versa.

Rossby Waves

Any given parcel of air has two kinds of spin or vorticity. One is the kind we see, like air rotating counter-clockwise in the Northern Hemisphere around a hurricane, called relative vorticity, and planetary spin, which is like potential spin that we can't see. Planetary vorticity is zero at the equator and largest at the poles. As air moves, it has to conserve the total or absolute vorticity it started with. As air moves northward from Position 1, its relative vorticity decreases as it gains planetary vorticity. This causes the air to begin to rotate clockwise more, which results in the parcel eventually moving southward at Position 2 as it is moving eastward. As it overshoots the latitude it started at, the parcel now loses planetary vorticity but gains relative vorticity. This eventually causes the parcel to move back north again at Position 3. Thus, a wavelike motion is traced out by the parcel—forming a Rossby wave (green ribbon). The wave itself can move slowly in either direction or remain stationary, depending on its wavelength and how fast wind moves through the wave.

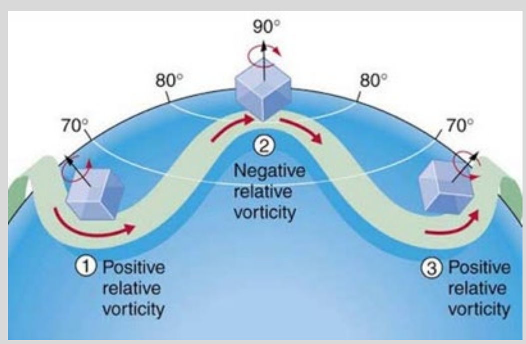

Figure 2. Vorticity and Rossby waves (Adapted from homework.uoregon.edu)

Atmospheric Blocks

Sometimes, these interactions can cause an *atmospheric block*—one that prevents weather systems from proceeding along their "preferred" path. An omega block is one example. It is so called because the flow of air around the resulting Low-High-Low or L-H-L pressure pattern resembles the Greek letter (capital) omega Ω. The rightmost panel in Figure 3 illustrates such a block (designated by H) preventing a low-pressure system (designated by L) from moving directly eastward. Instead, the low is forced to go northward around the block.

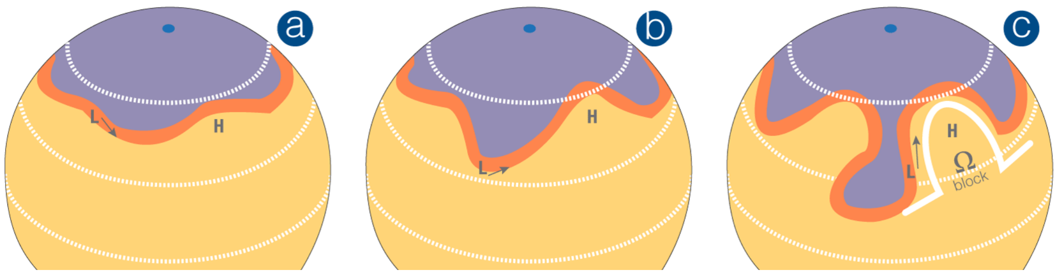

Figure 3. An example of Rossby wave evolution over a three-day period. Long Rossby waves amplify as they migrate very slowly eastward (left panel); a smaller wave of low pressure (L) moves more quickly through the Rossby wave configuration (middle panel); an example of what a block looks like and what it does to the path of approaching weather systems (right panel). (Adapted from wikipedia.org**.)**

Blocks can develop quickly, within a matter of days, but can last for weeks to months. Although they typically form in spring and fall as residual cold or warm air is injected equator- or poleward, respectively, blocks are most impactful in winter and summer because temperatures are more extreme then. Blocks can also affect hurricane tracks; they were influential for Hurricane Sandy's famous left hook into New Jersey in 2012 and for Hurricane Harvey's stall over Houston in 2017.

Extreme Events

The formation and dissipation of blocks are difficult to forecast and they can be self-sustaining. A good example is what happened over Australia in late 2019, just ahead of the **2019-20 bushfires**, which destroyed 10 million hectares (~25 million acres) and destroyed more than 2,600 homes, mainly in Australia Capital Territory, New South Wales, Queensland, South Australia, Tasmania, and Victoria. Earlier in the summer of 2019, one of the strongest positive phases of the **Indian Ocean Dipole** developed to cause a drought in western Australia. Less rain meant fewer clouds, enabling more sunlight to heat the ground. As surface temperatures rose, the wind pattern aloft responded accordingly—causing a southward bulge in the steering currents so that any rain systems approaching from the west were diverted to the south of the continent. That allowed yet more sunlight to reach the ground, perpetuating the pattern. An intense sudden stratospheric warming episode further exacerbated atmospheric conditions for bushfire.

What happened over Australia during their springtime is an example of an extreme climate pattern: complex and large in scale, it influenced weather and weather-related phenomena for months. To accurately model extreme weather phenomena, such as the 2019-20 Australia bushfires, one has to be able to model extreme weather patterns and the dynamics that cause them.

Another way blocks, or stagnant weather patterns, can form is when Rossby waves get trapped, stall, and amplify. Some recent studies have suggested a **link to climate change**, that is, as the pole-to-equator temperature difference decreases, the trapping and resonance of waves will occur more frequently. Not only drought, as in the case of the bushfires, but extreme flooding can also occur, as wet weather systems follow one after another over the same region, like railroad cars on a train track. When these tracks set up from southwest to northeast across oceans, *atmospheric rivers* can develop to transport up to 15 times more moisture than the Mississippi River can—to the west portions of continents. These high-amplitude, nearly stationary (in time) patterns have been used to explain the devastating 2003 heat wave in Europe that claimed 70,000 lives, as well as more recent events such as the 2010 Pakistan flood/Russian heat wave, the 2011 Texas drought, the **2013 European floods**, the 2015 California wildfires, and the 2016 Alberta wildfires. A prolonged atmospheric river event in mid-December 2010 fueled strong winter storms that battered the U.S. West Coast with up to two feet of rain and provided 75% of the Sierra Nevada's annual snowpack before winter even started. These are all examples of extreme climate patterns—extreme in the sense that, historically, they haven't happened very often and they resulted in significant consequences. Global climate models suggest, however, that the sort of resonance, or increased amplitude, that gives rise to such events will happen 50% more often.

While altered (e.g., blocked) flow patterns can lead to extreme weather events reaching areas that don't normally experience them, the intensity of the weather event can also reach unprecedented levels. As previously noted, as waves interact, they can resonate and grow in amplitude. Thus, a short wave moving through a long wave can grow in intensity dramatically, or explosively, especially as clouds are forming and releasing latent heat. The term *explosive deepening* applies to a low-pressure system whose central pressure decreases by as much as 24 mb in 24 hours. Such **"bomb" cyclones** are not new (the term was coined back in the 1970s), but they may be occurring more frequently because climate change is allowing more moisture in the atmosphere and more latent heating to occur.

Ocean-Atmosphere Interactions Add Complexity: The El Niño/La Niña Southern Oscillation (ENSO) Example

So far, we've discussed how atmospheric motions on different scales can interact to yield extreme weather. The long wave patterns themselves can be influenced by other dynamics that involve both the atmosphere and the ocean. Some of these dynamics repeat themselves with some frequency, which has led to the identification of a large number of climate signals, or climate factors. They are so called because they influence atmospheric motions on large (continental) scales and may do so either continuously or with some quasi-periodicity. The climate signals can also interact with weather.

The interaction of climate and weather is demonstrated most readily by considering a familiar climate phenomenon. The El Niño Southern Oscillation (ENSO) is an atmospheric-oceanic phenomenon that affects weather worldwide. The Southern Oscillation is the atmospheric part, but the oceanic part (El Niño) is what gets more attention. ENSO manifests on a time scale of three to seven years, meaning that in that period at least one weak El Niño and its counterpart La Niña will probably have affected Earth. The strength of each phase is typically measured by how anomalously warm (El Niño) or cool (La Niña) a region of the central-to-eastern tropical Pacific is over a three-month period.

During an El Niño for example, the warm water that is normally largely confined to the western Pacific shifts eastward, sometimes all the way to the west coast of South America. (The phenomenon got its name, which means "the boy" or "the Christ Child" in English, owing to its appearance being most noticeable at Christmastime.) The eastward transport of warm water implies that the ocean gets involved. It also implies that the normally reliable easterly trade winds, which are what typically hold the warm water in place in the western Pacific, relax or even change direction. The absence of offshore winds along the west coast of Peru during El Niños can halt the fishing industry there; the absence of offshore winds means no upwelling of cooler water, no bait fish, and hence no big fish.

But the water doesn’t just shift eastward on its own. Intense thunderstorm activity (weather) in the western Pacific is the precursor to an El Niño. This intense convection triggers the eastward transport of warm water just below the surface along the equator. Once the warm water reaches Peru, it spreads northward and southward along the west coasts of the Americas. El Niños have significant impacts on tropical cyclone activity, severe thunderstorm activity, and overall precipitation patterns over much of the world.

For example, during an El Niño, hurricane activity is typically reduced over the north Atlantic because of increased vertical wind shear and increased over the eastern and western Pacific owing to a larger expanse of warm ocean water. Because warm water is the fuel for not just tropical cyclones but precipitating weather systems in general, the west coasts of the Americas typically get increased rainfall and even flooding conditions. In contrast, much of Asia as well as eastern Australia can experience drought conditions. Food stocks can decrease and famines start in eastern Africa as crops start to fail, potentially leading to unrest. La Niñas can also bring such spatially correlated weather but in a different way. Atlantic hurricane activity is typically above normal, for example, while the Pacific tropical cyclone seasons are below normal. Heavy rains characterize eastern Asia, especially.

Climate in the Context of the Catastrophe Model

The periodicity of ENSO suggests that an appropriate time scale with which to define climate should allow for at least several such events. Indeed, the standard length over which weather is averaged to reflect climate is 30 years. While this length of time might at first glance suggest the appropriate data vintage for developing catastrophe models, one must consider that other climate factors, such as the Atlantic Multidecadal Oscillation (AMO) and the Pacific Decadal Oscillation (PDO), have even longer time scales. The former has a period of about 44 years, although the last full AMO cycle began in 1962 with a negative phase and the complementing positive phase that started in 1995 has not yet been completed. Despite the "decadal" in its name, the PDO is typically associated with a 50-year cycle. None of these numbers, of course, is as precise as suggested here; there is a natural variability in the periodicity of these climate signals.

Perhaps, then, 60 years is a sufficiently long period with which to define climate? Probably not, considering that climate signals interact. The convective activity in the western Pacific that can trigger ENSO is caused by the Madden Julien Oscillation (MJO), which operates on a time scale of 30 to 60 days. In fact, it’s quite possible that a particular weather regime that a climate factor *combination* might generate naturally has not yet been observed. For a catastrophe model, the appropriate record of historical data will depend on the variable of interest. This also speaks to the reason why the last few years of extremes may not be a predictor of what is to come: short-term trends can be misleading.

Pulling the question of climate time scale in a completely different direction is the existence of anthropogenically induced climate change. The fact that the Earth’s temperature has increased by more than 1° Celsius in just the last 70 years as a result of increased greenhouse gases means we cannot just take data over a very long period to try to account for all possible natural climate interactions without somehow also accounting for climate change because not only is climate itself complicated, the very definition of it is too.

What Can General Circulation Models (GCMs) Do?

The climate's atmospheric and oceanic motions and the extreme weather they can generate—specific combinations of which may have yet to occur—are complex and interconnected. The factoring in of climate change as it evolves makes for a convincing case that building catastrophe models from historical data alone, or even the statistics of that data, may leave gaps that affect a model's ability to appropriately represent extreme weather events, if not in the near future then certainly in the more distant future. How might the climate factors of most importance to the catastrophe modeler, for example, be altered by climate change?

One way to fill such gaps is through the more explicit use of general circulation models (GCMs) to model the actual physics of Earth's climate system and thereby account for outcomes that have not yet been observed—such as a Hurricane Harvey stalling over Miami, or even Manhattan. To see where and how far GCMs can take us, we must examine in some detail what such models can and cannot do, and why, which is what will be discussed in Part III.

Part III: Anatomy of a Climate Model

Henk Dijkstra, Ph.D., Professor of Dynamical Oceanography, Institute for Marine and Atmospheric research, Utrecht; and Director, Centre for Complex Systems Studies, Department of Physics, Utrecht University

Boyko Dodov, Ph.D., Vice President and Director, AIR Worldwide

As discussed in Part II, planetary-scale atmospheric and oceanic motions are the drivers of local weather extremes—the very extremes that result in large financial losses. To understand the dynamics between weather and climate, and how they may be altered as a result of climate change, scientists employ climate models, which numerically simulate planetary circulations and their interactions at various scales.

The first mathematical models of the atmosphere were developed decades ago. In 1904 physicist and meteorologist Vilhelm Bjerknes proposed that the principles of fluid dynamics could be used to predict atmospheric flows and hence the weather. In 1950 the first successful numerical weather prediction was made by a team under the leadership of the mathematician and physicist John von Neumann. The first 24-hour forecast took nearly an entire day to compute.

Today, a range of numerical models is available for the forecasting of weather and climate phenomena, each of which has been developed for specific forecasting or research purposes; they can differ in terms of spatial domain and resolution as well as the time period for which the forecast is valid. Numerical models are based on *primitive equations*, which include terms for the conservation of mass; a form of the **Navier-Stokes** equations governing fluid flow; and thermodynamic terms. Examples of numerical models include numerical weather prediction (NWP) models and global (and regional) general circulation models (GCM).

NWP models are atmospheric models that currently have a horizontal resolution of about 10 km; they may be regional or global in scope. With these models, only simulations over a short period of time (typically 14 days) can be performed. Because of the chaotic nature of atmospheric flows, the forecasting capability decreases dramatically beyond a two-week period. Furthermore, to produce a reasonably accurate forecast, using accurate initial conditions are important. These initial conditions represent a "known state" of the atmosphere and are constructed using observation data available through **data assimilation** techniques.

To account for uncertainty in the initial conditions, many simulations (typically 50) with different initial conditions are performed to create a statistical forecasting product (Figure 3). From this product, the probability of specific extreme events, such as severe thunderstorms, hurricanes, heat waves, and extreme precipitation, can be assessed. The forecasts that result from this process allow people, governments, and organizations to not only plan daily activities but also prepare for disasters, up to 14 days in advance.

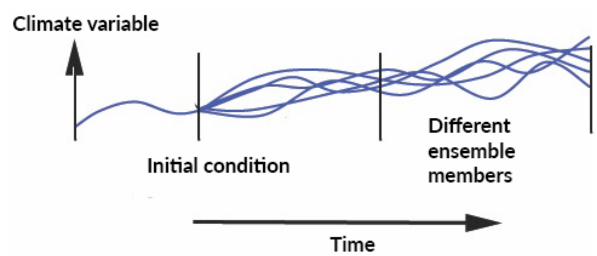

Figure 4. Principle of initial condition ensemble simulations. By choosing slightly different sets of initial conditions, equally likely realizations of a climate variable (such as temperature) are created.

The Promise of General Circulation Models

High-resolution NWP models are used primarily for weather forecasting. They cannot be used to make climate predictions on longer (seasonal to decadal) time scales with accuracy, not only because their resolution would require vast computing resources to do so, but also because they are not coupled with an ocean model. As we learned from **Part II**, weather patterns can be influenced by dynamics that involve both the atmosphere and ocean. The El Niño-Southern Oscillation (ENSO) is just one example of an atmospheric-oceanic phenomenon that affects weather—including tropical cyclone activity—worldwide. Enter the general circulation model, or GCM.

GCMs attempt to simulate a coarse-grained approximation of Earth's entire climate system. The most complex and resource-intensive component of a GCM is the atmospheric module, which uses the primitive equations to simulate the evolution of wind direction, wind speed, temperature, humidity, and atmospheric pressure, hereafter denoted as U, V, T, Q, and P, respectively. GCMs also include equations describing the oceanic circulation: how it transports heat and how the ocean exchanges heat and moisture with the atmosphere. Another important component is a land surface model that describes how vegetation, soil, and snow or ice cover exchange energy and moisture with the atmosphere. And yet another component captures the interactions between the atmosphere and the cryosphere (sea and land ice).

To solve the equations on a computer, GCMs divide the atmosphere, oceans, and land into a three-dimensional (3-D) grid (Figure 5). The equations are then numerically evaluated in each grid cell at successive time steps throughout the simulation period. The number of cells in the grid determines the model's resolution, or granularity. Each grid cell is characterized by the average value of each variable; therefore each cell effectively has a uniform velocity, temperature, etc.

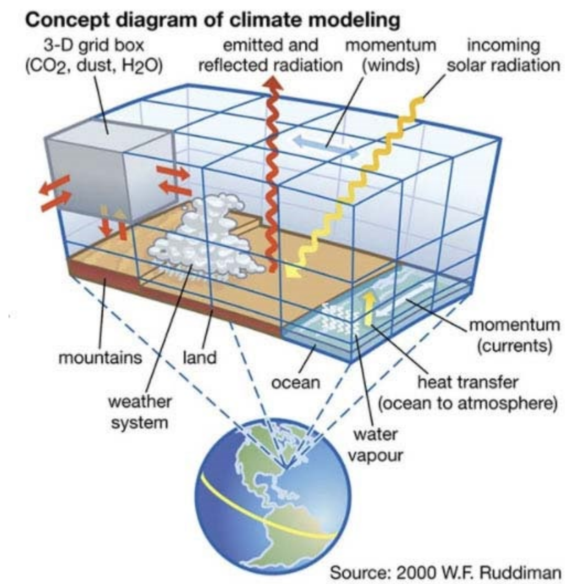

Figure 5. Schematic of a General Circulation Model (Source: Climate Information.)

While the most recent, state-of-the-art atmospheric GCMs might have a horizontal resolution of 25 km, the more typical GCM used for seasonal predictions and El Niño forecasting, for example, will have a horizontal resolution of about 100 km; a vertical resolution of about 1 km; and a time-stepping resolution of about 10 to 30 minutes. The horizontal resolution of the atmospheric component of most GCMs included in the World Climate Research Programme's Coupled Model Intercomparison Project (CMIP) is ~100 km, or 10 times coarser/lower than current NWP models. If they were to operate at a more granular/higher resolution, GCMs would represent some processes with more realism; however, the computational time required to do the calculations would increase substantially. For example, a doubling of resolution requires about 10 times more computing power because the time step must be halved; not only are there four times as many grid points to evaluate, but twice as many time steps are also required for the model to get to the same point in the future. Thus for most of the world's climate modeling centers, modeling a spatial resolution beyond 0.5 degrees (60 km at the equator) is not practicably feasible at present. The choice of model resolution is driven both by the available computer resources and by what physical, chemical, and biological processes are relevant to the model's unique purpose, which dictates the length and number of simulations to be conducted.

The part of a GCM that solves the primitive equations for U, V, T, Q, and P is called the dynamical core. Climate processes represented by this dynamical core are referred to as being "resolved" by the model. But with uniform values within grid cells available, the typical GCM is too coarse to solve important small-scale processes, including those that govern the extreme weather events of interest to catastrophe modelers, such as thunderstorms, tornadoes, and extreme rainfall events. Such "unresolved" sub-grid processes are therefore parameterized—and the parameterization formulas employed (which vary with the scientist or scientists involved) introduce uncertainty and potential bias. The issue of uncertainty and bias and how to reduce it in the context of catastrophe modeling will be discussed in further detail in **Part IV.**

GCMs and Climate Projections

Today, GCMs (sometimes combined with higher-resolution regional climate models for region-specific results) are providing forecasts of Earth's climate up to the end of this century. Uncertainty in the initial conditions doesn't play an important role here because the effect of model error dominates. Therefore, in addition to ensembles with different initial conditions, ensembles with different model parameters are used to evaluate the model error. Because many different GCMs are used, each with their specific biases, a multi-model analysis also provides a measure of uncertainty in the projections.

It's important to note that, despite the inherent model errors and biases, GCMs still do a reasonably good job of simulating general climate behavior: storms develop and move in realistic ways; temperatures change according to time of day and day of year in realistic ways; and precipitation falls where and when it should—generally. But the details—how intense storms will become, exactly where they will track or stall, and how heavy the precipitation will be—are not captured well enough to satisfy the catastrophe modeler.

In addition to model error, a major source of uncertainty when making climate projections over decades is the radiative forcing (the difference between energy in the form of sunlight absorbed by Earth and the energy radiated, or reflected, back into space) induced by anthropogenic greenhouse gas (GHG) emissions. GHG emissions depend largely on the usage of fossil fuels and thus human behavior. To cope with this uncertainty, the climate research community makes use of a suite of several emissions scenarios called Representative Concentration Pathways, or **RCPs**. Each RCP represents a potential trajectory of atmospheric GHG concentrations over the coming decades, culminating in a specific excess radiative forcing at the year 2100. The RCP 8.5 scenario, for example, assumes high and growing emissions that will lead to an 8.5 W/m^2 extra radiative forcing in the year 2100 and an average increase in surface temperatures of between 2.6°C and 4.8°C (at the 90% confidence level).

In the CMIP5 (CMIP Phase 5) projects undertaken in support of the Intergovernmental Panel on Climate Change (IPCC), the participating GCMs perform a set of predefined simulations resulting in an ensemble of climate projections. Each model first performs a century-long simulation under preindustrial initial conditions, which consist of prescribed solar forcing, aerosol forcing, and greenhouse gas forcing as of the year 1850. This serves as a control simulation. At the end of the control simulation, a so-called historical simulation is performed from 1850 to (usually) 2005. Next, the simulation is continued under one of the RCP scenarios.

For many models, more than one simulation is performed by starting the historical simulation from a different year than the one for the start of the control simulation. The reason being that, although the atmosphere responds quickly to various forcing conditions, the ocean takes much longer; thus at different years in the control simulation the ocean will be very different. These efforts then lead to probabilistic projections of, for example, the global mean surface temperature up to the year 2100 (Figure 6) for the different RCPs. Note that due to the uncertainties in both models and forcing, the behavior of the actual climate system can deviate from these results substantially, and even be outside the estimated range of possibilities illustrated by the shading in Figure 6.

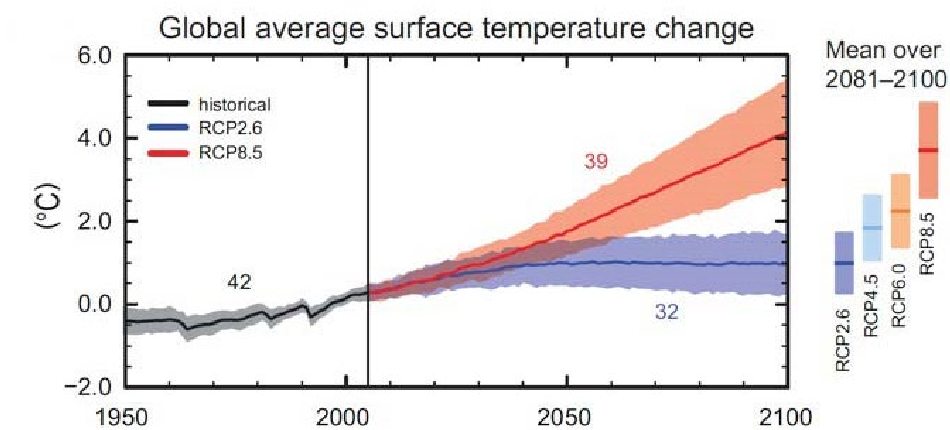

Figure 6. Global mean surface temperature as projected by the CMIP5 model suite, for different RCP scenarios. (Source: Source: IPCC, 2013)

In the realm of climate science, the principal goal of GCMs has been to forecast changes in average surface (land and ocean) temperatures. To determine whether the occurrence of extreme events will change over the next decades, large ensembles are needed, as such events are by definition rare. Several studies indicate that the probability distribution of extreme events will change as Earth's climate continues to warm. In the case of annual maximum temperature, for example, in many continental-scale regions the mean shifts to higher temperature and the amplitude of the positive tail of the distribution increases. Probabilities of occurrences of these extremes and the return period of specific amplitude extremes can be calculated from these results.

The Next Frontier: Overcoming Uncertainty and Bias in GCMs for Use in Catastrophe Models

GCMs are powerful tools built for purpose; however, none has been built with the catastrophe modeler in mind. While the IPCC's periodic Assessment Reports speculate on the likely increase in precipitation, for example, they are sparing in their commentary around the frequency and severity (and regional variation) of convective storms, tropical cyclones, and wildfires.

The central question then is: What confidence do we have in the model results as they relate to unobserved extreme events? While we must accept the uncertainties around greenhouse gas forcing, which will always have to be handled by scenarios, there are ways that we can reduce the spectrum of uncertainties and biases that arises through model error and parameterization. Part IV will explore these issues and point to AIR's solution for overcoming them.

Part IV: Climate Models in a Catastrophe Modeling Context: Opportunities and Challenges

Henk Dijkstra, Ph.D., Professor of Dynamical Oceanography, Institute for Marine and Atmospheric research, Utrecht; and Director, Centre for Complex Systems Studies, Department of Physics, Utrecht University

Boyko Dodov, Ph.D., Vice President and Director, AIR Worldwide

Climate Modeling at AIR: A Brief History

At their inception in the late 1980s, catastrophe models were constructed based largely on locally reported observation data. In the case of hurricanes, for example, these data might include central barometric pressure at landfall, forward speed, and angle of storm track. Probability distributions were fit to the data and simulated events were created by randomly drawing storm parameters from these distributions, taking care that the resulting draw was deemed physically plausible by meteorologists. It was a purely statistical exercise; the construction of these simulated storms was otherwise divorced from the global atmospheric and oceanic flows that give rise to the regional climate conditions that spawn and propel actual hurricanes.

This approach remains quite useful and largely valid for regions of limited domain where observational data is abundant and for perils that are physically cohesive and well defined, such as hurricanes. It is a less robust approach for regions where data are relatively scarce and for more amorphous weather systems characterized by considerable internal variability at fine scale—systems that cannot be neatly defined by a handful of parameters. Extratropical cyclones and severe thunderstorms are good examples of the latter.

Numerical Weather Prediction (NWP) and Reanalysis Data: A Regional Approach

To overcome the challenges of a purely statistical approach, AIR first introduced climate modeling—and, in particular, numerical weather prediction—into the AIR Extratropical Cyclone Model for Europe.

As we learned in **Part III**, numerical weather prediction (NWP) models are used for weather forecasting over relatively short time frames—typically, not more than 14 days. Just as the validity of an NWP forecast depends on the accuracy of the inputs, so too does the quality of catastrophe model output depend on the quality of the input data. In building a large catalog of simulated storms, AIR used NOAA **reanalysis** data of the environmental conditions (sea surface temperature, air temperature, wind speed, humidity, and atmospheric pressure) present at the time of roughly 1,500 historical "seed" storms affecting Europe over the last 40 years. These storms were then perturbed stochastically by employing robust statistical algorithms to create tens of thousands of potential future storms.

One potential limitation of such an approach is that the resulting catalog comprises what are, in effect, "siblings" of their historical counterparts. They are different, but the approach raises the question of whether we can be confident that we have captured all potential extremes. If we perturb the storms too much in an attempt to free ourselves from historical constraints, we may end up with results that no longer consistently and coherently represent the dynamic nature of the atmosphere.

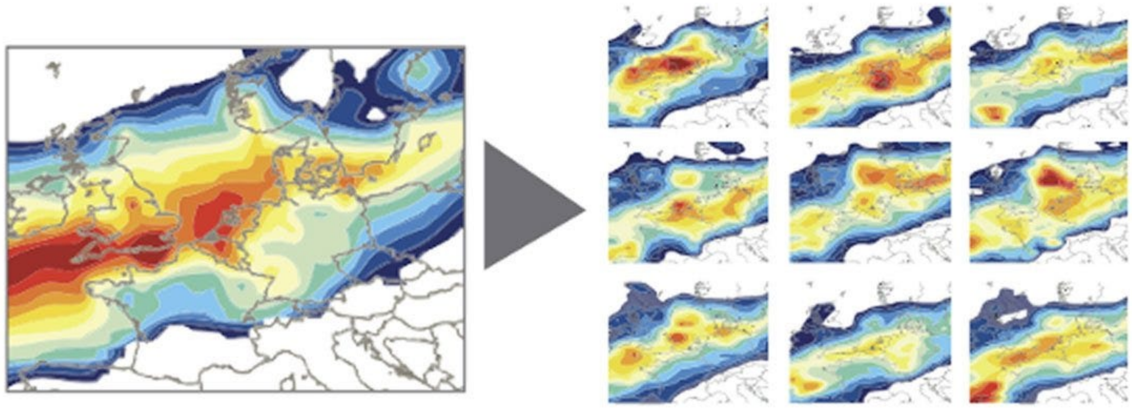

Figure 7. A historical seed storm is perturbed to create a set of possible realizations of such storms. (Source: AIR)

Another limitation of this approach is that NWP models, which currently run at a horizontal resolution of about 10 km, are typically regional in scope; they lack any relationship to other regions. Yet we learned from **Part II** in this series that *planetary-scale* motions are the drivers of *local* weather extremes.

The Promise of General Circulation Models

If the goal is to produce *global* catalogs that capture all types of dependencies—from global teleconnections to local correlations across all weather-related perils and across all regions—it would seem intuitive that the next step is to use a global *general* circulation model (GCM). Unfortunately, GCMs come with limitations, too.

While some of the processes in GCMs, as in all numerical models, are based on the laws of physics, the *primitive equations*, which include terms for the conservation of mass; a form of the **Navier-Stokes** equations governing fluid flow; and thermodynamic terms, as we discussed in **Part III**—there are other key processes in the model that are approximated, some of which are not based on physical laws. Recall from Part III that the "dynamical core" of a GCM is the part of the model that numerically solves the equations for wind speed and direction, temperature, humidity, and atmospheric pressure. The key to their solution is their spatial and temporal discretization by using various numerical methods. Depending on the spatial discretization, there are two major types of dynamical cores: (a) spectral, where the discretization is on waves of different length (i.e., bands in the frequency spectrum); and (b) gridded, working on spatial grids of various geometries. Both types are common in the climate modeling community and each has its pros and cons in the representation of the "true" continuous equations.

Climate processes represented by the dynamical core are referred to as being "resolved" by the model, as we discussed in **Part III**. Because of the relatively coarse spatial and temporal resolutions of the GCM grids, however, there are many important processes in the climate system that occur on scales that are smaller than the model resolution and contribute significantly to extreme weather on small scales. Examples include thunderstorms, tornadoes, convective clouds, and rainfall (Figure 8). Such *unresolved* sub-grid scale processes are represented by "parameterizations," which are simple formulas based on observations or derivations from more detailed process models. The parameterizations are "calibrated" or "tuned" to improve the comparison of the GCM's outputs against historical observation, and the parameterization formulas employed (which vary with the scientist(s) involved) introduce uncertainty and potential bias.

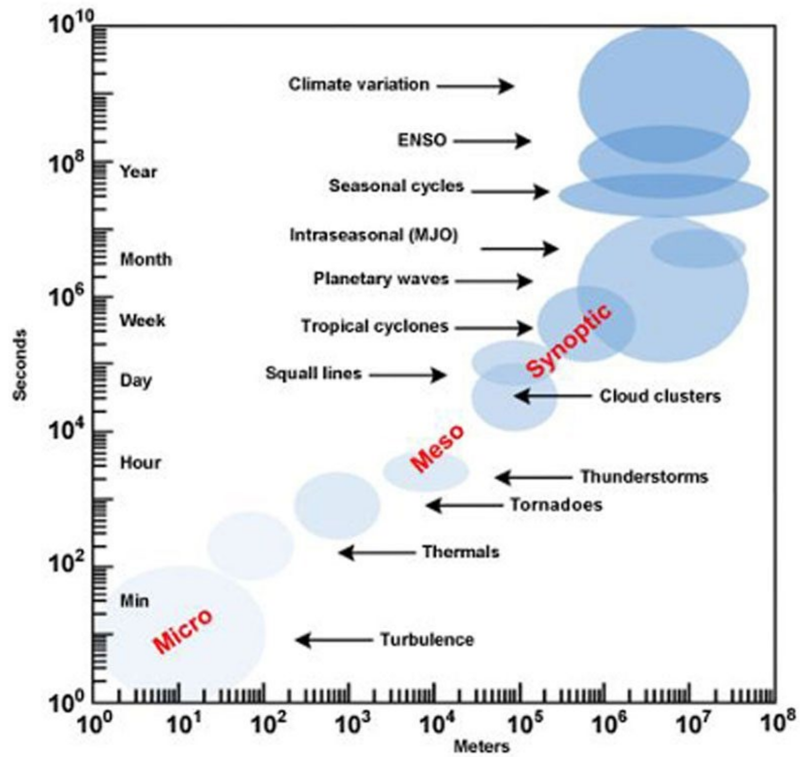

Figure 8. Resolved (dark blue) and unresolved (light blue) phenomena and processes in a GCM. (Adapted from Climate Change in Australia.)

Considering the potential biases introduced by discretization and parameterization, and the fact that the equations describing the resolved processes are to a large extent a limited view of reality, it is important to stress that GCMs can only approximate the physical processes they are designed to represent. While the large-scale dynamics are resolved in a GCM, the inaccuracies at smaller scales and their feedback on larger scales lead to some of these biases.

There are more than 20 international climate modeling groups, and there are thousands of different choices made in the construction of a GCM (resolution, type of dynamical core, complexity of physics parameterizations, etc.). Each set of choices produces a different model with different sensitivities and, most importantly, different statistics of the model output. Furthermore, different climate modeling groups focus on different interests—for example, long paleoclimate simulations, details of ocean circulations, nuances of the interactions between aerosol particles and clouds, or the carbon cycle. Given these different interests and many others, limited computational resources are directed toward one aspect of simulating the climate system in each case, at the expense of others.

To date, no GCM has been built to simulate the small-scale processes that produce the extreme weather events that the catastrophe modeler is interested in. In fact, most of the parameterizations tend to replace highly non-linear natural processes with their "average" response. As a result, the natural variability of the climate system tends to be lessened, thus missing the extremes. And while the resolution of GCMs has increased greatly over the last 10 years, the computational cost of generating very large (million-year) global catalogs of extreme weather events remains prohibitive.

Biases in GCMs: Examples

As we have described, GCMs have strong biases in simulating large-scale atmospheric phenomena relevant to the genesis of extreme events. Although visually these phenomena may look reasonable in GCMs, their statistics are often incorrect. For example, in evaluating the period 1961-2000, GCMs generally underestimate the frequency of wintertime blocking events over Europe. (Atmospheric blocks were discussed in **Part II**.) Blocking frequencies at lower latitudes are generally overestimated.

It’s important to note that, despite the inherent model errors and biases, GCMs still do a reasonably good job of simulating general climate behavior: storms develop and move in realistic ways; temperatures change according to time of day and day of year in realistic ways; and precipitation falls where and when it should—generally. But the details—how intense storms will become, exactly where they will track or stall, and how heavy the precipitation will be—are not captured well enough to satisfy the catastrophe modeler.

Regarding the polar jet (between 45°N and 50°N), most of the GCM models can reproduce seasonal variations of the jet latitude, but many overestimate the amplitude of the maximal wind speed. Figure 9 compares the daily mean wind speed in the polar jet as simulated by 11 GCM models from the World Research Programme’s Coupled Model Intercomparison Project Phase 5 (CMIP5) to historical (reanalysis) data shown in the bottom right panel. In most cases, the CMIP5 GCM models produce greater variability than the historical.

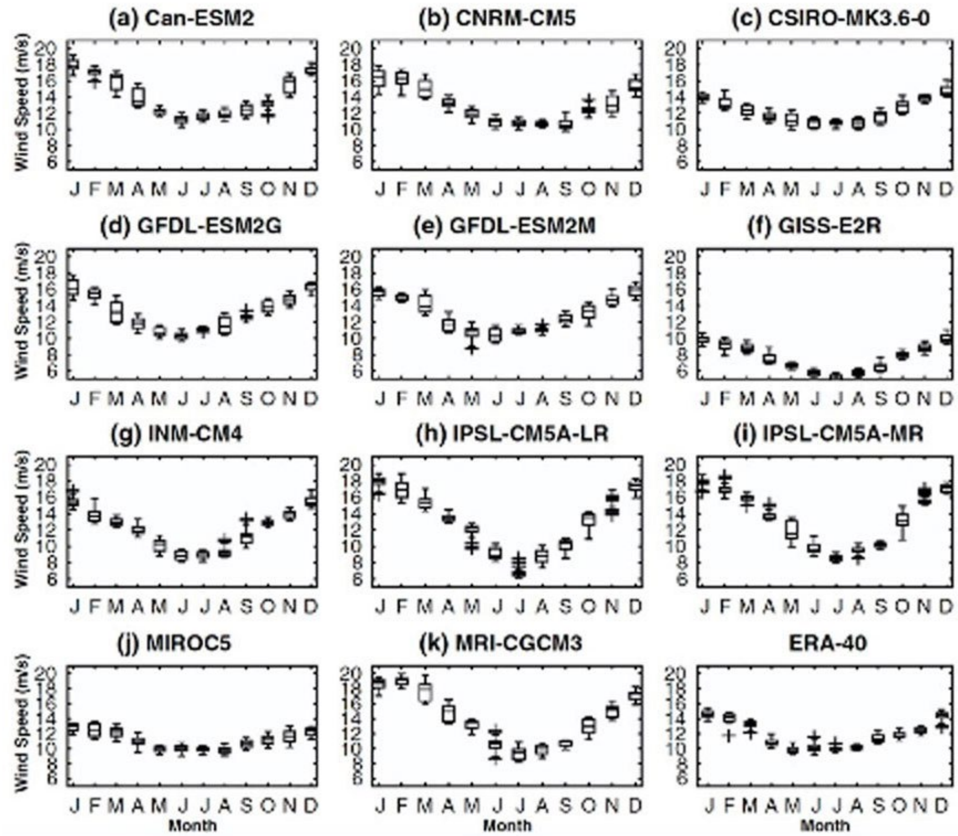

Figure 9. Boxplot of daily mean wind speed (m/s) of the polar jet from simulations of 11 CMIP5 models over the period 1980-2004 (acronyms above each plot) and historical (reanalysis) data over the period 1957-2002 (ERA40 in bottom right panel). (Source: Iqbal, W., Leung, W.-N., and Hannachi, A. (2018). Analysis of the variability of the North Atlantic eddy-driven jet stream in CMIP5. Climate Dynamics 51:235-247.)

GCMs also do not provide the necessary detail desired for extreme event forecasting on longer time scales, such as the prolonged periods of drought in many parts of Australia ahead of and during the [bushfire season of 2019-2020](). To capture that detail, regional general circulation models (RCMs) are often used. As their name would suggest, RCMs represent the climate over a limited region (such as Australia), and their resolution is typically much higher (down to 1 km) than a GCM's. These RCMs are connected to (nested within) the coarse-resolution GCM at the boundaries of the region. While these models provide more detail over the region of interest, the biases from the GCM cascade through to the RCMs. A GCM bias in the polar jet, for example, has a large effect on the regional atmospheric flow and can destroy the validity of a regional long-term forecast—in particular, regarding local extreme events.

In recent years, AIR has employed a hybrid solution for building atmospheric catastrophe models, one that nests a regional NWP model within a GCM. The high-resolution NWP model is connected to the coarser GCM at the boundaries of the region, then downscaled to a very high resolution using statistical algorithms followed by local climatological adjustment. For scientific questions on climate change, the GCM biases may not be a serious problem, as one is often interested in the difference between a future projection and the current climate simulation. For addressing questions related to the occurrence of extreme events, however, these biases pose a problem, as they can materially influence the spatio-temporal statistics—the patterns—of these events. Such pattern biases are critical in the context of loss occurrence when aggregated at a portfolio level.

Extreme Event Modeling for a Future Climate

Over the last decade, new ideas to better represent the unresolved sub-grid processes that drive extreme weather have emerged, such as stochastic parameterization (a probabilistic approach to unresolved processes) and super parameterization (building in a simplified high-resolution sub-model for cloud formation, for example). Although these approaches may improve the underlying climate model output, they do not provide explicit representation of the small-scale processes—a key requirement in a catastrophe modeling framework. From a catastrophe modeling perspective, the best approach for developing global simulations within which we can model the extremes may be using a GCM that has been debiased.

Some very recent research in the field of machine learning provides a promising solution in compensating for the biases introduced by the missing unresolved climate processes in a GCM, thus serving as a sophisticated parameterization scheme that can narrow the gap and render a coarse GCM output close to the reanalysis at the GCM resolution. Similarly, recent attempts have been made to use machine learning in downscaling—that is, in explicitly simulating unresolved processes in terms of the resolved ones. These ideas, when combined, have the potential for being implemented in efficient high-resolution climate simulations—and they represent the solution that AIR is developing as the foundational framework for our atmospheric peril models. They will also be the subject of **Part V**.

Part V: A New Framework for Global Climate Simulation, Purpose-Built for the Catastrophe Modeling Community

Themistoklis Sapsis, Ph.D., Associate Professor of Mechanical and Ocean Engineering, MIT

Maneesh K. Singh, Ph.D., Head, AI R&D, Verisk

Gautam Kunapuli, Ph.D., R&D Lab Manager, Verisk

Boyko Dodov, Ph.D., Vice President and Director, AIR Worldwide

The challenge for the catastrophe modeler can be simply stated: Simulate extreme but relatively localized weather events driven by global atmospheric circulations that obey extremely complex physical laws. To meet this challenge, the catastrophe modeling industry employs predominantly parametric statistical approaches and perturbations of historical events to produce large catalogs of simulated events. For some perils, regional numerical models may be employed to generate those perturbations—an advancement enabled by recent increases in compute power. Still, we inevitably fall short in representing the full complexity and non-linearity inherent in Earth's weather and climate.

At the other end of the spectrum of ways to meet our challenge are state-of-the-art, high-resolution Global Circulation Models (GCMs), which come the closest to a realistic representation of atmospheric circulation in all its complexity beyond the observed climate. Yet, as discussed in **Part IV**, even the most sophisticated GCM will have substantial biases related to the purpose for which it was built—and thus far no GCM has been built with the catastrophe modeler in mind. Perhaps even more importantly, implementing a state-of-the-art GCM in a catastrophe modeling context is simply infeasible because of its computational demands.

Therefore, we need to meet our challenge with a solution that harnesses the best of both approaches. At AIR, we see an opportunity to blend our traditional hybrid physical and statistical approaches with a new set of tools that come from the world of artificial intelligence—specifically, machine learning. This approach realistically and robustly represents the full complexity of atmospheric circulation but does so with computational efficiency, enabling us to generate very large catalogs of globally correlated events across multiple perils to explore the extreme limits of current and near-future climate.

The Contribution of Advanced Machine Learning to AIR's New Framework

Recent advancements in **machine learning** for weather and climate applications spawned the idea at AIR that emerging deep learning algorithms are the critical link between physics and statistics—the link that will enable a consequential upgrade of our catastrophe modeling framework.

Very generally, machine learning (ML) uses a universe of programmed algorithms to quickly learn and identify dependencies and rules from data—particularly "big data"—based on which they can make decisions or predictions that they were not specifically programmed to make. The algorithms themselves evolve as they learn.

The example shown in Figure 10 is very relevant to our framework because the equations governing fluid flow for a rising plume of smoke are the same as those governing atmospheric motions. The left-hand panel shows a snapshot of a coarse-scale physical model output. The right-hand panel is a fine-scale reconstruction, where small-scale fluctuations are simulated with ML algorithms and added to the coarse simulation to create realistic high-resolution model output at low cost.

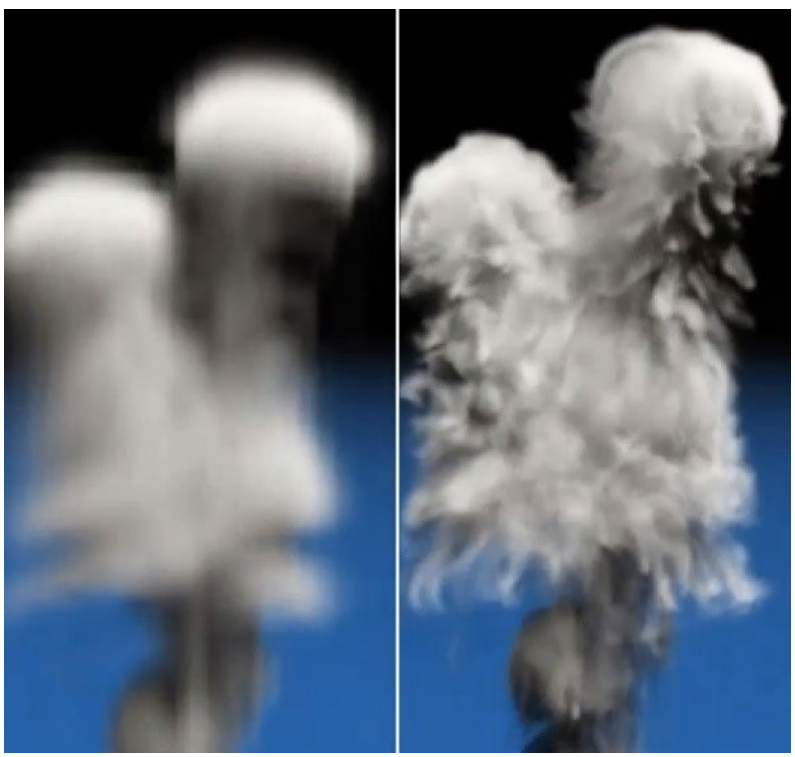

Figure 10. Interpolated physical model output of a smoke plume (left); a fine scale component of the flow is added by physically constrained ML algorithms to create realistic high-resolution model output (right). (Source: CS, Cornell University)

AIR's new framework builds on similar ideas and many years of modeling experience with a sophisticated approach. It combines a coarse GCM for the large scales with advanced ML algorithms at fine scales to obtain a physically consistent and statistically robust (i.e., data-driven) view of global risk at low computational cost. It is both efficient and purpose-built for catastrophe modeling.

Perhaps the most significant innovation is the development of a very specific flavor of machine learning algorithm designed for fluid flow simulations. The algorithm is the result of an ongoing collaboration among AIR, the Verisk AI Lab, the Massachusetts Institute of Technology, the Otto von Guericke University in Germany, and the University of Utrecht, Netherlands.

A distinguishing feature of AIR's deep-learning algorithm is that it can detect and simulate the propagation of waves, vortices, and other coherent dynamical structures in a fluid, making it ideal for atmospheric flow applications. The algorithm also learns the joint distribution of all output variables as it evolves in time as a function of the input variables, providing a host of opportunities for uncertainty treatment and data assimilation applications in a catastrophe modeling context.

Flood Modeling: The Current State of the Art at AIR

If the task is to model flood risk over a single small river basin, one can use a simple statistical model (uniform precipitation, for example) to accurately represent the local precipitation intensity time series. But if the task is to simulate events over large river basins (such as the Mississippi or Danube), which may last for weeks, continental-scale precipitation needs be integrated over space and time. A physical model, such as a GCM, is required.

Currently, most of AIR's flood models rely on a coarse global climate model coupled with regional numerical weather prediction models, followed by statistical downscaling to the resolution of a few kilometers. This framework does an excellent job at simulating precipitation at a continental scale to produce robust 10,000-year catalogs. With the industry's growing demand for even larger (million-year), globally unified, and multi-peril catalogs, however, the coupling of global climate and NWP models becomes computationally too expensive.

AIR's New Framework: Combining Three Critical Ingredients in a Two-Step Process

AIR's new framework for modeling atmospheric perils makes use of machine learning trained on reanalysis data to reproduce what would otherwise be computationally expensive fine-scale atmospheric circulations as a function of computationally inexpensive coarse-scale circulations. This requires three ingredients, or building blocks:

(1) **historical reanalysis data,** which serves as a benchmark to debias; (2) a **coarse-resolution GCM with simplified physics**; and (3) the cutting-edge machine learning algorithm resulting from our ongoing collaboration, as discussed—one specifically designed to simulate fluid flow. These three ingredients are combined in a two-step process that we discuss in the next section.

Reanalysis as Benchmark

Atmospheric **reanalysis** is a long-standing and ongoing project that uses data assimilation techniques to combine all available instrumental observations of past weather with simulations from numerical models to produce a complete and statistically, physically, and dynamically consistent recreation of the history of Earth's weather and climate. The reanalysis used in the development of AIR's framework is the ECMWF's **ERA5**, a fifth-generation reanalysis product covering the period from 1950 to the present, at a 0.25º spatial resolution and hourly timestep.

AIR has divided the data into several frequency (or, conversely, wavelength) bands, from fine to coarse scale. The schematic in Figure 11 shows two such bands. These bands provide benchmarks for machine learning to determine the dependencies between the coarse and fine scales, as well as for debiasing our coarse-scale GCM. Note that reanalysis data is characterized by many hundreds of variables. Our goal is to resolve the full complexity of atmospheric circulation *only* in the context of catastrophe modeling, so only those weather and climate descriptors relevant to catastrophe models need to be considered, making the task more focused.

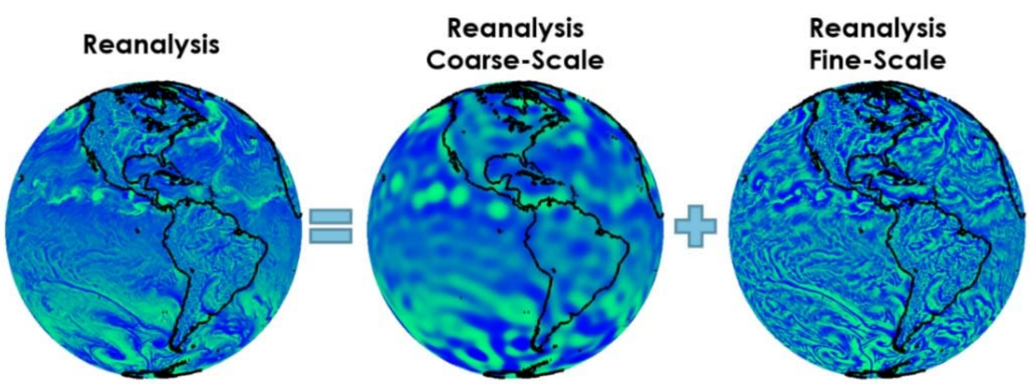

Figure 11. Illustration of splitting the ERA5 reanalysis data into coarse- and fine-scale frequency components. (Source: AIR)

Step 1: Debiasing the Coarse-Scale Climate Model (GCM) Output

As just noted and discussed in some detail in **Part IV**, the currently available state-of-the-art (high resolution) GCMs are not a good fit for the future of catastrophe modeling, both because of their biases and their computational cost. When *run* at a coarse resolution, however, the output from a state-of-the-art climate model is comparable to the output from a coarse one with simplified physics. A coarse and simple climate model has the advantage of speed; thus, we can use it to generate large catalogs of physically based atmospheric flow very quickly.

But given that these catalogs will be biased, our framework first requires us to debias the GCM output to ensure that we're producing realistic frequencies of atmospheric blocking, the polar jet, and other dynamical phenomena, as discussed in **Part II**. Otherwise, we cannot count on getting an unbiased representation of the frequency of stalled hurricanes like Sandy and Harvey, prolonged droughts like those ahead of the 2019-2020 Australia bushfires, or atmospheric rivers like the one that resulted in the Great Flood of 1862 that devastated Oregon, Nevada, and California. To do that we need to debias the climate model output by benchmarking it to "reality" in the form of the coarse-scale component of the reanalysis data (thus pairing the resolutions of the two data sets).

It's worth pointing out that this first step in our approach is a research project in itself, involving the application of cutting-edge machine learning techniques to perform the complex and high-dimensional mapping illustrated in Figure 12. This debiasing corrects both the local intensities of the model output parameters, as well as the patterns of these parameters evolving over time—that is, the atmospheric dynamics. By thus correcting the atmospheric dynamics of our simple climate model, we can get storm tracks, for example, and their frequency right. At the end, we have a fast-running, albeit still coarse, GCM without statistical and dynamical biases in the model output.

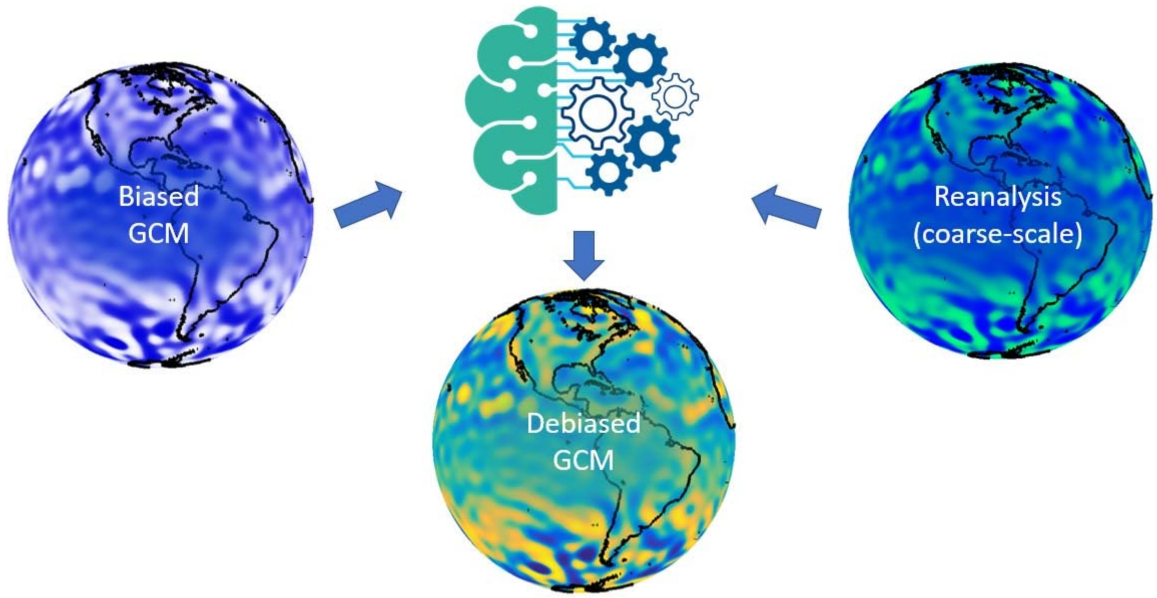

Figure 12. Illustration of the coarse climate model output by "mapping" the dynamics and statistics of the model output to the coarse-scale reanalysis data. (Source: AIR)

Step 2: Learning the Behavior of Fine-Scale Features as a Function of the Coarse-Scale Ones

The second step in the process achieves our goal of effectively replicating the output of a state-of-the-art GCM without actually employing one. That is, we reproduce fine-scale atmospheric circulations as a function of coarse-scale ones. We describe how we achieved this through our deep-learning algorithm in the following paragraphs.

The reanalysis data split into coarse and fine frequency bands acted as our training data, which was fed into the probabilistic machine learning algorithm. The algorithm learned two sets of rules from these data: (1) The primary set of rules governed the learning of deterministic dependencies—that is, the expectations of the fine-scale variables as functions of the coarse-scale ones; (2) the second set of rules is where we machine learned the statistics of the residuals of fine-scale variables, after the first set of rules was applied. These sets of rules are denoted as (1) $\boldsymbol{f_{ML}}$ and (2) $\boldsymbol{\varepsilon_{ML}}$ in Figure 13; they are the end product of Step 2 in our framework.

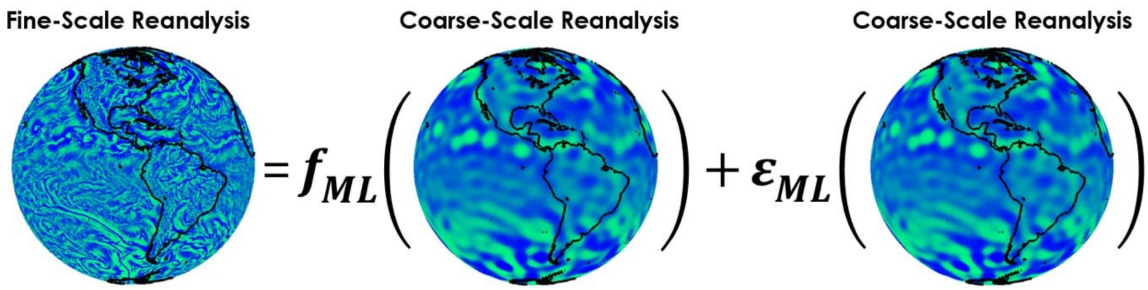

Figure 13. Application of our ML algorithm resulted in rules governing dependencies (middle) and statistics of the residuals (right). (Source: AIR)

This probabilistic machine-learning framework presents many exciting opportunities for using these functions in catastrophe modeling applications. We can employ them to quantify hazard uncertainty and create multiple versions of each stochastic event. In the context of historical and real-time events, the probabilistic ML framework can also be used to more faithfully calibrate the modeled hazard footprints to observations through data assimilation. While there are any number of additional benefits that can result from this research, our focus in the next section is on combining steps 1 and 2 to build global, multi-peril catalogs.

Combining Steps 1 and 2 into a Robust and Efficient Global Simulation of Atmospheric Dynamics

Recall that in Step 1 we debiased the output of our coarse-scale GCM. Step 2 gave us two sets of rules for simulating fine-scale dynamics from the coarse-scale GCM output, thus achieving the results of a state-of-the-art GCM but at low cost. Figure 14 illustrates how putting the solutions from Step 1 and Step 2 together allows us to efficiently simulate large global physics-driven *catalogs* with realistic dynamics and robust statistics at all scales.

To do that, we plug the adjusted large-scale variables from the simple climate model output from Step 1 into the machine-learned sets of rules obtained from Step 2 to simulate ensembles of fine-scale time series. We then add the large-scale and the ensemble small-scale variables to obtain the final full-scale catalog.

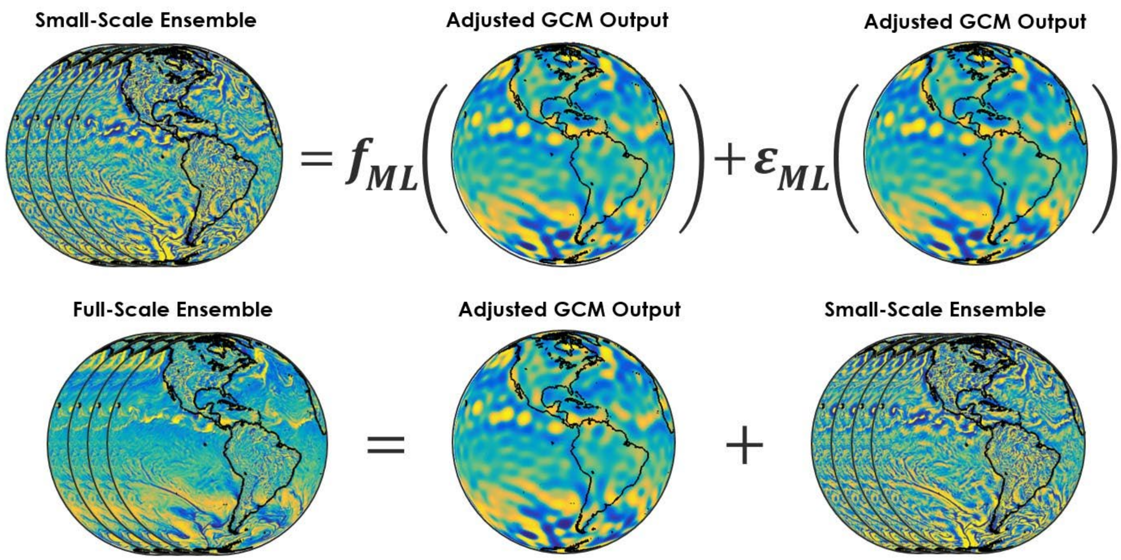

Figure 14. Illustration of the implementation sequence: Creation of fine scale-components from the adjusted simple climate model output (top row); combining the adjusted simple climate model output with the small-scale ensemble to obtain a full-scale ensemble, or catalog (bottom row). (Source: AIR)

Next Step: Fundamentally Change How Weather and Climate Are Simulated

The most sophisticated atmospheric peril models available to the insurance industry currently rely on perturbations of reanalysis data that are restricted to fine scales. This is because perturbing the large scales will make the event footprints unrealistic. Perturbing only the small scales, however, results in large-scale patterns that too closely resemble historical events, which are then repeated many times over in a stochastic catalog. The same would happen to parametric models that use large-scale reanalysis data for conditioning—that is, we may never see large-scale patterns capable of producing events far more severe than Hurricane Harvey, even though we know they are physically possible. Without the ability to explicitly simulate the large-scale patterns, we will never understand the probability of occurrence of atmospheric dynamics that give rise to the most extreme events—more extreme than historical events and in locations other than we've observed.

The new framework under development at AIR will address these issues and will fundamentally change the way weather and climate are simulated in the industry. Those tasked with managing risk will have access to large, robust global catalogs that capture all types of dependencies—from local correlations to global teleconnections and across *all* atmospheric perils, including tropical and extratropical cyclones, floods, severe thunderstorms, and droughts.

The coarse model output will have the correct frequency and magnitude statistics of large-scale phenomena, such as blocking, cyclogenesis, and **Rossby waves**, as well as the fine-scale signatures of tropical cyclones and convective storms. We can use the framework such that it can move around such a signature to simulate tropical cyclone wind and precipitation. Alternatively, we can train our algorithm with regional reanalysis data and add small scale detail to a specific region, such as the U.S., or Japan (Figure 15).

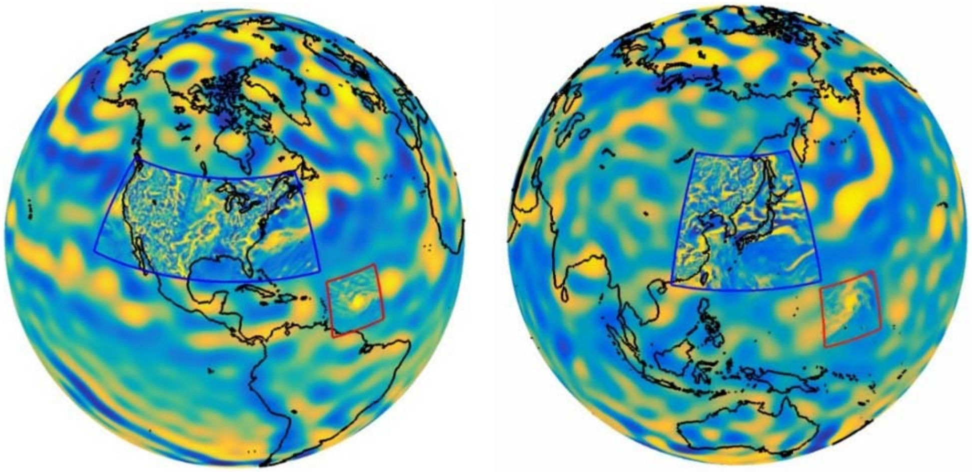

Figure 15. Illustration of the framework implementation for different geographic regions. (Source: AIR)

Finally, we can leverage the stochastic nature of our catalogs for efficient catalog versioning, where the versions can be conditioned on climate change or climate index oscillations. Part VI will describe how our new modeling framework can be used to provide insights into the potential impacts of climate change on the locations, frequencies, and intensities of extreme events.

Part VI: Using AIR's New Modeling Framework to Understand Climate Change Impacts

Peter Sousounis, Ph.D., Vice President and Director of Climate Change Research, AIR Worldwide

Boyko Dodov, Ph.D., Vice President and Director of Atmospheric Peril Models, AIR Worldwide

Jayanta Guin, Ph.D., Chief Research Officer, AIR Worldwide

The motivation to create a new paradigm in the way the insurance industry quantifies financial risk due to atmospheric perils was straightforward: How can we capture not just the physics of local or regional weather systems but also the physics of the larger, planetary-scale circulations that often operate on longer time scales and typically have a strong influence on the evolution of the small-scale weather phenomena that cause the damage? Only by understanding (and modeling) these "climate dynamics" will we achieve a robust representation of the most extreme events across multiple regions and perils. Not only will we be better positioned to capture atmospheric blocking events—such as the one that caused Hurricane Harvey to dump more than 50 inches of rain on Houston in 2017—but we'll be able to more confidently estimate the future frequency and geographic distribution of Harvey-like events worldwide. Because the physics of the atmosphere are not quick to change, models developed using the new framework should provide a robust view of risk for a time frame of 10 years or more.

But it is not only today's climate dynamics that concern us. The world finds itself at a tipping point when scientific investments in addressing climate change can still make a material difference. It was therefore important to choose a framework that could also be used to evaluate climate change risk on short- and long-term time horizons.

Moving from the Current to a Future Climate: Short- and Intermediate-Term Time Horizons

Part V described our groundbreaking framework, which blends AIR's current hybrid of physical and statistical approaches with machine learning; the result is that we can reap the benefits of general circulation models (GCMs) while circumventing their computational cost and reducing model bias. The framework makes million-year physically based catalogs possible—global catalogs that capture all types of dependencies, from global teleconnections to local correlations across all weather-related perils and across all regions. For certain perils on which the impact of climate change is more certain and quantifiable, such as extreme precipitation, we can make explicit adjustments to create forward-looking views.

But how do we make the leap from modeling the current climate to a future one? AIR's machine learning algorithm is trained on reanalysis data. While these data represent a statistically, physically, and dynamically consistent recreation of the history of Earth's weather and climate, they are, in the end, historical data. One solution is to create a climate change–conditioned catalog by subsampling from the inventory of simulated years (samples) that comprise the catalog representing the current climate. We are currently using a subsampling strategy in many of our climate change studies using existing AIR models. It is an approach that, for a couple of reasons, is particularly appropriate for short- to intermediate-term time horizons (10 to 30 years into the future). First, anthropogenic warming is happening slowly; this means that a climate influenced by it will likely not differ so much from the current one for at least the next one to three decades. Second, the new framework will produce a catalog large enough to enable us to choose a sufficient number of samples that could occur in a short-term time horizon, albeit with event frequencies representative of a climate change "target"—that is, a selected future year given a selected greenhouse gas emissions scenario, or **RCP**.

There are, of course, some differences between using our existing models and using the global catalog produced by our new framework to subsample a climate change–conditioned catalog. For the studies AIR has conducted to date, the future climate catalogs are, like our models, peril- and region-specific. Thus, in creating a climate change target, one has only to be concerned with how a single peril (and associated sub-perils) may change over a single region, such as a country, ocean basin, or continent, but not the entire planet. While there may be a temptation to subsample by randomly selecting individual events that are reflective of our target, the end result may not be physically consistent. For example, it may not make sense to draw hurricane events from both positive and negative Atlantic Multidecadal Oscillation (AMO) indices and combine them in the same year. More combinations that are similarly egregious could also result.

It thus makes more sense to draw entire years at a time so that global teleconnections are preserved, although this is also not without complication. For example, currently observed intra-annual correlations between Atlantic hurricane activity and winter storm activity over Europe might be different in a future climate. Because our strategy would only apply to short- to intermediate-term time horizons, however, that will likely not be the case. The tremendous advantage of the new global catalog is that it would represent the relevant parts of the actual physics of the atmosphere. Thus, in specifying a climate change target for one region and one peril, the other perils for other regions would, by default, appropriately reflect the climate impact as well.

In addition to creating climate change–conditioned catalogs to reflect short- to intermediate-term climate change targets, subsampling can also be used to reflect climate variability for future climate scenarios. For example, we could create an Atlantic hurricane catalog that reflects the impact of the La Niña phase of the El Niño–Southern Oscillation (ENSO) in conjunction with a positive AMO index, without first having to determine the change in hurricane frequency and then drawing years. That's because, in our new catalog, large-scale circulations would already be consistent with particular phases of climate oscillations. Indeed, each year in the catalog could be tagged with corresponding indices for a variety of climate oscillations/signals. Presumably, the appropriate hurricane frequencies, intensities, trajectories, etc., would then be captured as well.

One other way that subsampling can be used with the new model is that after the "current climate" version of the model has been in operation for 10 years or so, the extent to which climate change has had an influence on weather systems can be evaluated in much the same way AIR does it now—by analyzing historical data for physical (and statistically significant) trends to quantify the climate change effect. We can then compare the new current climate to the one that the model represents from 10 years or so ago. If we see signals—e.g., storm tracks are showing systemic changes—we can dive into our superset of events and subsample annual seasons that reflect the changes in underlying conditions. This is not an easy task, but relative to current practice we have a physical basis to conditionally create a new set of event simulations. This approach is sustainable for many years, circumventing the need to build a new model almost from scratch.

Moving from the Current to a Future Climate: Longer-Term Time Horizons

At some point, though, we will likely have to revisit building a new model almost from scratch. In **Part V**, we described the key components of our new framework as being: 1) a coarse resolution general circulation model (GCM) debiased to represent the observed climate; 2) a source of high-resolution information representing fine-scale weather features, and; 3) a machine learning algorithm trained on the second component to "learn" how fine-scale weather features depend on coarse-scale ones; the rules and dependencies learned are used to debias the GCM. For the model representing the current climate, the second component is reanalysis data.

To simulate a future climate, the coarse-resolution physics-based component could still be provided by a GCM, but one that is simulating a future climate. Any one of the models from the **Coupled Model Intercomparison Project** Phase 5 or 6 would be suitable. But what of the second component? Reanalysis data, which is in essence observational, can be thought of as "ground truth" (ocean and atmosphere truth, as well). But there is no ground truth for a future climate; the observations don't yet exist.

Hope is not lost, however, if we assume that the bias that exists under current climate conditions between the reanalysis (ground truth at fine scale) and the GCM (modeled truth at coarse scale) is the same for the future climate. For example, if the model runs 0.2° Celsius too warm for the historical period, can we reasonably assume that it will continue to run 0.2° Celsius too warm in the future? The answer is yes. This is a typical strategy employed by thousands of climate scientists: evaluate impacts from climate change by considering the delta(s) between future and current climate GCM runs to make statements such as, "Model X shows that, under RCP 4.5, global atmospheric temperatures will increase by 2° Celsius."

In the absence of reanalysis data, however, the same GCM output used to define the coarse-scale component of our framework must also be the source of the fine-scale information. Ideally, we would want the same high-resolution output as we do for historical reanalysis data covering about the same length of time. That means having 40 years' worth of future climate GCM output at a resolution of 30 km.

At this point one might ask why, if we have that kind of future GCM output, would we even need to build a future climate version of our new framework? The answer is the same as for the present climate. Forty years' worth of data is simply not enough to capture the full range of possibilities that might occur under that climate. It is the whole motivation for generating 10K, 100K, or even million-year catalogs. Indeed, it is the whole motivation for catastrophe modeling. But while generating 40 years' worth of high-resolution future GCM output may well be computationally inexpensive in 10 years, that will likely not be the case for generating 10K or 100K (or more) years of such data. The promise of AIR's new modeling framework is that it circumvents these costs while taking care of any biases.

Our second assumption is that the dependencies between the coarse and fine scales identified by our machine learning algorithm for the current climate also hold true for the future climate. This allows us to, in effect, back out "future reanalysis" data. We can assume, for example, that if there is a large-scale high-pressure ridge over the U.S. Pacific Northwest and deep low pressure over Nova Scotia in winter, then the possibility exists for a powerful Nor'easter to impact the eastern U.S. Again, it's a reasonable assumption. The Nor'easter may have a different strength in the future than it does in the current climate, it may move faster or more slowly, and the pattern may have a different frequency, but the basic configuration/orientation of large- and small-scale features should be the same. Any biases in the small-scale future climate GCM output may be corrected by using information obtained by comparing reanalysis data with high-resolution GCM output from a high-resolution re-simulation of the historical climate.

It's important to note that once we begin evaluating longer-term (>30 years) climate change impacts using our new framework, we must be circumspect. We should not read more into the data than is there or overextend the state of science. It is easy to be seduced by solutions that offer false precision, even while knowing that consensus in the scientific community is transient at best, particularly as it relates to the impact of climate change on individual atmospheric perils.

The Way Forward: Building a GCM that Captures Extreme Events Under Future Climate Conditions

While we may have solved our problem conceptually, much has yet to be done. When we undertook this project nearly two-and-a-half years ago, the goal was to create a model that would, for the first time, capture the planetary-scale atmospheric waves that can drive small-scale local extremes *under current climate conditions*. When complete, the model will be physically consistent across multiple regions and perils, so stakeholders can evaluate the global risk to their assets and portfolios for the next 10 years. It is also worth noting that in a time horizon of up to 10 years, the occurrence of extreme events will be driven more by natural **climate variability** than by climate *change*—a circumstance that will continue to be the case in future decades. Each new decade of data on climate variability, which would be used to update the model, would include the effects of climate change that have already taken place.

But it has become increasingly apparent over the last couple of years that clients want that kind of knowledge now to make business decisions at longer time horizons. Just as we have engaged in a collaborative effort with the scientific community to build a current climate version of our new model, so too will we likely need to rely on support from the community on a grander scale. Right now, AIR is ahead of the curve, but as new studies are conducted and published by academia and other research organizations, we will assimilate the findings and capabilities into our evolving plan to build a global climate model that captures extreme weather events from different perils in different locations under future climate conditions. At AIR, we are confident that our new framework will serve a multitude of purposes on a multitude of time scales.

About the Authors

Henk Dijkstra, Ph.D., Professor of Dynamical Oceanography, Institute for Marine and Atmospheric research, Utrecht; and Director, Centre for Complex Systems Studies, Department of Physics, Utrecht University

Dr. Dijkstra is an internationally renowned scientist, expert in the field of dynamical systems methods to problems in climate modeling, climate variability, and climate change. He is the author of several books and numerous articles covering different aspects of climate dynamics and climate change in the biosphere, hydrosphere, and atmosphere. Dr. Dijkstra is a member of an international team of leading scientists supporting AIR in our effort to build a new generation of catastrophe models capable of providing a global view of all weather-related perils in a manner efficient from a technological point of view.

Boyko Dodov, Ph.D., Vice President and Director of Flood and Atmospheric Peril Modeling, AIR Worldwide

Dr. Dodov is Vice President and Director of Flood and Atmospheric Peril Modeling in AIR's R&D group. He is one of the key developers of all AIR flood models and is now overseeing and contributing to the development of the global suite of all models related to weather and climate phenomena. Under Dr. Dodov's leadership, AIR has embarked on a novel approach to global weather and climate simulation, which is expected to unify and streamline the company's catastrophe modeling technology and allow an efficient global multi-peril risk assessment delivered to our clients. He received his B.S. in Hydrogeology from the Higher Institute of Mining and Geology in Sofia, Bulgaria, and his Ph.D. in Hydrology from the University of Minnesota. He is the author of numerous peer-reviewed articles related to the statistical representation and the modeling of extreme weather-related natural phenomena.

Jayanta Guin, Ph.D., Chief Research Officer, AIR Worldwide

Dr. Guin is responsible for the strategic direction and management of the Research and Modeling group. Under his leadership, the group has developed a global suite of models for assessing and managing extreme risk from both natural and man-made perils, continually pushing the boundaries of the state-of-the-science. Of late, he is leading the development of a set of solutions for analytics in casualty and cyber risks. Dr. Guin has more than 20 years of experience in probabilistic risk analysis for natural catastrophes worldwide and is well recognized in the insurance industry for his expertise in the financial risk posed by natural perils.

Gautam Kunapuli, Ph.D., Program Manager, Verisk

Dr. Gautam Kunapuli is a Program Manager at Verisk Innovative Analytics (VIA), Global Research. His current responsibilities include AI/ML R&D projects in document understanding and information extraction, spatio-temporal and weather modeling, pandemic modeling, insurance analytics, and explainable and ethical AI. Dr. Kunapuli's primary research efforts have been in machine learning and optimization, specifically human-in-the-loop and knowledge-based learning, efficient and scalable learning algorithms for hard machine-learning problems and neuro-symbolic learning. He received his Ph.D. in Mathematics from Rensselaer Polytechnic Institute in 2008.

Themistoklis Sapsis, Ph.D., Associate Professor of Mechanical and Ocean Engineering, MIT

Dr. Sapsis is Associate Professor of Mechanical and Ocean Engineering at MIT. He received a diploma in Ocean Engineering from Technical University of Athens, Greece, and a Ph.D. in Mechanical and Ocean Engineering from MIT. Before becoming a member of the faculty at MIT, he was appointed Research Scientist at the Courant Institute of Mathematical Sciences at New York University. He has also been visiting faculty at ETH-Zurich. Prof. Sapsis' work lies on the interface of nonlinear dynamical systems, probabilistic modeling, and machine-learning methods. A particular emphasis of his work is the formulation of mathematical methods for the prediction and statistical quantification of extreme events in complex engineering, physical, and environmental systems.

Maneesh K. Singh, Ph.D., Head, AI R&D, Verisk

Dr. Maneesh K. Singh leads the R&D efforts in AI and ML at Verisk Analytics and is the head of Verisk Innovative Analytics (VIA), Global Research. His current responsibilities include building efficient, collaborative, trustworthy and explainable AI and ML models, systems and platforms that draw knowledge and insights from unstructured and structured data sets with applications in areas spanning the insurance, financial services, and the energy sector. The research interests of VIA research span computer vision, audio, speech and language understanding, financial applications, efficient human-machine collaborative workflows, ethical and explainable AI. Dr. Singh's primary research efforts have been in the areas of computer vision with over 50 research articles and 20 patents. He received his Ph.D. in Electrical and Computer Engineering from the University of Illinois at in 2003.

Peter Sousounis, Ph.D., VP and Director of Climate Change Research, AIR Worldwide

Dr. Sousounis' responsibilities include ensuring that current and future catastrophe model development at AIR accounts for climate change, identifying products and tools to help clients address their climate change concerns, assisting with global resilience projects, and providing thought leadership in various forms of oral and written communication. He has been responsible for overseeing all global atmospheric model development, including hurricanes, extratropical cyclones, and severe thunderstorms, as well as for building the first numerically based storm surge model and the first ever tsunami model. He has authored nearly 100 publications on various topics of weather, climate, climate change, and catastrophe modeling.

About AIR Worldwide

AIR Worldwide (AIR) provides risk modeling solutions that make individuals, businesses, and society more resilient to extreme events. In 1987, AIR Worldwide founded the catastrophe modeling industry and today models the risk from natural catastrophes, supply chain disruptions, terrorism, pandemics, casualty catastrophes, and cyber incidents. Insurance, reinsurance, financial, corporate, and government clients rely on AIR's advanced science, software, and consulting services for catastrophe risk management, insurance-linked securities, longevity modeling, site-specific engineering analyses, and agricultural risk management. AIR Worldwide, a Verisk (Nasdaq:VRSK) business, is headquartered in Boston, with additional offices in North America, Europe, and Asia. For more information, please visit www.air-worldwide.com.

About Verisk

Verisk (Nasdaq:VRSK) provides predictive analytics and decision-support solutions to customers in the insurance, energy and specialized markets, and financial services industries. More than 70 percent of the FORTUNE 100 relies on the company's advanced technologies to manage risks, make better decisions and improve operating efficiency. The company's analytic solutions address insurance underwriting and claims, fraud, regulatory compliance, natural resources, catastrophes, economic forecasting, geopolitical risks, as well as environmental, social, and governance (ESG) matters. Celebrating its 50th anniversary, the company continues to make the world better, safer and stronger, and fosters an inclusive and diverse culture where *all* team members feel they belong. With more than 100 offices in nearly 35 countries, Verisk consistently earns certification by Great Place to Work. For more: Verisk.com, LinkedIn, Twitter, Facebook and YouTube.

Verisk