Institution: University of Sheffield

Unit of Assessment: 10 - Mathematical Sciences

Title of case study: Managing uncertainty in computer models: aircraft engine design and food safety risk assessment

1. Summary of the impact

Pratt & Whitney (one of the world's largest makers of aircraft engines) has developed a process, "Design for Variation" (DFV), that uses Bayesian methods developed at Sheffield for analysing uncertainty in computer model predictions within the design, manufacture and service of aircraft engines. The DFV process significantly improves cost efficiency by increasing the time an engine stays operational on the wing of an aircraft, so reducing the time that the aircraft is unavailable due to engine maintenance. DFV also saves costs by identifying design and process features that have little impact on engine performance, but are expensive to maintain. Pratt & Whitney estimate the DFV process to generate savings, for a large fleet of military aircraft, of [text removed for publication].

The UK Food and Environment Research Agency (Fera) has used these methods in their risk analyses, for example in assessing risks of exposure to pesticides.

2. Underpinning research

The research is concerned with statistical methods for handling uncertainty in computer models. By "computer model", we mean a deterministic mathematical model of a physical system, implemented on a computer. Uncertainty can arise from not knowing the true values for the model inputs, and an imperfect model structure; we may not understand the underlying physics perfectly, and/or it may not be feasible to implement our best description of the physical process on a computer. Such models can take a long time to run. Reinman et al. [**S1**] describe computationally expensive models used at Pratt & Whitney: finite element models, as used for heat transfer and mechanical stress modelling, can take hours to run at one choice of input values, and computational fluid dynamics models based on Navier–Stokes equations can take days to run just once.

Kennedy and O'Hagan [**R1**] is concerned with calibrating computer models to data. They consider a computer model with two types of inputs: uncertain, fixed "calibration" inputs, and variable, known "control" inputs. Physical experiments are conducted at different values of the control inputs, and the aim is to find values of the calibration inputs so that the computer model outputs match the physical data as closely as possible. An important development is the inclusion of a "model discrepancy" function: a mechanism for learning the error in the model structure and correcting the model prediction at new input settings. The authors use a Bayesian framework, and show how to quantify all sources of uncertainty when predicting with the model.

This research was started at the University of Nottingham by O'Hagan as Principal Investigator and Kennedy as Research Associate, funded by an EPSRC grant with support from the National Radiological Protection Board [**G1**]. O'Hagan and Kennedy moved to the University of Sheffield in January 1999, where they continued and developed their research, leading to their 2001 publication.

Oakley and O'Hagan [**R2**] is concerned with identifying the most influential inputs in a computer model. They consider a computer model with multiple inputs, the true values of which are uncertain. The aim is to quantify how each uncertain input contributes to the uncertainty in the model output. A variance-based approach was taken, which quantifies how much the output variance can be reduced by learning the true value of an input.

The key development is to use a Gaussian process emulator to speed up computation of the variance-based measures of input influence. The emulator is a statistical approximation of the computer model, which is constructed from a relatively small number of computer model runs, and can then be used as a fast surrogate. Previous computational methods required large numbers of runs of the computer model, and were infeasible for computationally expensive models. (Faster computers have not solved this problem. As computing power increases, model users may choose to run their models at higher resolution, improving accuracy, but requiring more computational effort).

This research was conducted at the University of Sheffield, funded by an EPSRC grant with O'Hagan as Principal Investigator and Oakley as Research Associate (now Lecturer at the University of Sheffield) [**G2**].

3. References to the research

Papers:

R1 Kennedy, M.C., O'Hagan, A. (2001). Bayesian calibration of computer models (with discussion). *Journal of the Royal Statistical Society,* Series B 63, 425-64. doi: 10.1111/1467-9868.00294

R2 Oakley, J.E., O'Hagan, A. (2004). Probabilistic sensitivity analysis of complex models: a Bayesian approach. *Journal of the Royal Statistical Society,* Series B 66, 751-69. doi: 10.1111/j.1467-9868.2004.05304.x

Grants:

G1 Engineering and Physical Sciences Research Council, £105,685 (1995–98) Bayesian uncertainty analysis and computer model inadequacy, with support from the National Radiological Protection Board. PI: Anthony O'Hagan

G2 Engineering and Physical Sciences Research Council, £86,940 (2000–02). Realising Our Potential Award: Bayesian elicitation of expert opinion. PI: Anthony O'Hagan.

4. Details of the impact

The Sheffield research has changed the way Pratt & Whitney designs and manufactures aircraft engines, and the way the Food and Environment Research Agency (Fera), part of the Department for Environment, Food and Rural Affairs (Defra) conducts risk assessments.

Commercial impact

The Design for Variation initiative was led by Grant Reinman, a statistician at Pratt & Whitney. Pratt & Whitney learned of Kennedy and O'Hagan's work from a literature search. A key dissemination route for Oakley & O'Hagan (2004) was the software package GEM-SA, which implements the methodology in this paper, written by Marc Kennedy during his time at Sheffield. It was made available for free download, and Pratt & Whitney have used it in their design processes (though they have now built on it to develop their own software). Kennedy is now a risk analyst at Fera, and so disseminated the research within Fera directly.

Pratt & Whitney's Design for Variation process has five steps: (i) define probabilistic design criteria; (ii) use computer models and physical experiments to identify causes of performance variation and uncertainty; (iii) find the optimum design to satisfy the design criteria; continue data collection to (iv) validate the models; and (v) ensure the models remain consistent with the real world. Methods in Oakley & O'Hagan (2004) and Kennedy & O'Hagan (2001) play an essential role in steps (ii) and (iii), and hence have contributed to what Al Brockett, a former vice president of engineering module centres at Pratt & Whitney, describes as a "*paradigm shift*" and a "*high-visibility strategic priority*" in the way they design and manufacture aircraft engines [**S1**].

An illustration is given in Reinman et al. [**S2**]. In the design of a jet engine turbine airfoil, a computer model predicted the life expectancy of the airfoil, given its design. There was variability in airfoil life expectancy due to part-to-part variation, engine-to-engine variation, and environmental variation, and the designers wanted to know how to reduce variation in life expectancy. A variance-based sensitivity analysis was used: the analysis told them how much of the output variance was caused by each source of input variation. As the model was computationally expensive, the analysis could not have been done without Oakley & O'Hagan (2004). The designers used the results to assess the most cost-effective way of reducing variability in life expectancy, by targeting the most important sources of input variation (and not wasting resources by reducing unimportant input variation).

Pratt & Whitney calibrate their computer models to data using Kennedy & O'Hagan (2001). This method allows them to account for all sources of uncertainty in their model predictions – in particular, uncertainty due to a model not representing reality perfectly. Reinman et al. explain the benefits: "*Significant insight can be gained from the calibration results. In a recent study, assumptions typically made about boundary conditions near the airfoil surfaces were found to be over 20% higher than what the calibration process revealed them to be. Part temperatures were being over-estimated, and correspondingly airfoil life was being under-estimated*" [**S1**].

To quantify the financial benefits of DFV, Pratt & Whitney did a Business Case Study to assess the value of quantifying and managing uncertainty over the entire life cycle of an engine (from design through to service), using sensitivity analysis and calibration methods within their DFV process. The published saving in sustainment costs from doing this, for a large fleet of military aircraft, was approximately [text removed for publication] [**S3**]. The company also estimates that its component-level DFV initiatives "*have yielded a 64% to 88% return on investment by reducing design iterations, improving manufacturability, increasing reliability, improving on-time deliveries, and providing other performance benefits*" [**S1**].

Change to professional practice in environmental management

The Sheffield research has also changed the way the Food and Environmental Research Agency (Fera) conduct probabilistic risk assessments. Kennedy et al. [**S4**] report an analysis funded by the UK Health and Safety Executive's Chemicals Regulation Directorate (Defra project no. PS2005), investigating risks of exposure to pesticide from the spray drift of an agricultural boom sprayer. A computer model predicted the level of exposure to bystanders and residents after a crop-spraying event. The model had uncertain and variable inputs, such as the height of the boom, distance of a bystander from the source, wind speed, etc. Using the sensitivity analysis method of Oakley & O'Hagan (2004), implemented in GEM-SA, they quantified the contribution of each uncertain/variable input to the output uncertainty, to give risk managers information on how best to manage risks by reducing output uncertainty. Due to the computational expense of the model, this would not have been feasible without Oakley & O'Hagan (2004). The analysis in this case suggested reducing boom height and variation in boom height has the potential to reduce exposure.

Other ongoing projects at Fera are using GEM-SA for contaminated land and assessing the impact of recycling pesticide containers: Defra research project PS1010 – Development of Category 4 Screening Levels for Assessment of Land Affected by Contamination; and Defra research project PS2808 – Recycling of Home and Garden Pesticide Containers.

5. Sources to corroborate the impact

S1 ANSYS (2013). ANSYS Advantage, 7, 2, p. 18. Available at http://tinyurl.com/q63h646

S2 Reinman, G. et al. (2012). Design for variation. *Quality Engineering*, 24: 317–45 doi: 10.1080/08982112.2012.651973

S3 Statistician, Pratt & Whitney, letter on file corroborating savings in sustainment costs.

S4 Kennedy M.C., Butler Ellis, M.C., Miller, P.C.H. (2012). BREAM: A probabilistic bystander and resident exposure assessment model of spray drift from an agricultural boom sprayer. *Computers and Electronics in Agriculture*, 88: 63-71.