PURE

Probability, Uncertainty & Risk in the Environment

NERC SCIENCE OF THE ENVIRONMENT

Transforming the quantification and communication of natural hazard risk

Contents

Introduction

NERC’s £7m Probability, Uncertainty & Risk in the Environment (PURE) has helped UK Government and industries be better prepared for natural hazards

The number of major disasters resulting from natural hazards has risen dramatically since the 1970s, and is still increasing. From floods and ash clouds, to droughts and earthquakes, natural hazard events claim thousands of lives every year, and financial losses amount to billions of dollars. The Great Tohoku Earthquake (2011) alone left almost 19,000 people dead or missing, and hundreds of thousands homeless as it wiped out entire towns. The World Bank’s estimated economic impact was $235 billion, making it the costliest natural disaster in world history.

While events such as these are often confined to hazard prone areas of the world, the UK is not immune to the impacts of natural hazard events. December 2015 was an extraordinary month in both meteorological and hydrological terms, leading to widespread and severe flooding. In early December, Storm Desmond brought severe gales to the UK and localised flooding to the north west of England. An estimated 5,200 homes were affected by flooding, while in the town of Lancaster, 55,000 people were left without power for several days when an electricity substation was flooded.

NERC’s £6.8m Probability, Uncertainty & Risk in the Environment (PURE) action comprised a research programme and knowledge exchange network. Its aim was to transform the way that uncertainties associated with risk assessment for natural hazard events are dealt with.

PURE was ultimately designed to help the UK government, and industries as diverse as finance, energy and aviation, to prepare for such events – thus decreasing associated financial losses and increasing societal resilience.

This document outlines its successes.

Network

The PURE Network is a knowledge exchange network managed by the Smith Institute for Industrial Mathematics and System Engineering. Its aim is to ensure the science arising from the latest research is translated for use by decision makers in industry, business, government and the third sector. Collaboration between academia and these sectors is encouraged meaning researchers gain feedback on which aspects of their work have greatest impact. Meanwhile, users of science in industry and government can inform strategic directions of future research. The Network encourages inter-disciplinary and inter-sectoral collaborations and organised over 50 events and projects over the life of the research programme. Highlights include over 20 PURE Associates projects that were set up to put a scientist at the heart of a user organisation, so they could implement the latest scientific thinking to solve real business and policy problems.

Research programme

The aim of the PURE research programme was to improve the assessment of risk from natural hazards by developing new methods for the quantification of uncertainty, and demonstrating their applicability to enhance the uptake of state-of-the-art natural hazards science.

Uncertainty is ubiquitous in natural hazards, arising both from the inherent unpredictability of the hazard events themselves, and from the complex ways in which these events interact with their environment and with people. Scientists, governments and businesses have recognised this for a long time.

The PURE research programme sought to:

- Promote the tailoring and uptake of statistical techniques that allow some of the uncertainty to be represented explicitly, reducing the need to rely on a ‘margin of error’
- Develop methods that allow us to assess less quantifiable aspects of uncertainty
- Improve the visualisation and communication of uncertainty and risk

The PURE research programme brought together the expertise needed to achieve these objectives, creating greater consistency and rigour, and making risk assessment more transparent. Two research consortia were funded: the Consortium on Risk in the Environment: Diagnostics, Integration, Benchmarking, Learning, and Elicitation (CREDIBLE) and the Robust Assessment and Communication of Environmental Risk (RACER) consortium.

Portal

Running alongside the main PURE research activities is the development of the PURE web portal - www.pureportal.org - designed to facilitate the sharing and fusion of models and data from a variety of different sources.

The portal has been developed by the British Geological Survey and is for use by the whole user community. The site provides free access to metadata for model codes and their applications (*model instances*). In the future NERC-funded researchers who have produced environmental model outputs will supply metadata for their models to be uploaded to the Portal. This has the advantage of massively increasing the potential impact of NERC-funded environmental modelling, and also enhancing the PURE portal as a central place to access information about environmental models.

CASE STUDIES:

CREDIBLE

The Consortium on Risk in the Environment: Diagnostics, Integration, Benchmarking, Learning, and Elicitation

CREDIBLE brought together a combination of scientific expertise from the universities of Bristol, Exeter, Lancaster and Oxford, with leading scientists and practitioners from industry and government. CREDIBLE's project partners included government agencies (Met Office, Environment Agency), reinsurance and catastrophe modelling companies (RMS, Willis Re, XL Catlin), environmental consultants (HR Wallingford, JBA Consulting), and infrastructure owners (Thames Water). The project was led by Thorsten Wagener, Professor of Water and Environmental Engineering at the University of Bristol.

A key aim of our project was to develop and implement generic methods that work across natural hazards, both current and emerging, and to apply them to a wide range of case studies. Highlights include:

- A Bayesian severe weather warning system developed and currently tested with the MetOffice, providing transparent and tailored warnings for individual end-users.

- Two new software toolboxes for sensitivity analysis and uncertainty assessment, already being used by hundreds of academic users worldwide and trialled by our industrial partners.

- Hazard case studies of sensitivity, uncertainty, risk analysis and communication in floods, droughts, lahars, windstorms, tsunamis, volcanic ash clouds, algal blooms, landslides, bridge scour and earthquakes.

- A new approach to rapidly bounding the exceedance probabilities of high aggregate losses for the insurance industry.

- Downscaled climate model simulations for the UK, generated with the weather@home computing facilities, providing a unique spatially coherent dataset for hydro- meteorological risk analysis.

- Over a hundred researchers and practitioners trained through our annual one-week summer school on risk and uncertainty in natural hazards.

Using weather forecasts to optimally issue severe weather warnings

Warning systems play a major role in reducing economic, structural and human losses from natural hazards such as windstorms and floods

The challenge

A weather warning system is a tool by which imperfect forecasts about the future are combined with potential consequences to produce a warning in a way that is deemed optimal.

A warning system is only useful if well defined and thus understood by stakeholders. The challenge here was to improve the current severe weather warning system used by the UK Met Office, by making it more transparent and tailored for the various end-users.

What was achieved

Based on sound mathematical theory, we produced a tool that combines predictions of future weather with user-attitude towards false alarms/missed events to produce bespoke warnings that are optimal for each user.

Below are examples of rainfall warnings in 15-31 October 2013 for two such users: 1) an end-user who is tolerant towards false alarms (left), and 2) a forecaster who issues warnings and thus less happy about false alarms as they might affect her credibility (right). There are 4 increasing levels of warnings: green, yellow, amber and red, and 4 rainfall intensity categories: very low, low, medium and high. The height of the bars indicates the forecasted rainfall intensity (1 for no rainfall, 8 for high rainfall) whereas the symbols at top of each bar indicate what actually happened. Clearly the end-user and the forecaster have very different views about what warnings they want to see, which is what our framework is designed for: bespoke warnings to all end-users, with minimal user input regarding false-alarm appetite.

"Weather forecasts and warnings are only useful if people use them to make decisions which help to protect lives, livelihoods and property. As forecasts become more sophisticated and include information on probability and risk they are potentially more valuable, but interpretation needs to be tailored to the vulnerabilities of particular decision-makers. This project has significantly advanced our capability to apply this in the field of severe weather warnings and opens the possibility of warnings tailored to the needs of individual users."

Ken Mylne
Met Office

Met Office first-guess warning
Using the Met Office rule

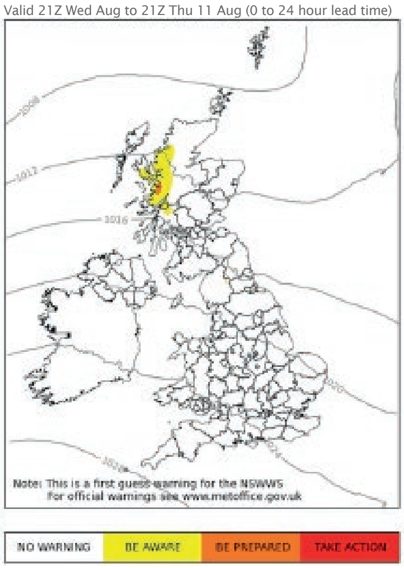

This warning area is most similar to the areas issued by the forecasters.

Met Office first-guess warning
Minimising loss for a forecaster

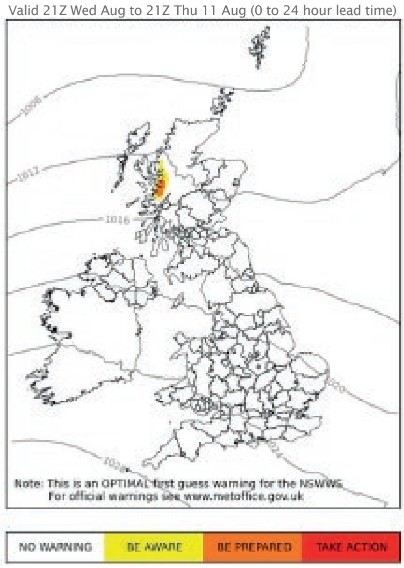

This is the smallest warning area. Forecasters want to minimise loss from too many false alarms.

Met Office first-guess warning
Minimising loss for a householder

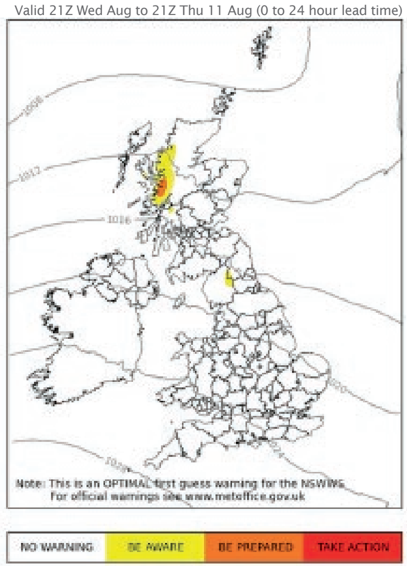

Similar to the Met Office rule but with a larger amber area. Householder is happy to receive a higher magnitude warning.

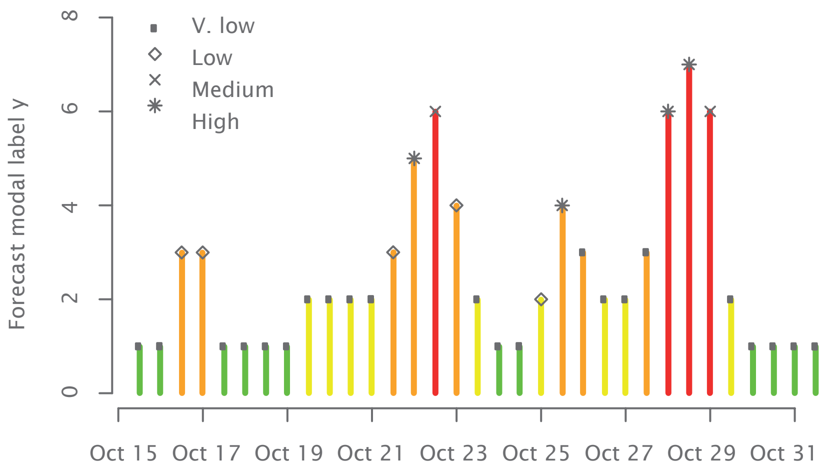

BAYES' WARNINGS, END-USER

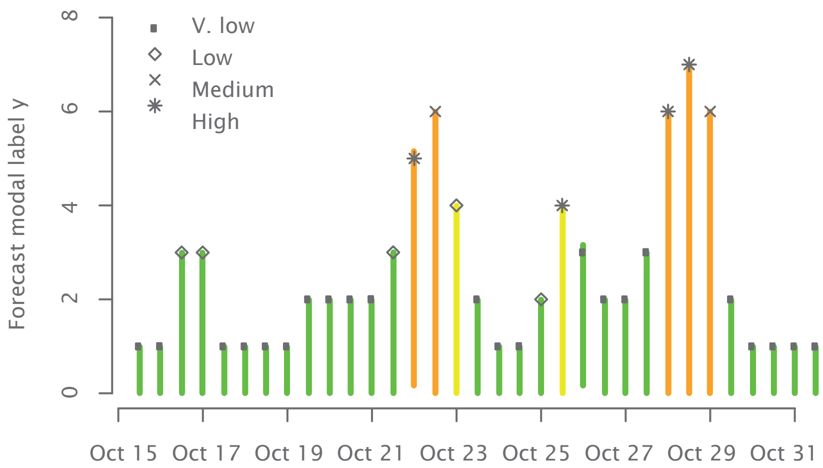

BAYES' WARNINGS, END-USER

The team

Theo Economou

University of Exeter
t.economou@exeter.ac.uk

David Stephenson

University of Exeter
d.b.stephenson@exeter.ac.uk

Ken Mylne

UK Met Office
ken.mylne@metoffice.gov.uk

Rob Neal

UK Met Office

Evaluating the importance of spatial complexity and uncertainty in predicting and managing floods

Floods are a major natural hazard that have affected millions of people throughout the world

The challenge

Accurately predicting current flooding is critical in helping to mitigate against the impacts that flooding may have in the future.

Computational hydraulic models are powerful tools that allow flood events to be simulated and mapped to understand areas that are at the greatest risk of being flooded. However, these models can take a long time to run, meaning that their use for real time decision making can be limited, especially when running multiple simulations to understand the uncertainty of the prediction.

One solution is to run models at a coarser spatial resolution, which reduces the run time of a simulation, but the effect of this alongside other uncertainties is unknown. This project applied novel approaches developed elsewhere in CREDIBLE to investigate how uncertainty and spatial complexity trade off against each other in making probabilistic flood inundation predictions.

What was achieved

Our results allow modellers to identify the simplest and most efficient model capable of achieving specific levels of accuracy. By using the simplest acceptable model users then have the time to undertake more testing and analysis, and to better characterize how likely errors in model input data will affect the quality of the predictions.

By applying our methods to the Carlisle 2005 flood event we were also able to demonstrate how complex hydraulic models can be utilised more quickly to allow their predictions to inform real time decision making.

How we did it

We took a state of the art flood inundation model developed at the University of Bristol and applied this to simulate a flood in the Imera River in Sicily for which we had high quality terrain data and observations of the maximum water height at various locations within the basin. We developed multiple models of this site at grid scales of 10m to 500m and looked at how the importance of different sources of error changed with grid resolution. We were able to show that for the very highest modelling scales increasing grid resolution gave no improvement in predictions when other errors were taken into account. In effect we showed that errors in model inflow meant that the apparent benefits of increased resolution were spurious beyond a particular threshold.

We used a new Global Sensitivity Analysis technique developed elsewhere in CREDIBLE to explore how influential spatial resolution and resampling of a fine scale Digital Elevation Model (DEM) when compared to uncertainties in the Manning's friction coefficient parameters and the inflow hydrograph providing the

boundary conditions. We found that the sensitivity of flood inundation models to these different factors was far more variable in space and time than was previously known.

In combination, these methods allow us to identify the simplest and most efficient model capable of achieving specific levels of accuracy taking all model errors into account.

Finally, we used a flood inundation model that was trained on a simulation of the Carlisle 2005 flood and built up a catalogue of possible flood events by running many more model simulations. We then used Bayesian Belief Networks to provide very quick estimations of flood hazard across Carlisle. When coupled with decision theory we showed how these probabilistic predictions could be used to identify whether roads at risk of flooding should be closed in advance of a flood event occurring. Similar approaches could enable flood inundation models to be more widely adopted for real time flood decision making.

References

Savage, J., Bates, P., Freer, J., Neal, J. and Aronica, G. (2016) When does spatial resolution become spurious in probabilistic flood inundation predictions? Hydrological Processes 30 (13), 2014-2032.

The team

James Savage

University of Bristol

Paul Bates

University of Bristol
Paul.Bates@bristol.ac.uk

Jim Freer

University of Bristol
Jim.Freer@bristol.ac.uk

Thea Hincks

University of Bristol

National scale drought risk and adaptation modelling for Great Britain: Infrastructure strategies to 2050

The challenge

Protecting drinking water, farming, energy and industry systems from drought in the context of increasing demand and a changing climate will require sustained investment. Even within a national picture where overall supply exceeds demand, heterogeneity of water availability in space and time leads to expensive and damaging water shortage.

In order to cope with this pressure, the natural systems that provide water supplies can be augmented through engineered infrastructure. The challenge for drought management in the 21st Century lies in the implementation of appropriate solutions that can meet changing requirements without compromising the environment and other users of water, or placing excessive financial burden on citizens.

Traditional storage and transfer technologies are required in parallel with new ideas for water supply and efficiency of water use. Establishing the appropriate combination of these options is a complex engineering and economic systems problem. With growing pressure on water supplies and major investment decisions ahead, a more strategic national approach is required.

What was achieved

We found that investment in new technologies and system efficiencies were most likely to result in a secure, low-carbon water resource system. Strategies which follow conventional engineering approaches, including those with a long term perspective, were environmentally more consequential and led to lower supply/demand margins by 2050 due to their implicit inefficiencies.

However, strategies based entirely on efficiency or new technologies could not meet demand across the country and therefore mixed strategies were required. Climate change and demographic scenarios have complex interaction, with dry climate scenarios leading to localised problems in the south and east, and population growth scenarios causing impact more broadly across England and Wales. Generally, vulnerability to drought is prevalent across Southern England, with the North of England and Wales at risk in more extreme scenarios.

How we did it

Our drought model relies on a simplified reconstruction of the existing regional water resource management arrangements in Britain. Mainland Great Britain is divided in 130 Water Resource Zones (WRZ), or Megazones in Scotland. Models of water availability of each WRZ determine values for potential water supply (deployable output). This includes river intakes, reservoir intakes and groundwater. Flow values are taken from the National River Flow Archive, with recorded flow scaled to the representative river intake flow or reservoir watershed on the basis of sub-catchment area. Reservoirs are modelled as a single storage with maximum capacity set to the total capacity for reservoirs in the WRZ.

Water demand for each WRZ is the sum of domestic and non-domestic demand. Domestic demand for a WRZ is determined by multiplying the projected population by the average per capita water demand for the WRZ. Non-domestic demand is calculated as a percentage of domestic demand for a given WRZ based on current region figures.

Alternative water supply options include desalination, aquifer recharge, demand reduction, effluent re-use, leakage reduction, inter-company transfers, and new ground water supplies. Investment and operating costs for each of these options are based

2050 WATER SUPPLY-DEMAND BALANCE (ML/DAY) IN A "NO BUILD" STRATEGY

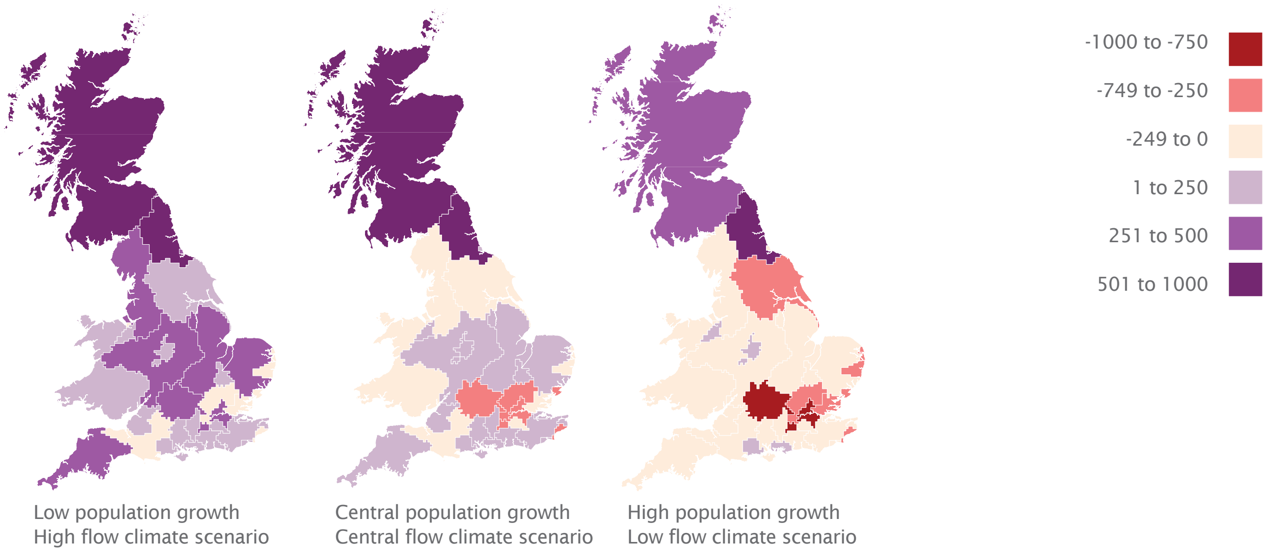

on a series of regression analyses of estimated costs against yields taken from 800 examples of water infrastructure projects. Climate change and population scenarios are used to modify supply and demand figures. Strategies are implemented using weighted preferences for water resource options.

References

Simpson, M., Ives, M. C., Hall, J., Kilsby, C., 2016, Water Supply Systems Assessment, In Hall, J., Tran, M., Hickford, A. and Nicholls, R. (eds.), The Future of National Infrastructure: A Systems-of-Systems Approach, Cambridge University Press: Cambridge

Simpson, M., James, R., Hall, J., Borgomeo, E., Ives, M., Almeida, S., Kingsborough, A., Economou, T., Stephenson, D., and Wagener, T. (Accepted for publication 2016), Decision Analysis for Management of Natural Hazards, Annual Review of Environment and Resources

The team

Mike Simpson

Environmental Change Institute, University of Oxford
michael.simpson@ouce.ox.ac.uk

Jim Hall

Environmental Change Institute, University of Oxford
jim.hall@eci.ox.ac.uk

Matt Ives

Environmental Change Institute, University of Oxford
matt.ives@ouce.ox.ac.uk

A tool for understanding and quantifying model uncertainty: Sensitivity Analysis For Everybody (SAFE)

The challenge

Our goal was to provide a suite of tools and methods to support uncertainty and risk assessment in natural hazards – to improve the transparency and defensibility of risk management decisions.

Good modelling practice requires an assessment of the confidence in a model. Predictions arising from numerical models are affected by potentially large uncertainties due to a range of factors, such as observation errors and uncertainty in model parameters and structure. Depending on the model application, such uncertainties can undermine the robustness and credibility of the modelling results to the extent that the usefulness of environmental models for supporting decision-making must be called into question.

Global Sensitivity Analysis (GSA) provides a valuable tool for both model developers and users to quantify the uncertainty in model outputs, estimate the relative contribution of the different input factors to such uncertainty, and thus prioritize efforts for its reduction.

Sensitivity analysis is recommended practice when model predictions are the substance of regulatory analysis or policy appraisal.

What was achieved

The SAFE tool allows those who use and develop models – in academia, government and industry – to investigate the potential for model simplification by identifying model components that have little impact on model behaviour and therefore can be omitted in a simplified version of the model. It also supports model calibration by identifying the parameters that most influence model accuracy (i.e. the ability to reproduce observations) and therefore need to be properly calibrated.

Further, SAFE supports model validation by checking consistency between the model response and our understanding of the system represented by the model, and model debugging by identifying combinations of input factors (model parameters, initial or boundary conditions, input data, etc.) that cause the model to fail. Finally, SAFE enables the user to identify the major sources of model uncertainty e.g. errors in input data or uncertainty in parameters, and thus to prioritize efforts for uncertainty reduction.

SAFE currently has over 300 academic users in a wide range of application areas. It is being trialled by industrial users, including catastrophe risk modelling companies, environmental consultants, technology and manufacturing companies.

“The SAFE toolbox, and the thinking behind it, are helping us to get more out of the models we use in assessing strategies for cost-effective river water quality monitoring, ultimately supporting decisions about catchment management to meet the requirements of the European Water Framework Directive”

Rob Lamb
JBA

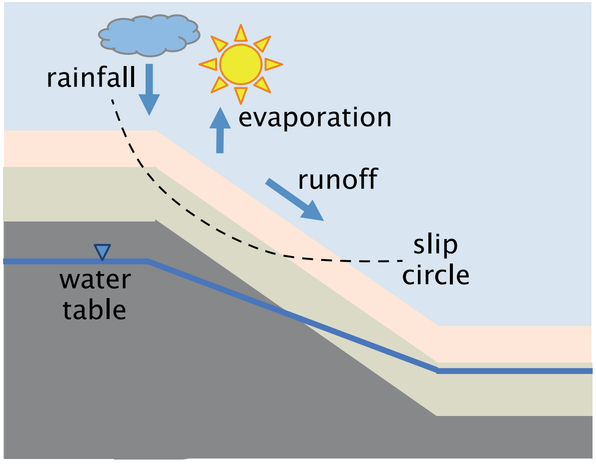

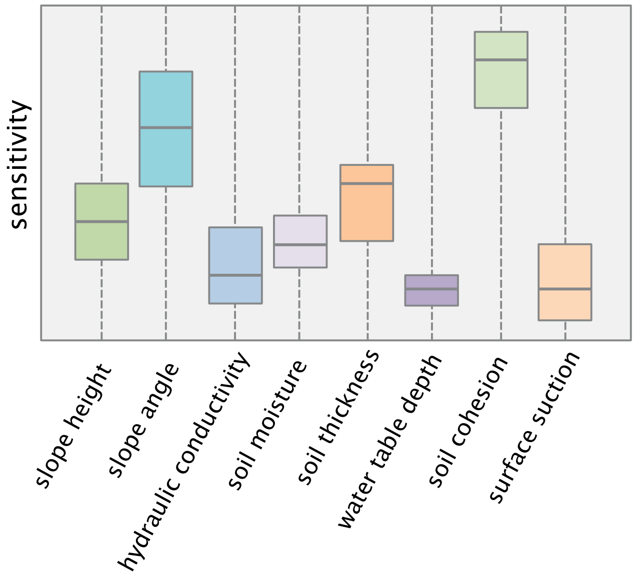

How we did it

The SAFE Toolbox provides a set of functions to perform Global Sensitivity Analysis in Matlab (or equivalently in the free software Octave) and R.

GSA methods implemented in SAFE:

- Elementary Effects Test (or method of Morris)
- Regional Sensitivity Analysis (RSA)
- Variance-Based Sensitivity Analysis (Sobol')
- Fourier Amplitude Sensitivity Test (FAST)
- DYNamic Identification Analysis (DYNIA)
- Density-based PAWN Sensitivity Analysis (Pianosi and Wagener, 2015)

The unique features of SAFE:

- Modular structure to facilitate interactions with other computing environments
- Set of functions to assess the robustness and convergence of sensitivity indices
- Several visualization tools to investigate and communicate GSA results
- Lots of comments in the code and workflow examples to get started

A general introduction to the rationale and architecture of SAFE is given in Pianosi et al. (2015).

A literature review of global sensitivity analysis techniques was undertaken. This formed the basis for the techniques incorporated in the SAFE software package. In performing the literature review it also became clear that there was a need for additional methods to be developed. Following from this, a novel approach to global sensitivity analysis was produced based on cumulative distribution functions. The method is named PAWN (Pianosi and Wagener, 2015).

References

More information is available from the SAFE webpage:
www.bris.ac.uk/cabot/resources/safe-toolbox/

Pianosi, F. and Wagener, T. (2015) A simple and efficient method for global sensitivity analysis based on cumulative distribution functions. Environmental Modelling & Software 67: 1-11.

Pianosi, F., Sarrazin, F. and Wagener, T. (2015) A Matlab toolbox for Global Sensitivity Analysis. Environmental Modelling & Software 70: 80-85.

Sarrazin, FJ, Pianosi, F & Wagener, T, 2016, 'Global Sensitivity Analysis of environmental models: Convergence and validation'. Environmental Modelling and Software, vol 79., pp. 135-152

Pianosi, F, Beven, K, Freer, JE, Hall, JW, Rougier, J, Stephenson, DB & Wagener, T, 2016, 'Sensitivity analysis of environmental models: A systematic review with practical workflow'. Environmental Modelling and Software, vol 79., pp. 214-232

The team

Francesca Pianosi

University of Bristol
Francesca.Pianosi@bristol.ac.uk

Thorsten Wagener

University of Bristol
Thorsten.Wagener@bristol.ac.uk

Tailloss: Evaluating the risk of insolvency by estimating the probability in the upper tail of the aggregate loss distribution

The main aim of this project was to evaluate the risk that annual loss will exceed the company's current operating capital

The challenge

One of the objectives in catastrophe modelling is to assess the probability distribution of losses for a specified period, such as a year.

From the point of view of an insurance company, the whole of the loss distribution is interesting, and valuable in determining insurance premiums. But the shape of the right hand tail is critical, because it impinges on the solvency of the company. A simple measure of the risk of insolvency is the probability that the annual loss will exceed the company's current operating capital. Ensuring an upper limit on this probability is one of the objectives of the EU Solvency II directive.

How we did it

We explored six different approaches to evaluate the annual loss distribution: Monte Carlo simulation and Panjer recursion, both of which are widely used but also very computationally expensive, and four upper bounds which are effectively costless to compute. We demonstrated the conditions under which the upper bounds are tight, and the appropriate size for Monte Carlo simulations.

We presented a numerical illustration of all the methods proposed using a large Event Loss Table (with 32,060 events) concerning US hurricane data, and used this challenging application to assess the various methods, carried out in our 'tailloss' package for the R computing environment.

What was achieved

Using results from Probability Theory, we derived four upper bounds and two approximations for the upper tail of the loss distribution that follows from an Event Loss Table. We argue that in many situations an upper bound on this probability is sufficient. For example: to satisfy the regulator, in a sensitivity analysis, or when there is supporting evidence that the bound is quite tight. Of the upper bounds we have considered, we find that the Moment bound offers the best blend of tightness and computational efficiency. In fact, the Moment bound is effectively costless to compute, based on the timings from our R package.

EXCEEDANCE PROBABILITY CURVE

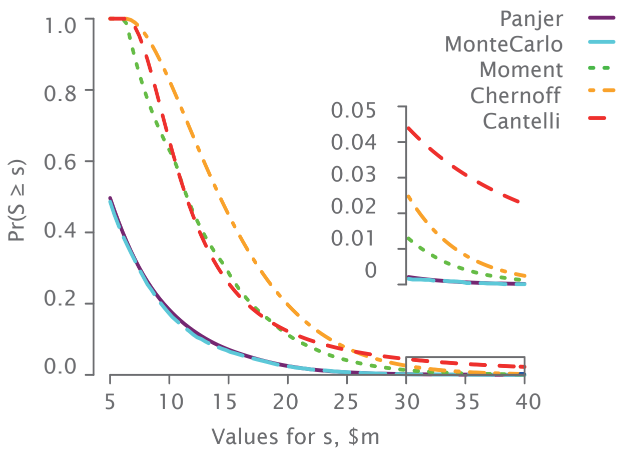

References

Gollini, I., and Rougier, J., (in press) “Rapidly Bounding the Exceedance Probabilities of High Aggregate Losses”, The Journal of Operational Risk. Preprint: http://arxiv.org/abs/1507.01853

Gollini, I., and Rougier, J., (2015) “Tailloss: An R package to Estimate the Probability in the Upper Tail of the Aggregate Loss Distribution” https://cran.r-project.org/web/packages/tailloss/index.html

http://www.risk.net/type/journal/source/journal-of-operational-risk

The team

Isabella Gollini*

Department of Economics, Mathematics and Statistics, Birkbeck, University of London
i.gollini@bbk.ac.uk

Jonathan Rougier

Department of Mathematics, University of Bristol
j.c.rougier@bristol.ac.uk

*This research was developed while Isabella Gollini was on a postdoctoral fellowship in the CREDIBLE project at the University of Bristol.

CREDIBLE project Uncertainty and Robustness Estimation toolbox (CURE)

The challenge

There is a general trend for increasing inclusion of uncertainty estimation (UE) in environmental and hazard modelling because the effective use of model results in decision making requires a level of confidence to be established.

Another requirement is the assessment of the implicit and explicit choices and assumptions made during the modelling and UE process (Rougier and Beven, 2013). In particular, assumptions made regarding the nature of epistemic uncertainties (uncertainties relating to a lack of knowledge), and how they are taken account of, should be recorded and communicated to stakeholders such that the meaning of the results and any subsequent analysis is put into context.

Good practice in this respect is still developing and is integral to the CREDIBLE Uncertainty and Robustness Estimation Toolbox (CURE), which aims to represent best practice in applying UE methods as well as best practice in being explicit about modelling choices and assumptions. In this way CURE will contribute to the ongoing development and testing of UE methods and good practice in their application.

What was achieved

The CURE toolbox consists of computer program functions which quantify uncertainties associated with simulation model results for a given application. It provides a broad range of different UE methods that users can apply that are demonstrated using various environmental model applications. Example applications are provided in the form of workflow scripts (sequences of computer program commands that link to CURE program functions) that demonstrate the different UE analysis methods and ultimately help users define and structure their own workflow scripts for their particular applications.

Apart from performing core calculations associated with individual UE methods, CURE functions compute additional modelling and UE diagnostics and results including visualisations that aid in the communication of uncertainty. The recording of modelling choices and assumptions, including treatment of epistemic uncertainties, is facilitated by a Graphical User Interface (GUI) where users can choose information for inclusion in a modelling audit trail log. The audit log is an important component for the communication of the meaning of the uncertainty estimates, as it sets the context for the UE.

How we did it

CURE is a fully open source toolbox written in the MATLABTM programming language. As it is aimed at simulation models, it employs a range of different Monte Carlo methods for forward UE (i.e. the forward propagation of uncertainties through a model: e.g. using prior estimates of input and parameter uncertainty) and conditioned UE (i.e. where UE results are conditioned on observations). The UE methods included span both formal statistical and informal approaches, underpinned by different philosophies, such that users are able to explore various approaches.

Each approach and modelling application will include its own implicit and explicit modelling choices and assumptions recorded via the GUI. The GUI takes the form of a number of simple, sequential dialogue boxes where the user is asked

AUDIT TRAIL
LOG FILE

WORKFLOW

Workflow	
Set up inputs and observations	Condition Tree GUI
Set up sampling ranges and distributions	Pre-Processing
Specify UE method and performance measure	Sampling Inputs/ Parameters Simulations
On-line and off-line processing	On-line Sampling/ Diagnostics
Diagnostics results visualisations	Post Processing

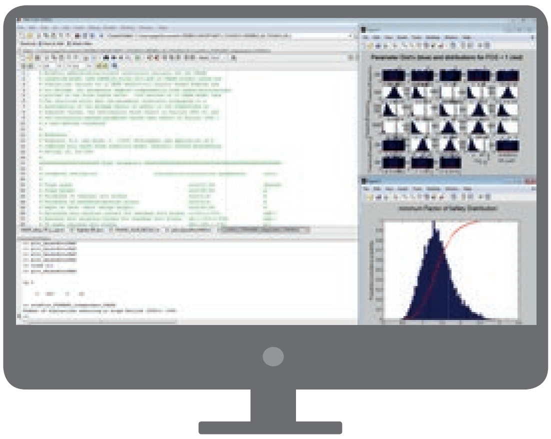

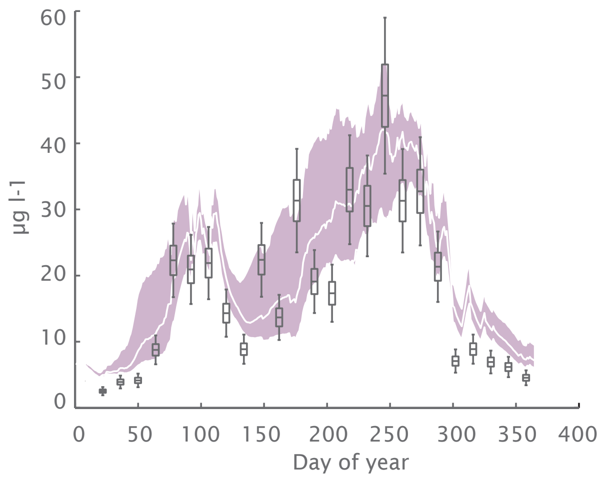

to enter information, as text, in a structured way and which can be iteratively edited during any modifications to analyses. The toolbox structure is such that new methods can be easily added and it will be subject to ongoing development and augmentation with additional workflow examples.

References

Rougier, J and Beven, K J, 2013. Model limitations: the sources and implications of epistemic uncertainty, in Rougier J, Sparks, S and Hill, L, Risk and uncertainty assessment for natural hazards, Cambridge University Press: Cambridge, UK, 40-63

NERC, PURE, CREDIBLE (Consortium on Risk in the Environment: Diagnostics, Integration, Benchmarking, Learning and Elicitation) Project: NE/J017450/1

The team

Trevor Page

Lancaster University
t.page@lancaster.ac.uk

Paul Smith, Keith Beven, Michael Hollaway

Lancaster University

Francesca Pianosi, Fanny Sarrazin, Suzanna Almeida, Liz Holcome, Jonty Rougier, Jim Freer and Thorsten Wagener

University of Bristol

CASE STUDIES:
RACER

Robust Assessment and Communication of Environmental Risk

RACER brought together nine different research institutions (Birkbeck College, British Geological Survey, Durham University, Edinburgh University, Fera, Imperial College, Oxford University, Reading University, UCL) and a range of different user organisations representing sectors including insurance and finance, government agencies and engineering. As well as expertise in the specific hazard areas of earthquake, extreme weather, flood, tsunami and volcanic ash dispersal, RACER included statisticians, psychologists and information designers to provide input into the quantification, communication and perception of uncertainties. It was led by Richard Chandler, Professor of Statistics at UCL. The RACER activities were focused around a collection of targeted case studies, each one aiming to address a specific problem involving risk assessment and communication relating to a particular hazard or class of hazards.

Highlights include:

- The development of an operational tsunami hazard assessment system for the northwest Pacific, running on the Oasis Loss Modelling Framework platform and supported by the (re)-insurance industry.

- Identification of what information will most effectively increase the accuracy of volcanic ash forecasts during an eruption, thereby supporting improved forecasting by providing clear and detailed guidance on what to measure.

- Development of methodology for combining different types of earthquake measurement onto a common scale, thereby providing access to longer and more consistent catalogues of earthquake data than have ever been available before.

- An assessment of the value of different weather forecasting products when considering forecasts of extreme wind speed up to two weeks ahead.

- Development of new tools and methods to make it possible to trace uncertainties throughout the entire data analysis workflow that generates flood risk maps and forecasting products.

Tracking uncertainties through flood risk assessments

We have developed methods to assimilate new types of information such as satellite-based precipitation observations, and to represent better the uncertainties related to imperfect understanding of hydrological processes

The challenge

Flood risk assessment and flood forecasting are still riddled with uncertainties.

One of the main problems is tracing the potential sources of error, which can be manifold: weather forecasts, computer models that convert rainfall into streamflow and flood extent, and vegetation and soil characteristics, among others, can all contain sources of error.

This becomes even more problematic under changing conditions such as climate change, increasing urbanization, and changes in the course of the rivers.

Tracing these sources of error and understanding how they propagate through the analytics workflow is therefore essential. At the same time, it is paramount to build predictive models that make best use of existing and new data sources, which are often very heterogeneous and error prone themselves, to reduce these uncertainties.

What was achieved

We have developed new tools and methods to make it possible to trace uncertainties throughout the entire data analysis workflow that generates flood risk maps and forecasting products. In particular, we have developed methods based on Bayesian statistics to assimilate new types of information such as satellite-based precipitation observations, and to represent better the uncertainties related to imperfect understanding of hydrological processes. For the latter, the use of multi-model ensembles is particularly promising.

They make it possible to represent hydrological processes in different ways, as "hypotheses" of processes that are not fully understood or cannot yet be quantified fully. Such frameworks are also an elegant method to incorporate uncertainties about future change, such as the impacts of global warming. Lastly, Bayesian statistical frameworks are very powerful for models to "learn" when new data become available, potentially rejecting some of the models within an ensemble and "promoting" models that are more consistent with the data.

How we did it

First we used new developments in cloud computing, web technologies, and "big data" to create more modular computer models. For this, we implemented existing flood models using emerging standards for data exchange and model coupling, such as those developed by the global Open Geospatial Consortium. In such a modular setup, individual components (for instance a data source or a particular algorithm) can be evaluated, compared, and replaced conveniently.

It also makes it possible to implement different simulation models that represent the same process in parallel, as a way to assess the impact of uncertainty in the model structure on the

final prediction. We then applied this model to a range of case studies, including the well-studied Plynlimon open air laboratory catchments, as well as a set of 135 catchments that are part of the UK's Future Flow hydrology dataset. Lastly, we used this experimental approach to test and refine specific methods to combine information from multiple model structure into a coherent flood risk estimation framework.

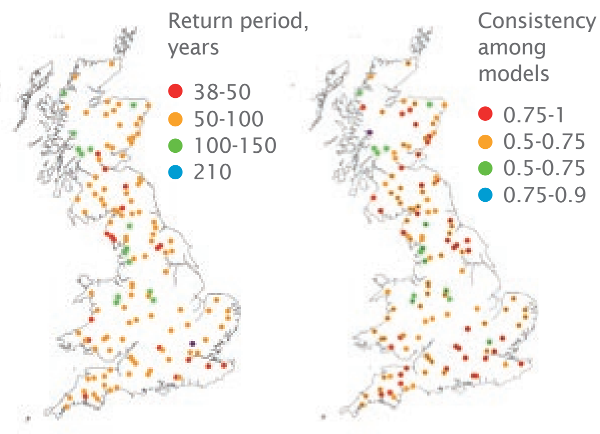

Figure 1: Typical (but flawed) predictions of change in flood risk for selected catchments in the UK using a multi-model ensemble (a) and agreement between the models within the ensemble, which is a good indicator for uncertainty (b). The typical protocol indicates a flood risk increase at the majority of considered catchments in the future.

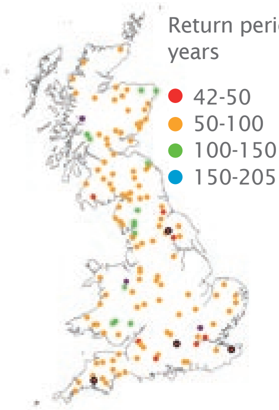

Figure 2: Predictions of change in flood risk for selected catchments in the UK using the newly developed multi-model combining protocol [Le Vine, 2016].

References

Almeida, S., Le Vine, N., McIntyre, N., Wagener, T., Buytaert, W., 2016. Accounting for dependencies in regionalized signatures for predictions in ungauged catchments. Hydrology and Earth System Sciences 20, 887-901

Vitolo, C., Elkhatib, Y., Reusser. D., Macleod C.J.A., Buytaert, W., 2015. Web Technologies for Environmental Big Data. Environmental Modelling & Software 63, 185–198

Le Vine, N. (2016). Combining information from multiple flood projections in a hierarchical Bayesian framework. Water Resources Research, pp. 1–18

The team

Wouter Buytaert w.buytaert@imperial.ac.uk
Nataliya Le Vine

Imperial College London

Claudia Vitolo ECMWF

Assessing and applying probabilistic forecasts of UK high windspeed

The forecasts produced by many national and international weather services offer skill up to 5-7 days ahead. However, "ensemble" forecasts are in general poorly calibrated in the sense that the spread of the ensembles does not represent uncertainty accurately

The challenge

The high windspeeds in storms which strike the UK and northwest Europe during winter are a substantial risk for the insurance industry and other business sectors. Between 1998 and 2013, wind damage was responsible for 45% of UK household weather-related claims compared to 30% attributable to flood.

Despite these impacts our knowledge is limited on how well high windspeeds can be forecast up to 10 days ahead, on how well such forecasts represent uncertainty, and on whether it is possible to use windspeed forecasts to usefully predict insurance losses with quantified uncertainty so that insurers can prepare for events before they happen.

The understanding of forecast performance is relevant to all forecast users and, of course, to the main PURE objective of offering improved assessment of uncertainty.

What was achieved

Focusing on high wind speed in northwestern Europe, we showed that the forecasts produced by many national and international weather services offer skill for lead times up to 5-7 days ahead.

However, "ensemble" forecasts – namely multiple forecasts issued in a way to represent uncertainty about the future – are in general poorly calibrated especially over land in the sense that the spread of the ensembles does not represent uncertainty accurately. It is important for forecast users to be aware of this. Based on performance for the period 2011-2014, of all the products considered the best-calibrated ensemble forecasts are those from the European Centre for Medium-Range Weather Forecasting.

We also developed residential and commercial damage curves of UK windspeed against insured loss for 10 UK regions, and showed that it is extremely challenging to predict insured loss from windspeed in the sense that the associated uncertainties are currently too large to be useful. The main obstacle to reducing these uncertainties is the relative scarcity of high-quality insured loss data.

How we did it

This challenge was undertaken using state-of-the-art data sets comprising the output from eight numerical weather prediction models with pan-European coverage, three wind speed verification data sets and UK insured loss data at 2-digit postcode resolution. The study period was March 2011 to February 2014.

A simple way to evaluate an ensemble forecasting system is to consider that if the uncertainty is represented accurately then, in the long run, the actual wind speed should fall in the top (or bottom) 50% of the ensemble range half of the time; in the top (or bottom) 25% of the range one quarter of the time; and so on. The Reliability Index (RI) uses this idea, and was one tool used to assess how well ensemble forecasts of European windspeed represented uncertainty – the RI will be close to zero if the forecast ensembles are well calibrated, and will be close to 100% if the observations always fall outside the range of the ensemble. The figure shows a typical example of how well a leading numerical weather prediction model represents uncertainty (in this case at a lead of 72 hours) in its windspeed forecasts. Much of Europe has RI values of at least 35-40 indicating that this model's ensemble forecasts of windspeed are poorly calibrated.

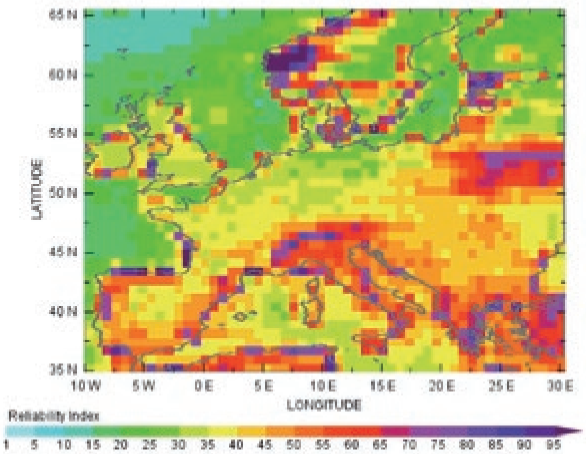

The team

Adam Lea
Mark Saunders
m.saunders@ucl.ac.uk
Richard Chandler

University College London

Probability and uncertainty in earthquake risk estimation

The aim of the earthquake strand of the PURE programme was to undertake a ‘root and branch’ review of some of the basic questions in seismic hazard analysis

The challenge

The UK has a low but not insignificant earthquake risk, as illustrated by the magnitude 4.3 tremor that hit Folkestone in 2007 and damaged hundreds of homes (see main photograph). Estimating hazard from a low seismicity area is difficult for a number of reasons, mainly associated with the small number of events on which to base the estimates.

A further challenge is that the measurement of earthquake magnitudes has changed over time so that it is hard to develop a long catalogue of consistent measurements. The aim of the earthquake strand of the PURE programme was to undertake a ‘root and branch’ review of some of the basic questions in seismic hazard analysis, paying particular attention to the quantification of uncertainties to ensure that risk can be estimated appropriately.

What was achieved

We made two major breakthroughs in the development of long and consistent earthquake catalogues by combining different types of observations. We solved the problem of how to use proxy data from the pre-digital era to derive modern physically-based estimates of earthquake characteristics, such as seismic moment (a measure of the amount of energy released by an earthquake). The most important sources of proxy data are seismograph traces, which are available back to the start of the 20th century; and, further back in time, anecdotal and newspaper reports of shaking and damage. We developed a rigorous statistical framework that links all of these data sources and incorporates knowledge of the underlying physical and measurement processes. The results can be used to determine a recalibrated catalogue for the UK (or indeed anywhere in the world) that accounts for all sources of uncertainty and all of the known constraints.

How we did it

We solved the problem of converting proxy data to modern physical measures by considering explicitly the mechanisms by which physical earthquake characteristics are translated into proxy measurements, and representing these mechanisms in the form of a graphical model (Figure 1). Probability calculus, using modern Bayesian computational machinery, then allows the relationship between the proxy and the underlying physical characteristics to be determined as a full probability density function. For the analysis of anecdotal and newspaper reports, we developed an innovative model based on the idea that damage severity decreases on average with distance from an event, so that the probability of a reported damage level exceeding some threshold can be mapped using roughly elliptical contours.
To constrain the results, we also incorporated

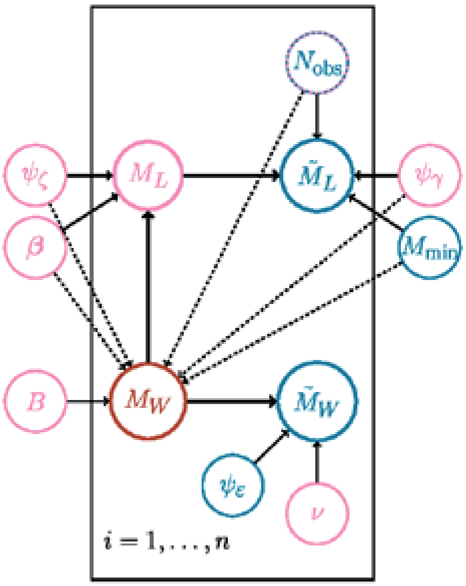

Figure 1 Graphical model showing the relationships between quantities involved in making proxy measurements of earthquake magnitude. Blue quantities are observed, other quantities are unobserved and the quantity of primary interest is denoted M_w

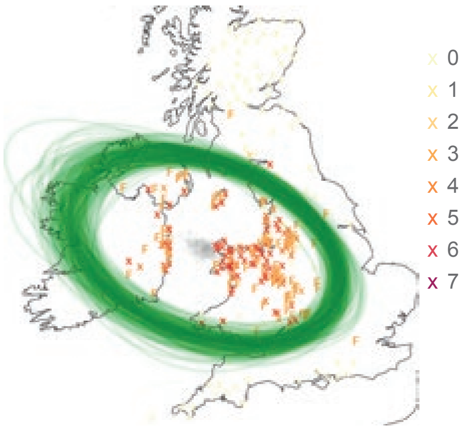

Figure 2 Elliptical contours defining the expected area in which damage / shaking of "intensity 3 and above" was experienced for the earthquake of November 9th 1852, based on anecdotal and newspaper reports. The blurring indicates uncertainty in the ellipse locations and in the earthquake epicentre (west of Anglesey).

information on the locations from which reports are likely to be received: for example, there are no reports from offshore locations, or from locations where nobody lives. Again, a Bayesian approach was adopted so that uncertainties could be incorporated explicitly into the results. An example is shown in Figure 2.

References

Touati, S., R.E. Chandler, I.G. Main, M. Naylor, R. Musson & S. Sargeant (2016). A Bayesian approach to calibrating the MW:ML relationship for British earthquakes, Geophys. J. Int. (in prep.)

The team

Ian Main (ian.main@ed.ac.uk), **Mark Naylor, Sarah Touati, Emily Kawabata**

University of Edinburgh

Roger Musson, Susanne Sargeant

BGS Edinburgh

Richard Chandler

UCL

Quantification of uncertainties in tsunami modelling

The uncertainties in tsunami wave impacts are enhanced when all relevant uncertainties in the input factors are considered together

The challenge

Tsunamis can cause catastrophic damage and loss of life. To understand tsunami risk at any location, complex computer simulations are used.

These simulations are subject to large uncertainties, which must be understood in order to avoid making costly wrong decisions. For example, tsunamis result from deformation of the sea bed, but the precise characteristics of seabed deformation cannot be known in advance.

To evaluate tsunami hazard therefore, it is necessary to explore the range of possible deformations, as well as uncertainties in other factors such as the bathymetry (i.e. the exact shape of the sea floor). Unfortunately, it is impossible to run enough simulations to explore fully the possible combinations of the uncertain factors. The challenge is thus to develop feasible methods for exploring these combinations, so as to obtain realistic assessments of uncertainty in tsunami hazard.

What was achieved

A fast approximation to one of the leading tsunami simulators (VOLNA[1]) was developed: this enabled us to explore the effects of combinations of uncertain factors much more quickly than before.

The computational efficiency was further improved by developing a new method for sequential design of computer experiments [2], which identifies the most informative settings for each simulation based on the results from earlier ones. A further refinement is to identify those factors that have little influence on the tsunami hazard, in order that these factors can be eliminated from the investigation so as to reduce the computational burden even further. Finally a method for the propagation of uncertainties in bathymetry [3] was investigated.

Some of these developments are being incorporated into an operational tsunami hazard assessment system for the region of Cascadia (NW Pacific), running on the Oasis Loss Modelling Framework platform and supported by the (re)-insurance industry [4].

How we did it

To develop our approximation to the VOLNA simulator, we used a technique called a statistical emulator (a similar technique was used for the volcanic ash study reported elsewhere in this brochure). Development of an emulator requires a modest number of simulator runs: these were carried out using state-of-the-art accelerated supercomputing facilities.

We then used geophysical principles to determine the possible variations in seabed deformation and other uncertain factors, coupled with statistical

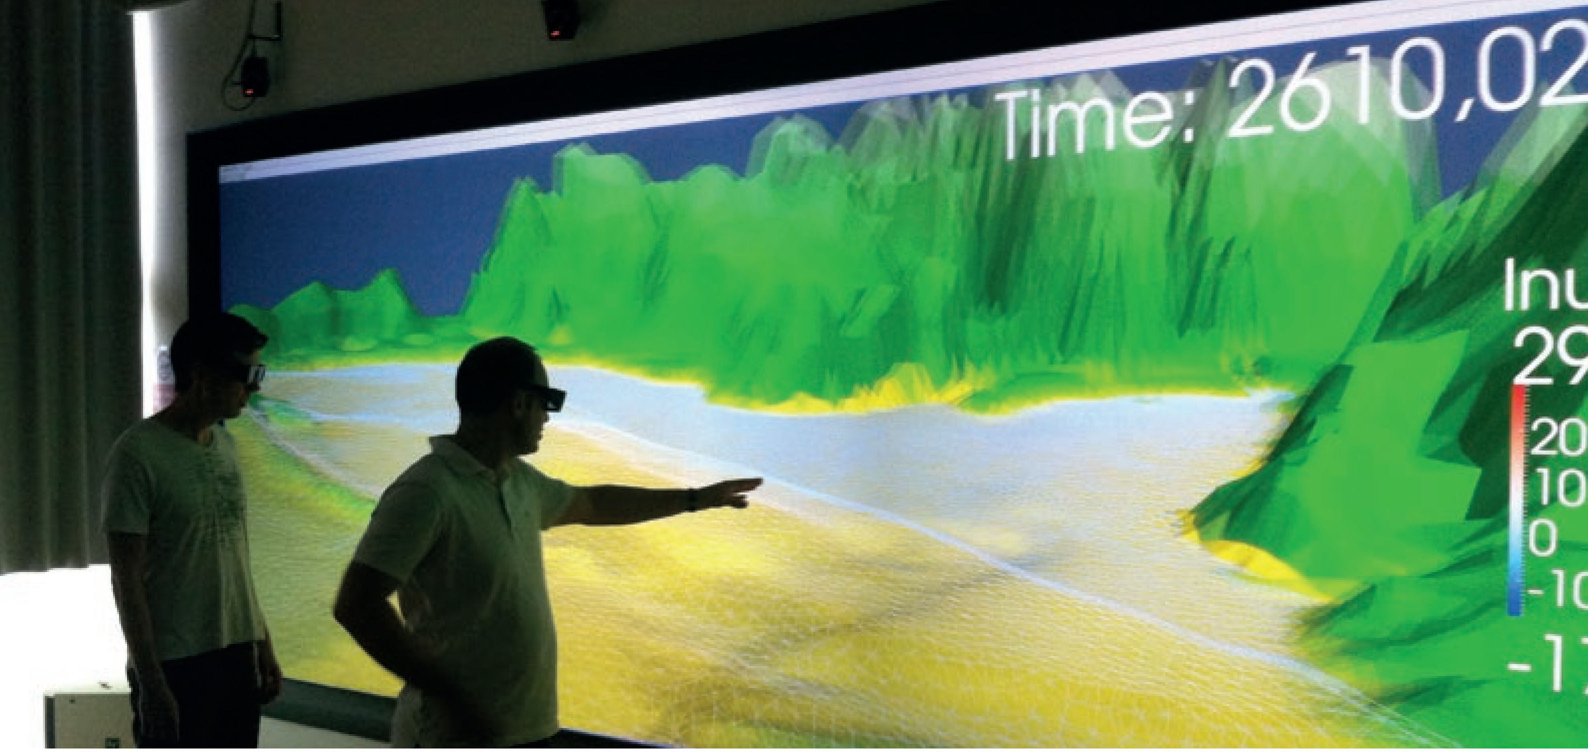

Figure 1: scientists studying a 3D visualisation of a tsunami simulation

expertise to find appropriate ways of representing these possible variations mathematically.

For the novel sequential designs of experiments we employed advanced mathematical approaches to prove near-optimality of the technique. We then used various test cases to illustrate the benefits of the technique in practice. For the VOLNA tsunami simulations, we found that the new design framework could save up to 20% in total computation time compared with existing approaches.

To account for uncertainties in the bathymetry, the main challenge is the sheer number of locations that must be considered. To deal with this, we used efficient dimension reduction techniques that are designed to reduce the size of the problem without compromising our ability to evaluate the tsunami wave impact at the coast. We illustrated the technique using a synthetic, but realistic test case. Unsurprisingly, we found that the uncertainties in tsunami wave impacts are enhanced when all relevant uncertainties in the input factors are considered together.

To help us to evaluate our work and communicate the results, we also developed visualization tools, and provided 3-D movies of simulations (see Figure 1): these enabled us to examine in detail the possible variations in tsunami flows and associated complexities.

References

[1] Dutykh, D., Poncet, R., & Dias, F. (2011). The VOLNA code for the numerical modeling of tsunami waves: Generation, propagation and inundation. European Journal of Mechanics-B/Fluids, 30(6), 598-615.

[2] Beck, J., and S. Guillas, (2016), Sequential Design with Mutual Information for Computer Experiments (MICE): Emulation of a Tsunami Model, SIAM/ASA J. Uncertainty Quantification, 4(1), 739–766.

[3] Liu, X., and S. Guillas, (2016), Dimension reduction for emulation: application to the influence of bathymetry on tsunami heights, submitted.

[4] Guillas, S., A. Sarri, S. J. Day, and F. Dias, (2016), Functional emulation, with an application to high resolution tsunami modelling over Cascadia, in revision, 2016.

The team

Serge Guillas, Simon Day, Joakim Beck

University College London
Contact: s.guillas@ucl.ac.uk

RACER Visualisation and communication

Current forecasts do not explicitly show uncertainty information. Studies suggest decision-makers infer their own level of uncertainty, with considerable individual differences

The challenge

The PURE vision to study risk and uncertainty in natural hazards would be incomplete if we did not consider how risk and uncertainty information is communicated to end-users.

We focussed specifically on people who use and create weather forecasts, including the public, decision makers in industry (e.g. insurers, energy companies, airlines, local flood managers, and first responders), forecasters, and academics. Our project addressed a number of questions relating to communicating risk and uncertainty: How can we help the producers of risk and uncertainty information communicate more effectively? Do users change their decisions if the same information is shown in different ways? Do users infer uncertainty if it is not explicitly communicated? Which method of communicating uncertainty is the most easily interpreted?

The answers to these questions can be used to provide guidance on how best to present uncertain information to different user groups.

What was achieved

We demonstrated that users' decisions are sensitive to the manner in which uncertainty is presented. For example, in a test case, we asked producers and users of volcanic ash forecasts for the aerospace industry to make a series of flight decisions based on different representations of fictitious conditions. We found that giving volcanic ash concentration and location information led to fewer unnecessary flight cancellations than giving just the location of an ash cloud, as was operational practice before 2012. Because current volcanic ash forecasts do not explicitly show uncertainty information, we were interested to see if users inferred uncertainty in the forecasts. Each respondent inferred their own level of uncertainty, with considerable individual differences.

We are now investigating how users respond to and interpret natural hazard forecasts where uncertainty information is explicitly shown. Our initial findings suggest that when users are given a median line on a graph explicitly showing uncertainty, their eyes are drawn to the median line and they are less aware of the extreme values. When no median line is shown, they are more aware of extreme values.

We have collated key points and guidelines into a leaflet, *Presenting data and uncertainty*, to help producers of natural hazard information communicate more effectively. This leaflet is freely available from the link given on the next page.

How we did it

We used stakeholder engagement as a springboard for our research, conducting initial interviews with industry users of risk and uncertainty information, in order to understand their process of decision-making and the information available to them.

We then conducted a series of focussed surveys with particular user groups, ranging across academia, insurers, energy companies, airlines, local flood managers, first responders, and the

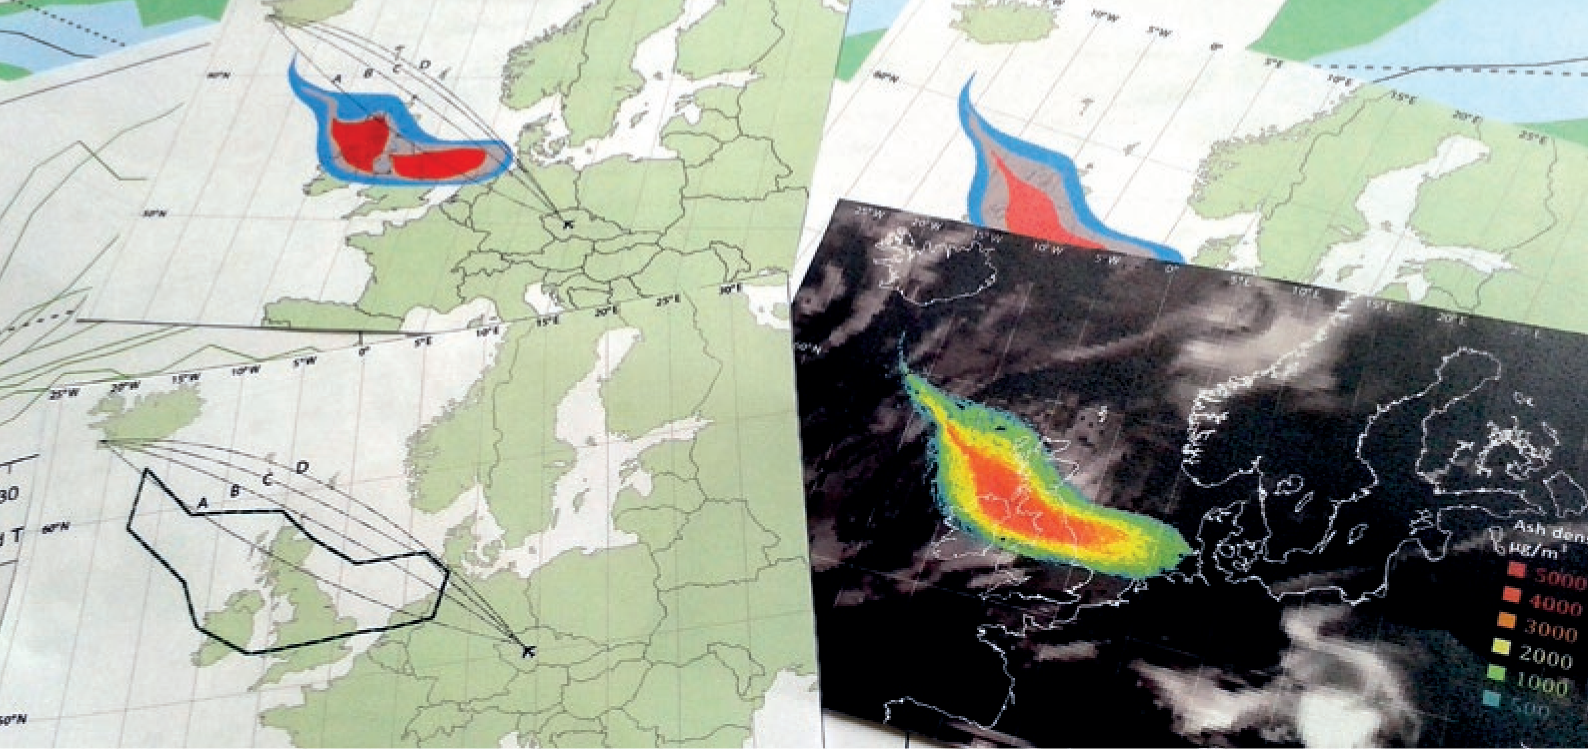

public, to understand how the characteristics of the most common forms of presentation – line graphs and maps – influence decision making. The surveys were designed as decision games, where respondents were given a hypothetical scenario and forecast that they used to make decisions. We varied the ways in which information was presented in the forecasts (see main photograph for examples of volcanic ash forecasts used in the study of flight decisions) and observed the impact of manner of presentation on respondents' decisions and confidence in their decisions. As our work continues, we will contribute to an understanding of how best to communicate uncertainty in order to support well-informed decisions that balance both business and safety factors.

Presenting data and uncertainty, available at https://connect.innovateuk.org/web/pure/quick-guide

References

Mulder, K. J., N. Harvey, M. Lickiss, A. Black, A. Charlton-Perez, H. Dacre, and R. McCloy: Visualizing volcanic ash forecasts: Scientist and stakeholder decisions for different designs and conflicting forecasts. Submitted to Weather, Climate and Society.

Mulder, K. J., M. Lickiss, A. Black, A. Charlton-Perez, and R. McCloy: Median Lines: Proceed with Caution. In preparation for Weather, Climate and Society.

The team

Kelsey Mulder, Andrew Charlton-Perez

Department of Meteorology, University of Reading
Contact: k.mulder@reading.ac.uk

Matthew Lickiss, Alison Black

Department of Typography & Graphic Communication, University of Reading

Rachel McCloy

Department of Psychology, University of Reading

Quantifying the uncertainty in volcanic ash forecasts

By identifying the factors contributing most to the overall uncertainty we were able to calculate confidence in volcanic ash forecasts

The challenge

Volcanic ash poses a significant hazard to aircraft: it can cause both temporary engine failure and permanent damage.

Flights are therefore restricted in ash-contaminated airspace, disrupting air traffic with the potential for large financial losses. For example, the 2010 Eyjafjallajökull eruption (see main photograph) grounded over 95,000 flights, costing the airline industry over £1billion. To ensure safe and optimised flight operations during volcanic eruptions, accurate forecasts of ash location and concentration are needed.

Current forecast models are, however, susceptible to many uncertainties, for example due to the imperfect representation of atmospheric processes and to incomplete knowledge of the volcano's eruption characteristics. At present, these uncertainties are not explicitly recognised in the operational forecast system: the challenge is to incorporate them to provide users (airline operators, regulators and air traffic managers) with more robust information to support decision-making during an eruption.

What was achieved

Through this project we have, for the first time, developed a way to understand the limitations of current ash forecast models and to account for the resulting uncertainties. This has enabled us to identify, from among the many possible options for improving the forecasts, which ones would be most effective in reducing uncertainty during a future eruption. For example, more accurate radar observations of volcano plume heights and new developments in the ability to measure the size of ash particles using satellite observations will reduce the uncertainty in ash forecasts.

We have also identified priority areas for model development. For example there is large model uncertainty associated with the representation of physical processes such as ash aggregation (ash particles sticking together) and turbulence (small-scale mixing). Future research in these priority areas will enable the UK to be more resilient to volcanic ash events in the future.

How we did it

It's difficult to perform a comprehensive uncertainty analysis with a volcanic ash model because there are so many uncertain factors to consider: as an example, Figure 1 shows how different forecast wind fields can affect the dispersion of ash particles. The obvious way to handle these uncertainties is to make multiple forecasts with different plausible values for each of these factors to provide a range of forecast scenarios. This is not feasible however, because the models used to produce the forecasts are complex and take a long time to run. We solved this problem using a novel statistical approach known as Bayes linear emulation. The key features of the approach are to combine expert opinion on which factors control the ash distribution with the output from our model, and to supplement the original forecasting model with a statistical approximation (an 'emulator') that can be run very fast.

As the emulators are fast, we could run them many times to find out which factors contribute most to the forecast uncertainty. We learnt that the most important factors are the height of the ash plume

and the size of the ash particles: therefore, better observations of these quantities should be prioritised in order to increase forecast confidence. We also found that uncertainties relating to turbulent mixing and ash particles sticking together (aggregation) were large: thus more work is needed to understand these processes. Finally, by identifying the factors contributing most of the overall uncertainty, we were able to go back to the original forecast model and design a small number of runs that fully capture the associated uncertainties, and hence calculate confidence in volcanic ash forecasts.

References

Dacre, H. F., A. L. M. Grant, N. J. Harvey, D. J. Thomson, H. N. Webster, and F. Marenco (2015), Volcanic ash layer depth: Processes and mechanisms, Geophys. Res. Lett., 42, 637–645, doi:10.1002/2014GL062454.

Harvey, N. J. and H. F. Dacre (2016) Spatial evaluation of volcanic ash forecasts using satellite observations, Atmos. Chem. Phys., 16, 861-872, doi:10.5194/acp-16-861-2016.

Dacre, H. F., N. J. Harvey, P. W. Webley, and D. Morton (2016), How accurate are volcanic ash simulations of the 2010 Eyjafjallajökull eruption?, J. Geophys. Res. Atmos., 121, 3534–3547, doi:10.1002/2015JD024265.

Harvey, N. J., N. Huntley, H. F. Dacre, M. Goldstein, D. Thomson and H. N. Webster (2016). Multi-level emulation of a volcanic ash transport and dispersion model to quantify sensitivity to uncertain parameters. Submitted to Natural Hazards and Earth Systems Sciences. Risk and Uncertainty Estimation Special Issue.

Huntley, N., Harvey, N. J., M. Goldstein, H. F. Dacre, D. Thomson and H. N. Webster (2016). TBC. In preparation to submit to Natural Hazards and Earth System Sciences. Risk and Uncertainty Estimation Special Issue.

The team

Helen Dacre, Natalie Harvey

University of Reading
Contact: h.f.dacre@reading.ac.uk

Nathan Huntley, Michael Goldstein

University of Durham

Helen Webster, David Thomson

Met Office

PARTICLE TRAJECTORIES 01Z 3/5/10 - 01Z 14/5/10

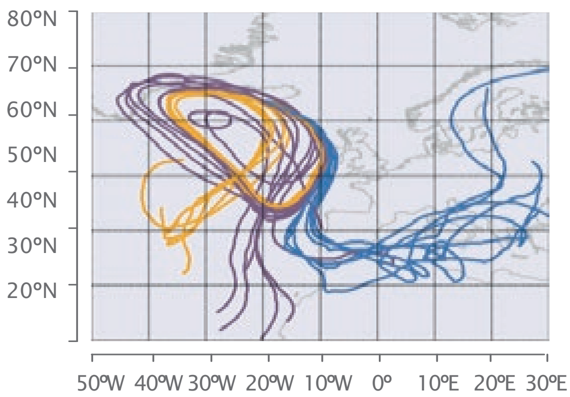

Figure 1 Model simulated trajectories of 20 volcanic ash particles emitted on 3 May 2010 from the Eyjafjallajökull volcano in Iceland, under 20 different plausible wind fields representing the meteorological uncertainty. Amber, purple and blue trajectories represent the spread in possible locations of ash particles, coloured according to their final location.

CASE STUDIES:

PURE ASSOCIATES

In 2013/14, the Natural Environment Research Council funded 20 PURE Associate projects through the PURE Network

It is vital that the knowledge and expertise arising from science is harnessed for the benefit of a wide range of end users. To facilitate this, NERC funded scientists from a number of leading UK Universities to work with businesses, policy makers and NGOs to apply their knowledge and skills to improve assessments of natural hazard risk and enhance decision making.

Each of the 20 projects which lasted for between three and six months, involved an academic being seconded to the user organisation to work on a specific problem or challenge. The programme was designed to be a fast-paced, flexible mechanism for knowledge exchange to generate impact. All case studies are available on www.pure-network.org and we chose to showcase several of them below to demonstrate the diversity of applications and impact.

Wildfires

Costs Droughts

Eruptions Geologic

Property Damage

Earthquackes

Cyclones

Losses

Vulnerable

Adverse

NATURAL

Storms

Floods

Event

DISASTERS

Tornadoes

Recover Tsunamis Damage

Meteorological

Catastrophes

Hurricane

Insurance Aid

Climatological

Geophysical

Property

Repair

Emergency

Loss of Life

Volcanic

Twisters

Earth

Weather

Can we improve heatwave and cold weather warnings?

Expanding and improving environmental heat and cold-stress forecasts

The challenge

There are clear links between extreme weather conditions and human health.

For example, during the 10 day, 2003 August heatwave, there were over 2,000 extra deaths compared to the same period in the previous year. The UK also has one of the highest levels of excess winter mortality in Europe. For example, during December to March 2013 there were 32,000 more deaths than the rest of the year.

Public Health England delivers warnings of extreme heatwave or cold conditions to the public with a range of different warning levels, aimed at reducing the risk to the public of these extreme events. The aim of this project was to investigate how the communication of these warnings could be improved using operational Met Office forecast systems.

What was achieved

We have proposed a revised health forecasting system for the UK which, in line with the National Severe Weather Warnings Service, produces colour-coded warnings of periods in which the impact of extreme weather on human health is predicted to be high. The revised system produces warnings based both on the potential impact of the extreme weather conditions and on their likelihood. By adopting this 'risk likelihood' approach we have been able to explicitly incorporate the uncertainty inherent in weather forecasts into the health warning system. The revised system also offers users a detailed view of regional variations in alert levels, to allow them to plan resources more effectively. Initially, users are presented with a UK map showing a single warning level for each of 11 regions. By clicking on a region with a given warning level, they are then able to examine variation in the alert level at much finer resolution.

We have tested the new system for three recent cases studies (two hot, one cold) which had significant health impacts. Discussion of these cases with users has suggested the additional information provided by the new system would be of operational benefit. To begin to test this, we have started a follow up project to implement the new system at the Met Office.

How we did it

Our re-design of the health forecasting system focussed on producing two risk matrices for heatwave and cold weather as shown in the figure on the right. The risk matrix is made up of thresholds for the impact of extreme weather on the x-axis and thresholds for the likelihood that those conditions will occur on the y-axis.

Thresholds for the impact part of the matrix are derived from papers in the epidemiological literature (Armstrong et al. 2011, Hajat 2013) which derive relationships between temperature and the mortality

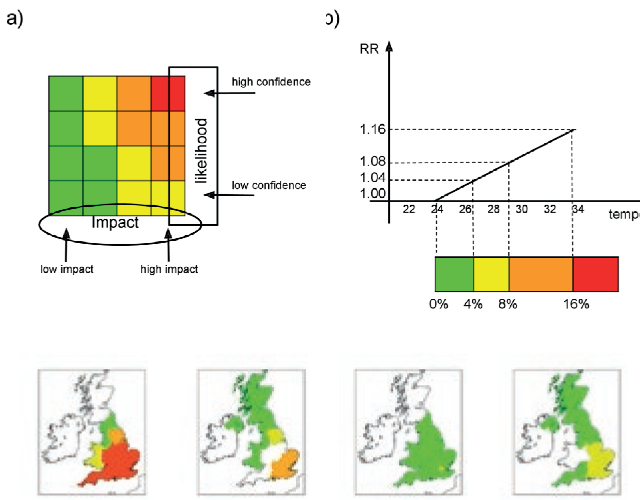

An example of the new system's output is shown here for July and August 2013, where the panels show forecasts, 1-4 days prior to the event.

risk ratio for 11 regions in the UK. The temperature thresholds for the risk matrix are then derived for fixed risk ratios from these studies as shown by the figure on the top right.

The likelihood axis is derived by considering the lead-time of the forecast for the case studies considered in this project (assuming that forecasts with longer lead time have less skill and hence lower likelihood), but we have also tested a case where a full probabilistic ensemble forecast is used to assign likelihood.

"We have found the NERC PURE Associates project immensely useful in exploring how we might increase the usefulness of hot and cold weather alerts, and so better protect health during adverse temperature."

Angie Bone
Public Health England

"This project has been a real personal success, it has been very instructive to see how the expertise I developed in my research career can be successfully applied (and be relevant) to a real world problem."

Giacomo Masato
University of Reading

The team

Giacomo Masato, Andrew Charlton-Perez

University of Reading

Virginia Murray, Angie Bone

Public Health England

The pest and disease risk to UK forest carbon projects

A model for standardised pest and disease risk assessment of UK forests for carbon buffer and insurance purposes

The challenge

Under the UK's Woodland Carbon Code (WCC), verified carbon sequestered in woodland can be sold as carbon credits.

The code includes provisions to ensure permanence and to protect the purchaser from potential recall of credits due to the loss of any woodland. A portion of carbon from each project is set aside in a pooled buffer to cover losses from a variety of causes e.g. wind, fire and pests and diseases (P&D).

Historic data is often used to estimate such risks but has little relevance to P&D as the highest risks often come from newly arrived P&D or established P&D expanding their range, which are not captured in historic data. This project was created to develop a new model to estimate the risk of P&D for such purposes and to see if it could be adapted for insurance purposes.

What was achieved

The project has:

- Generated greater awareness and discussion of a risk -based approach within the wider Forestry Commission/ Forest Research.
- Highlighted how current Pest Risk Assessments often do not provide the key information needed for risk assessments! ..and the key data needed.
- Provided a relatively simple tool to make sense of complexity i.e. capable of reducing 200+ P&D threats to a single risk factor per species.
- Provided a flexible scenario tool to support quantification and review of the adequacy of the WCC buffer.
- Confirmed that the key threats come towards the end of a WCC project's life and so can inform decisions on early release of buffer credits if withdrawals are low in the early years.
- Highlighted the main P&D threats to the carbon portfolio from a wide range of threats for further analysis.
- Potential to support other purposes e.g. broader plant health analyses, economic evaluations, insurance.

How we did it

The key species in the current Woodland Carbon Code portfolio (constituting >2%) were identified from an analysis of project documentation. A literature review and interviews with experts were used to identify the main pest and disease (P&D) threats to these species. A matrix was developed of the key metrics required to analyse risk and populated for a 25% sample with input from the Forestry Commission/Forest Research experts.

Example metrics were range, %mortality, % likelihood of entry and establishment. A scenario tool was developed to determine the worst case carbon loss for each P&D – a function of the year of arrival, age affected, rate of spread, and time for regrowth.
A model was then developed to combine this

information into a single risk factor for each key tree species according to which P&D affect them. This can then be used to assess individual project risk based on species weightings. Example data and outputs:

BIRCH	BETULA		Risk factor	Cumulative remainder
Bronze Birch Borer	Agrilus anxius	BETULA	3.57%	96.43%
	Anisogramma virgultorum	BETULA	0.34%	96.10%

Latin Name	English Name	Risk Factor	WCC	Weighted
ACER pseudoplatanus	SYCAMORE	5.4%	3.42%	0.2%
ALNUS	ALDER	5.4%	3.87%	0.2%
BETULA	BIRCH	9.4%	26.89%	2.5%
CORYLUS	HAZEL	5.4%	2.35%	0.1%
FRAXINUS	ASH	90.1%	8.53%	7.7%
PICEA sitchensis	SITKA SPRUCE	4.8%	5.40%	0.3%
PINUS sylvestris	SCOTS PINE	4.1%	10.86%	0.4%
PRUNUS avium ssp./var. avium	WILD CHERRY	3.2%	2.00%	0.1%
QUERCUS	OAK	2.8%	15.89%	0.4%
SALIX	WILLOW	5.4%	3.09%	0.2%
SORBUS	ROWAN	3.2%	4.19%	0.1%
			86.52%	12.3%

NB: Preliminary data for illustration ONLY. Worst case loss scenario: assume new P&D cause maximum potential loss and high likelihood of arrival: 1 in 20yr to 1 in 100yr.

The team

Susan Davies

University of Edinburgh

Pat Snowdon

Forestry Commission

Genevieve Patenaude

University of Edinburgh

"This project is enhancing our understanding of risk. Woodlands in the UK face a growing number of risks from pest and disease threats. Assessing these is essential in enabling us to develop resilient woodlands for the future and in underpinning important policy initiatives such as the Woodland Carbon Code."

Pat Snowdon
Forestry Commission

"The model was developed through an iterative process, combining expertise from both parties, and proved a valuable learning experience. The relationship will continue through a sister NERC project—the Forest Finance Risk Network—where the model will be exposed to a wider audience with potential for new insights and collaborations."

Susan Davies
University of Edinburgh

Flood forecasting challenges in Nepal

Operation probabilistic flood forecasting model of the Karnali river basin in Nepal

The challenge

Nepal is among the 20 most disaster-prone countries in the world.

In the past 30 years more than 4,400 deaths and over $10 billion of damage can be attributed to floods. Practical Action leads the development of community based early warning systems in Nepal. These benefit many thousands of vulnerable people allowing them to act to protect their lives and livelihoods. These systems use warnings based on observed data.

The steepness of many of the locations at risk often means these warnings cannot be given until very close to the flood occurring. Increasing the time before the event and the warning allows for a more proactive response to the flood risk. The integration of a robust operational flood forecasting system would allow this goal to be achieved.

What was achieved

This project has laid a basis for enabling flood vulnerable communities across Nepal to benefit from current best practices in flood forecasting. As part of the project, forecasting models have been developed for two basins and integrated into the existing, community based early warning systems. These are being operated by the Nepalese Government to assess the expected improvements in Early Warning, preparedness and risk reduction. Practical Action has been enabled by the provision of the knowledge and tools to replicate these case studies elsewhere in Nepal, and potentially in other countries (e.g. India, Bangladesh, Peru). A software toolbox suitable for both forecast model construction and real-time operation has been produced and made available to the Nepalese Government.

Training materials focusing on the use of probabilistic forecasts to complement community based early warning systems have been developed. This training will be rolled out by Practical Action to build national capacity on probabilistic forecasts in Nepal. The insights gained during this process will help drive research into future, novel forecasting methodologies.

How we did it

The outcomes of this project were achieved by close and careful cooperation of all the parties involved including the Nepalese Government. Only by utilising the complementary range of skills of the project team to consider the available data, forecasting methodologies and the usefulness of their outputs could the significant knowledge gaps be identified and solutions proposed. Forecasts are generated using a methodology based on physically interpretable time series models.

These are derived from the observed data. In real time operation they assimilate the most recent

observations to improve their probabilistic forecasts. They complement the existing real time data display systems currently used in Nepal for warning. To ensure the longevity of the system the software has been developed collaboratively in an open source framework.

To make the best use of the technological advances there was the need to upskill those involved in generating and utilising the forecasting. The training material developed focus on how good practice in operational flood forecasting from around the world can be given an appropriate local interpretation.

"The NERC PURE Associates project has allowed Practical Action Consulting to secure a partnership with Lancaster University, providing top quality UK technical expertise to tackle flood forecasting challenges in Nepal, enabling us to provide enhanced flood warning and protection for poor, flood vulnerable communities across Nepal."

Sarah Brown
Practical Action Consulting

The team

Paul Smith

Lancaster University

Sarah Brown

Practical Action Consulting

Keith Beven

Lancaster University

"Working with Practical Action Consulting to utilise cutting edge research has revealed many opportunities and scientific questions. Building close partnerships with end users offers a great opportunity to improve the quality and applicability of our research."

Paul Smith
Lancaster University

How frequent are hurricanes?

Incorporating parameter uncertainty in hurricane risk modelling

The challenge

RMS produces one of the leading models internationally for estimating future hurricane damages, based on the modelling of hazard, vulnerability and exposure.

In order for decision makers to manage the risks, there is a clear need to represent all sources of uncertainty in the model.

The objective of this project is to develop ways of incorporating parameter uncertainty into one component of the RMS hurricane model and assess the impact of that uncertainty on the estimated distributions of financial losses. In particular, the component of interest is to improve the frequency modelling of hurricanes in the Atlantic.

What was achieved

We have learned that neither the choice of probability distribution nor the estimation of its parameters represent a significant source of uncertainty when estimating the number of hurricanes in any particular year. However, these uncertainties become more important when they are propagated forward to derive estimates of financial loss.

We compared bootstrap and objective Bayesian methods for quantifying parameter uncertainty and found that they give very similar assessments. The project enabled RMS to improve upon its existing views of parameter uncertainty in its catastrophe models, and included exposure to new statistical techniques that could be used for modelling hurricane frequencies. Also, the methodology is generic, so potentially extending the scope over which this can be used to model other hazards as well.

The project allowed the associate and supervisor to tailor various statistical techniques in real applications in such a way that they can be integrated with the RMS hurricane model; this will prove useful well beyond this project. It is expected that this work will be submitted to peer-reviewed journals for publication.

How we did it

Existing models describe hurricane frequency with just a Poisson distribution, which captures the mean and variance of the number of hurricanes per year relatively well, but nothing more. However, in some cases the observed hurricane counts are more variable than would be expected under a Poisson model.

This led us to model the data using a Negative Binomial distribution which allows for overdispersion. Other models considered include the Conway-Maxwell-Poisson distribution and mixture models to allow for greater flexibility. Models were compared based on data visualization, model selection techniques, formal statistical tests and comparison of predictive distributions incorporating parameter uncertainty.

Parameter uncertainty was assessed separately using Bayesian schemes with objective priors, and bootstrap-based procedures. Objective priors for the Negative Binomial distribution were developed here.

The objective assessment of parameter uncertainty for the Conway-Maxwell-Poisson and mixture distribution seems too computationally intensive for routine use. Also, we provided some indications of how the different modelling choices and uncertainties influence loss calculations.

"This project has broken important new ground for us by developing practical ways to model parameter uncertainty in our hurricane model. The methods are general enough to allow application across our whole suite of mathematical models. Overall, that's a great outcome."

Steve Jewson
Risk Management Solutions

The team

Ken Liang, Richard Chandler

UCL

Steve Jewson

RMS

"I have greatly enjoyed my experience of working at RMS. This project provided me a fantastic opportunity to put into practice the statistical modelling skills I developed through my PhD."

Ken Liang
University College London

Dealing with the aftershock

A near real time aftershock forecasting tool for humanitarian risk assessment and emergency planning

The challenge

Concern Worldwide is an international humanitarian organisation that works with communities in the poorest countries of the world, including those facing significant earthquake risks.

We cannot predict earthquakes but we are increasingly able to give accurate near real-time forecasts of the intensity and spatial distribution of aftershocks.

This project will develop a set of software tools to assess aftershock hazard in near real-time during the emergency response phase and explore its potential to inform humanitarian emergency planning and response activities. In addition to developing and testing the tool, and developing guidelines and protocols for its application, the project will deliver a support programme to build humanitarian capacity in tool use.

What was achieved

For the first time it may be possible for aftershock hazard after a major earthquake to be integrated into humanitarian emergency planning and response processes. The University of Ulster Geophysics Group has conducted world class research on aftershock forecasting for many years and for this project developed a set of software tools to make a near real-time assessment of aftershock hazard during the emergency response phase of earthquake disaster.

Concern Worldwide and The University of Ulster Geophysics Group co-produced a training programme to teach humanitarian practitioners how to use the forecasting tool and a set of protocols and ‘rules of thumb’ to guide its use within Concern’s emergency planning and response processes.

They also explored the most appropriate user interface and ways of communicating risk and uncertainty. Together the partners explored the strengths and weaknesses of the tool (addressing for example the issue of data quality), and the potential mainstreaming of the approach (e.g. of integrating it into established humanitarian information systems).

How we did it

An initial workshop was held to develop project aims and structure. At this the University of Ulster clarified the potential and limitations of the proposed approach, and Concern Worldwide the needs and requirements of the humanitarian sector.

A phase of software tool development was conducted drawing on state of the art earthquake science, and creating the user interface and ways of communicating aftershock risk. A training programme and associated learning materials was

co-produced between project partners. This used synthetic earthquake data and the highest quality real-world data (e.g. from California) to illustrate key data quality issues and how they impact on the uncertainties in the aftershock forecasts. Central to this was the development of fictional scenarios to work through the strengths and limitations of the approach and how it might inform humanitarian practice.

A second workshop was held to deliver the training programme and explore initial guidelines and protocols for tool use and a website was developed to allow a wider network of humanitarian workers access to the learning materials, evaluate their usefulness and share their views. A final workshop was held to consolidate learning, critically evaluate the project and further explore the central issues of data quality, risk communication and the potential integration of the approach into humanitarian emergency response planning and response processes.

The team

Dominic Crowley, Dom Hunt

Concern Worldwide

Keira Quinn, Abigail Jimenez Lloret, John McCloskey, Max Hope, Paul Dunlop, Mairead NicBhloscaidh

University of Ulster

"Having seen the immediate aftermath of devastating earthquakes, it is clear that we do not know enough about the pattern of aftershocks and the risks they pose to vulnerable communities. This project, for the first time, gives us a tool that will help us make decisions about how and where to programme our responses."

Dominic Crowley
Concern Worldwide

"The project has created a unique opportunity for genuine knowledge exchange between science and humanitarian partners and to transform cutting edge earthquake science into something that could save peoples lives."

Keira Quinn & Abigail Jimenez Lloret
University of Ulster

WHAT THEY SAY

"The CREDIBLE project has been a tremendous success in bringing together experts from across academic and industry sectors to learn from each other about better ways to manage natural hazards. Through CREDIBLE we have been able to explore and test new approaches to dealing with risks caused by flooding and extreme weather, establishing productive collaborations that would not have grown without the project, and that will continue into the future."

JBA
- on CREDIBLE work

"Often, it feels there is a language barrier between the insurance industry and the academic community. Through our work with the CREDIBLE project, we have been incredibly pleased to be able to break down some of those barriers. Questions are being asked that would never have been thought of without the project. Working closely with the leading natural hazard scientists and statisticians involved, we are pushing proposals for future projects forward with increased energy and support."

XL Catlin
- on CREDIBLE

"PURE has provided a great opportunity to combine expertise in different aspects of environmental physics with expertise in statistics. From our point of view the combining of understanding of atmospheric physics, volcanology, satellite retrieval algorithms and statistical emulators has been particularly valuable in making progress on the problem of volcanic ash prediction."

Met Office
- on RACER's volcanic ash prediction work

"RNLI saves lives at sea in all weathers. The weather has an impact on the number of incidents we are called to, as well as the dangers our crews face. This tool [produced in the PURE Associate project] should help us be better prepared for busy times, and keep our crews safe."

RNLI
- on their PURE Associate project with LSE

"Education and raising awareness of flood risk is an integral part of the Scottish Government's approach to sustainable flood risk management. We're delighted to support Heriot Watt University in the development of such an innovative project to demonstrate in an enjoyable way that everyone can do something to reduce flood risk."

Scottish Government
- on their PURE Associate project with Heriot-Watt University

"Partnering with UEA really highlighted the benefits of engaging with academia. As a company heavily involved in research, we appreciate the opportunity to work on cutting edge climate risk issues and we look forward to implementing this knowledge through the delivery of products and services to our clients operating at the global scale."

Maplecroft
- on their PURE Associate project with UEA

ABOUT NERC

The Natural Environment Research Council (NERC) is the leading funder of independent research, training and innovation in environmental science in the UK

We invest public money in world-leading science, designed to help us sustain and benefit from our natural resources, predict and respond to natural hazards, and understand environmental change. We work closely with policymakers and industry to make sure our knowledge can support sustainable economic growth and wellbeing in the UK and around the world.

People around the world aspire to escape poverty and improve living standards: achieving this whilst living within the Earth's limits is a great challenge of the 21st century and NERC science has a critical role to play in meeting it.

We do this by investing in excellent, peer-reviewed science, often designed and delivered in partnership - with our academic community, business and the public and third sectors. Our research meets society's needs through discovery science - driven by curiosity - and strategic research that addresses some major challenges of the 21st century: how we can benefit from finite natural resources, build resilience to natural hazards and manage environmental change.

Our vision: To place environmental science at the heart of responsible management of our planet.

Probability, Uncertainty & Risk in the Environment
www.pure-network.org