European Research Council

ERC Advanced Grant 2013
Research proposal [Part B1]
(to be evaluated in Step 1)

Professor Norman Fenton

Queen Mary University of London

Effective Bayesian Modelling with Knowledge before Data

Short Name: BAYES-KNOWLEDGE

Proposal summary

This project aims to improve evidence-based decision-making. What makes it radical is that it plans to do this in situations (common for critical risk assessment problems) where there is little or even no data, and hence where traditional statistics cannot be used. To address this problem Bayesian analysis, which enables domain experts to supplement observed data with subjective probabilities, is normally used. As real-world problems typically involve multiple uncertain variables, Bayesian analysis is extended using a technique called Bayesian networks (BNs). But, despite many great benefits, BNs have been under-exploited, especially in areas where they offer the greatest potential for improvements (law, medicine and systems engineering). This is mainly because of widespread resistance to relying on subjective knowledge. To address this problem much current research assumes sufficient data are available to make the expert's input minimal or even redundant; with such data it may be possible to 'learn' the underlying BN model. But this approach offers nothing when there is limited or no data. Even when 'big' data are available the resulting models may be superficially objective but fundamentally flawed as they fail to capture the underlying causal structure that only expert knowledge can provide.

Our solution is to develop a method to systemize the way expert driven causal BN models can be built and used effectively either in the absence of data or as a means of determining what future data is really required. The method involves a new way of framing problems and extensions to BN theory, notation and tools. Working with relevant domain experts, along with cognitive psychologists, our methods will be developed and tested experimentally on real-world critical decision-problems in medicine, law, forensics, and transport. As the work complements current data-driven approaches, it will lead to improved BN modelling both when there is extensive data as well as none.

Section a: ***Extended Synopsis of the project proposal (max. 5 pages)***

a1. Aim

The aim of this project is to help doctors, lawyers, engineers and strategists who wish to properly quantify uncertainty to improve decision-making and risk assessment. A secondary aim is to reduce the enormous waste associated with studies and systems that place their trust in data-driven methods of risk assessment and prediction, but which are flawed as a result of failure to properly incorporate the knowledge of domain experts about the underlying problem. The proposal is radical in that it plans to produce risk quantification models in situations (common for critical problems in medicine, law, and systems engineering) where there is little or even no data, and hence where traditional statistics cannot be used. To help non-statisticians develop causal probabilistic models of the underlying problem without data requires us to solve a range of complex technical, empirical, and cultural challenges. These require major interdisciplinary collaboration, cases studies and experiments. This is risky and complex, which is why other well-qualified researchers have largely avoided the problem, despite its known importance. The Research Fellow is uniquely well qualified to tackle the challenges, given his: proven record of research and development in the technical subject area (Bayesian causal models), practical experience in building real-world models with medics, lawyers and engineers, interdisciplinary contacts (academic and commercial), and extensive record of outreach.

a2. Background and Motivation

The focus is on decision-making and prediction in situations involving multiple related uncertain variables. For example, in determining the most suitable treatment for a patient, a doctor must combine multiple types of uncertain information about: different symptoms they observe, the reliability of different test results, and the potential side effects that different treatments may cause. Professionals are often no better than lay people in their ability to produce rational decisions when confronted with such evidence [20]. Even worse is the tendency of decision-makers to ignore uncertainty completely in their decisions, often resulting in over-confidence in the wrong result [9]. These examples seem to suggest that an expert's subjective judgement cannot be trusted. But we must distinguish between two fundamentally different types of expert judgement:

1. That used (as above) to combine multiple pieces of uncertain information in one's head in order to arrive at a conclusion. We can consider this 'mental probability analysis' rather than expert judgement.
2. That used merely as inputs to the problem (clinician's understanding of the false positive rate of a particular test, probability a particular treatment can cause a particular side effect, etc). This kind of 'prior information' is more accurately referred to as 'expert knowledge' (we are especially interested in quantitative aspects of such knowledge, but purely qualitative knowledge as in [6] may also be relevant).

We should not – and indeed need not - trust an expert with respect to 1 because (as we explain in a.3) there are tools which perform these computations far better then humans. However, it is unnecessarily constraining not to trust an expert with respect to 2, especially if the objective is to quantify uncertainty.

Nevertheless, implicit distrust of 'expert judgement' has created a push towards data-driven decision models which, due to advances in statistical machine learning, require little or no expert involvement in their construction or application (the 'big data' movement [109_E][1] has seriously raised expectations in this respect). But, for many critical risk analysis problems decisions must be made where there is little or no direct historical data to draw upon. Clearly, in such situations we either need to exploit expert judgement or make decisions randomly. This point is increasingly widely understood [15],[19]. However, what is less well understood is that, for most real-world risk assessment problems, ***even when large volumes of data exist***, data-driven machine learning methods alone are unlikely to universally provide the insights required for improved decision-making [29_E]. Indeed, failure to incorporate prior expert knowledge will often result in a data-driven model that is superficially 'objective' but fundamentally flawed and incapable of making accurate and useful predictions. Suppose, for example, we have a large dataset of patients who entered hospital with head injuries. We have two types of data for each patient. The 'inputs' (collected on arrival at hospital, such as age, pupil dilation etc – shown in Fig 1a) and the 'outcome' ("OK" or "death/permanent brain damage"). In order to identify future cases requiring most urgent treatment we want to 'learn' which profile of 'input' data is most likely to result in the worst outcome. Standard statistical learning will typically produce a regression based model, or similar, such as that in Fig 1a, while Fig 1b is the 'causal' model learnt from a structural machine learning algorithm in a real case study [95_E]. Neither of these models take account of *actual treatment* that did take place for the patients and which *could* take place for future patients; nor do they take account of what might have happened if the treatment was different (this is called a *counterfactual* problem [15]). They also fail to recognize the difference between causal factors affecting the seriousness' of

[1] All references with subscript E can be found on www.eecs.qmul.ac.uk/~norman/projects/B_Knowledge.html

injury (such as delay in arrival) compared with those that result from its measurement (such as pupil dilation). They will therefore be unable to predict accurately from profile data which patients most urgently need treatment and how the chance of survival is reduced if treatment is delayed. In contrast, the expert causal model in Fig1(c) addresses all of these concerns by introducing the necessary hidden interventions and explanatory factors and distinguishing between cause and effect.

Similar problems - where purely data-driven 'solutions' miss explanatory or intervention information - occur not just in medicine, but in many areas of the social sciences, finance, forensic sciences, transport, and policy making. Such data-driven modelling is typically predicated on 'what data is available', rather than on 'what data and knowledge is required'. It enhances convenience but at the cost of accuracy. By accepting the data-driven model we are asked to defy our senses and experience and actively ignore the role unobserved factors might play. Although the problem is well known to statisticians, the standard proposed solution – to perform *randomised controlled trials* – is neither possible nor even feasible for most critical risk assessment problems [9$_E$, 14$_E$][19]. For these problems, it is often assumed that lack of data means that no probabilistic models can be used - forcing decision makers to use only gut instinct. The history of errors made by experts in such situations warns us of the dangers of such capitulation.

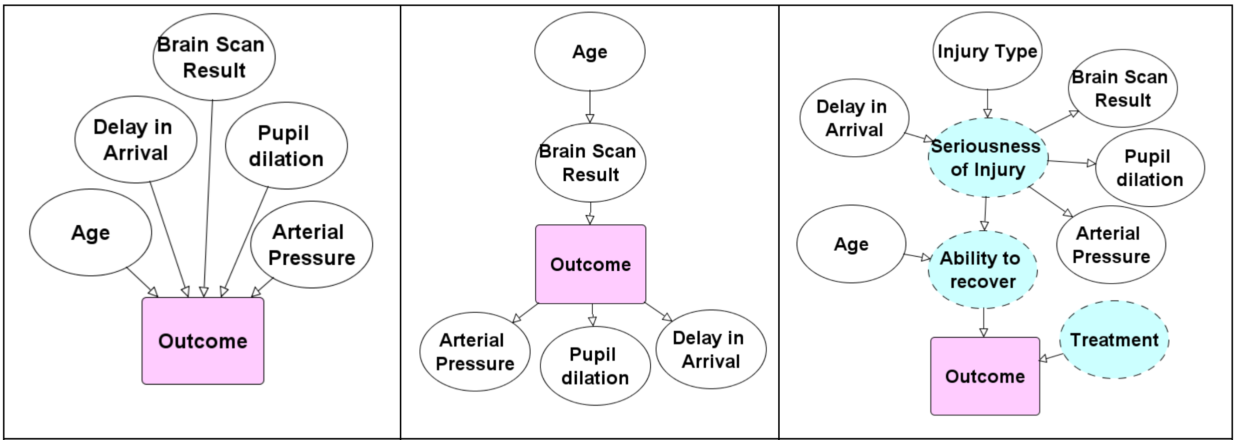

(a) Statistical regression models learnt form data: Outcome=*f*(inputs)	(b) 'Causal' model learnt purely from data	(c) Sensible causal model with missing/unobserved variables

Figure 1 The problem with data-driven models

a.3. State-of-the-art: the Bayesian approach

Fortunately, Bayes' Theorem [12] enables us to supplement observations with prior judgements to make more informed decisions with less data. It is the most widely accepted method of incorporating knowledge for probability assessment and it provides the only rational and consistent way to solve the problem of how to update a belief in some uncertain event (such as whether a patient has a particular disease) when you observe new evidence about the event (such as a symptom or positive test result). Unfortunately, the main rationale for using Bayes' – to exploit expert knowledge – is also at the root of most objections to its use; many decision-makers are reluctant to accept a probability based on subjective judgement [7]. Another impediment to using Bayes' is the complexity of the calculations. Even in the simplest case of just two variables (hypothesis and evidence) people without mathematical training find it difficult to compute and understand the result (indeed, numerous probabilistic fallacies of legal and medical reasoning would be avoided if people could understand Bayes' [32$_E$, 37$_E$]). When we have multiple related hypotheses and different types of evidence, such as in Fig.1c, the necessary Bayesian calculations are impossible for even mathematicians to do without special tools. Specifically, we need to use the Bayesian network (BN) representation where nodes represent the variables and arcs represent dependencies between them. Each node has a probability table (NPT), which provides the prior probability of each state of the node conditional (when the node has parents) on each combination of parent states. Efficient propagation algorithms for performing Bayesian inference for a broad class of BNs were produced in the late 1980s [16][59$_E$]. This led to widely available tools [72$_E$] that implement these algorithms and provide a relatively simple interface for users. These developments were the catalyst for an explosion of interest in BNs [1][5] [39$_E$]. However, despite significant benefits [5], and many successful applications [17], BNs have been under-exploited. This is especially true in areas where they offer the most potential for transformative improvements in quantified decision analysis (notably law, medicine and safety). There are three main reasons for this:

Problem 1: Good BN models are extremely difficult to build.

Constructing a BN involves two steps – determining the structure (what are the variables and dependencies between them) and determining the prior and conditional probabilities (i.e the NPTs). Ideally the domain experts (not probability experts) need to drive this process. However, despite developments in the last 15 years (see below), both of these steps remain a highly specialised task that is the preserve of a niche group of academics and consultancies with very strong mathematical qualifications.

Problem 2: There is strong resistance to BN models that involve subjective priors.
As for all Bayesian methods there is widespread resistance to relying on expert judgement and subjective probabilities [4], even though sceptics are often happy to rely on decisions made purely on subjective (but hidden) assumptions.

Problem 3: BN models are difficult to use and understand
We already highlighted above the problems with even simple Bayesian inference. For more complex BNs the challenge is especially daunting; there is no agreement on how end-users of a BN model should be expected to interact with it or how to present results to such users.

While there has been much fine theoretical BN research, especially in developing ever-improved propagation algorithms [1], relatively little research has directly addressed the three core problems listed above. One of the most popular ways to address the core of problems 1 and 2 is to rely exclusively on data so the role of the human expert can be minimised or even made redundant. The idea is that, if BNs can be 'learnt' automatically from data then this not only solves problem 1) but it also solves problem 2) since little or no expert judgement is used. The research in such learning algorithms is extensive [11][13][15_E, 17_E, 53_E, 96_E]. But learning BNs purely from data is somewhat ironic, since the rationale of Bayes is to improve predictions and decision-making by incorporating any type of 'data' and this includes expert judgement. BNs are no better able to discover true causal knowledge and insights from data alone than any other machine learning methods [48_E]. A BN learnt purely from data may suffer the same fundamental problems as described in Section a.2. Indeed, the causally 'irrational' model in Fig 1(b) ***was*** the BN model learnt purely from data.

There ***has*** also been extensive research on BN model building that has sought to *complement* available data with expert input [$2_{E,}$,12_E, 13_E, 22_E 40_E, 41_E 45_E, 46_E 50_E, 60_E, 84_E 82_E, 102_E 108_E], and even methods to learn certain types of missing variables [25_E, 7_E, 66_E]. These methods normally require a mature statistician to implement, for example using Winbugs [11_E] or the Application Programmer Interface (API) of a BN tool, and they also impose severe and often unrealistic constraints on model selection – for example conjugate, prior distributions are often enforced not because they are realistic but because they are mathematically convenient. Of most direct concern to us, however, is that while all this extremely rich research provides a range of powerful solutions to support BN modelling, ***none of it provides any help at all when the relevant data are not available***. The research that ***has*** explicitly focused on helping experts construct BNs without access to datasets, has been fragmented, unsystematic and at the margins of the field (indeed the absence of data is a typical reason for such work to go unpublished and unfunded). What we have is a set of partial solutions making tiny and slow inroads into an urgent problem.

To address the problem of defining the BN structure (from a 'blank canvas') the most promising work uses the idea of BN 'idioms' that capture some commonly occurring reasoning patterns [21_E, 34_E, 55_E $56_{E,}$ 76_E]. However, the generic idioms are incomplete and too abstract for direct use, with few domain specific instantiations. For managing and organising complex BN structures, the work on probabilistic relational models (PRMs) [110_E], Object Oriented BNs (OOBNs) [51_E, 76_E] and Dynamic Bayesian Networks (DBNs) [72_E], is helpful, but lacks a consistent framework for abstraction and efficient implementation.

Support for defining the probabilities in a BN is extremely limited. While the work in [14] is helpful for eliciting the unconditional prior probabilities from experts it is of minimal value for the much harder problem of defining the conditioned priors for nodes with multiple parent state combinations – a problem that is present in any non-trivial BN. In such situations it is commonly assumed that this cannot be done manually, and this is considered to be the most serious barrier to expert-built BNs. Solutions to this problem apply only for a limited class of nodes and relationships (namely certain Boolean nodes where logical operators such as noisy-OR can be effective [47_E]; ranked nodes using the approach defined in [31_E]; and numeric nodes using dynamic discretisation [54_E] [78_E] [77_E]). The perennial difficulty of defining NPTs also has a major impact on problem 2 (resistance to subjective prior probabilities) since decision makers are less likely to trust models where the NPTs were produced by unreliable ad-hoc means. One potentially promising way to address both problems is to exploit opportunities for minimizing the amount of probability elicitation required (and hence reduce dependence on subjective priors) through a) sensitivity analysis and/or b) identifying and eliminating redundancy. With sensitivity analysis [62_E] it may be possible to identify those nodes where fairly wide differences in the prior probabilities have little impact on model predictions. Unfortunately, despite promising work such as [8_E, 62_E, 104_E] there is currently no practical formal BN modelling mechanism or inference algorithm other than manual sensitivity analysis to compute the range of values pointing to any single conclusion within a specified threshold. The redundancy angle offers great potential for simplifying both structure and probability elicitation, because redundancy occurs in every model where there is asymmetry and where events have mutually exclusive causes [35_E]. These phenomena are especially common when modelling medical, legal and safety problems, but current BN theory and tools are neither able to identify such redundancy, nor provide alternative structuring. An extended BN theory and graphical notation is required (possibly based on the notion of chain graphs [98_E]) to provide a way forward for this important challenge.

Research addressing Problem 3 (using and understanding BN models) has been especially limited. In general, there is no easy way to make clear how the results of a BN model are arrived at. The methods that can work for simple Bayesian arguments [37_E, 44_E], do not scale up to non-trivial BNs [38_E]. Theoretical and algorithmic limitations also make it extremely difficult to enter anything other than very simple types of observations in a BN. For example, it is generally impossible to enter ranges (as opposed to point values) as observations, while existing tools do not correctly handle input of uncertain evidence [35_E]. Also some observations should ideally result in a change to the underlying BN model – but this would require a notion of dynamically restructured BNs, which is well beyond the capability of current BN technology. Fixing these problems will vastly improve the scope for more useful end-user interaction.

a.5: Research objectives

Our specific research objectives are to address the technical and cultural barriers that underlie the three problems listed at the end of Section a.3; in other words our objectives are to:

1) ensure that good BN models can be built (without enormous difficulty) in the absence of data;

2) counter resistance to BN models that involve subjective priors;

3) make BN models easy to use and understand.

Meeting these objectives will also require us to develop an extended BN theory with appropriate notation and algorithms implemented in an easy to use toolset, with extensive domain-specific support. The state-of-the-art shows that previous attempted solutions have been highly constrained, since they offer little in situations where data are absent.

a.6 Impact of proposed research

The proposed research has the potential to both reduce at source much unnecessary data collection *and* improve the results of analysis of data that is collected. Indeed, it has the potential to provide rigorous, rational, auditable, visible and quantified probabilistic arguments to support decision-making and recommendations in areas where currently only 'gut-feel' is possible. This could lead to: more rational and defensible strategic policy making by decision makers in government, financial, and other organisations; better medical diagnostics; better understanding of the impact of different types of legal and forensic evidence. Direct commercial beneficiaries will be any BN tool providers who wish to exploit the open source solutions packaged in the prototype tools the project will deliver. Academic beneficiaries include *non-statisticians* who undertake empirical and data-driven studies that can be improved and simplified by the method in the project, and *statisticians and mathematicians* who already use Bayesian methods. More generally this project will enable all scientists, statisticians, medics and lawyers, to be better able to reason about probability and understand the role and limitations of data, making better decisions with less data.

a.6 Methodology

The method is a balance between a) necessary theoretical work on extending BN theory, notation and algorithms; and b) empirical case studies and formal experiments in which end users drawn from law, forensics, medicine, and transport are directly involved from early on to ensure the theoretical work is relevant, usable, and validated. The theoretical issues we need to address fall into two categories – ***structure*** (WP1) and ***probability*** (WP2) which each span the 4 years of the project. WP1 has two major tasks: 1) develop a method for problem framing as causal models using idioms and 2) develop an extended unified graphical notation of BNs to support full abstraction, extending the work on GPMs, OOBNs and DBNs (including handling multiple levels of nested components). WP2 groups together the technical tasks that focus on minimising the effort required for defining NPTs without data (resulting in NPTs that are more easily justified and auditable, and with minimal redundancy), and also for entering observations in a BN.

WP1 and WP2 will produce extensions to the way BNs are specified and executed, as well as extensions to the standard BN inference algorithms, and new interface requirements for building and using BNs. Hence, we will have a workpackage WP3 that will first provide a consistent notation suitable for any graphical implementation, and second develop working prototype implementations of these new algorithms and interfaces. We will implement the ideas generated from WP1 and WP2 as soon as we have stable specifications and algorithms (delivering prototype tool increments each 6 months from year 2) so that the case studies can access working toolsets early. We will leverage as much as possible from existing free and commercial BN toolsets. To ensure the maximum possible exposure of the toolset, as well as a level playing field for future commercial BN tool exploitation, we will adopt an open source Java approach.

The empirical work is undertaken in WP4 and has two major objectives. One (using model-building case studies) is to validate the methods that result from WP1 and WP2 and the associated toolset from WP3; the other (using formal experiments) is to determine how best to understand and gain trust in the assumptions and results of BN models. Because there is no consensus on what 'validation' means for an expert built BN model, WP4 also contains a crucial task to define a validation method that will be applied to the case studies. The case studies involve leading scientists, practitioners and influencers in the disciplines of medicine, forensics, law and transport as well as participants in the BayesLaw consortium [2]. Initial case studies include: decision support for limb amputation; management of muscloskeletal injuries; age estimation from dental records; probabilistic evidence in medical negligence cases; analysis in the forensics of mycology; improved station safety through incident analysis; evaluating legal evidence pre-trial. The experiments will be developed by cognitive psychologists using the crowdsourcing tool Amazon Experimental Turk, to identify and refine which methods are most effective for communicating the results of BN computations. In addition to testing existing alternative presentation methods, they will help develop where necessary new presentation techniques (relevant to the assumptions and results of larger BNs). Two final work packages WP5 and WP6 deal with outreach and project management respectively.

The team: In addition to the Senior Fellow **Fenton** (50%), **Prof Martin Neil** (10%) will provide support for the underlying algorithmic work, **Dr William Marsh** (10%) will support the medical and transport case studies, **Dr Anne Hsu** (10%) (cognitive psychologist) will direct the experimental work, and **Dr Amber Marks** (5%) (barrister and lecturer in criminal law and evidence) will support the legal and forensic case study work. We request funding for three full-time Research Assistants (two to work on WP1-3 for all 4 years and one to work on WP4 in years 2-4)) and a programmer (25%) in years 2-4.

References (see www.eecs.qmul.ac.uk/~norman/projects/B_Knowledge.html for full list)

[1] BAYES-KNOWLEDGE : Effective Bayesian Modelling with Knowledge before Data, www.eecs.qmul.ac.uk/~norman/projects/B_Knowledge.html
[2] Bayes and the Law : Improving Expert Evidence Presentation and Analysis through Cross Disciplinary Collaboration, https://sites.google.com/site/bayeslegal/
[3] Darwiche, A. (2009). Modeling and Reasoning with Bayesian Networks Cambridge University Press.
[4] Druzdzel, M. K. and L. C. van der Gaag (2000). "Building Probabilistic Networks: Where Do the Numbers Come From?" IEEE Transactions on Knowledge and Data Engineering 12(4): 481-486.
[5] Fenton, N.E. and M. Neil, Risk Assessment and Decision Analysis with Bayesian Networks. 2012, London Chapman and Hall.
[6] Fox, J. (2011). Arguing about the evidence: a logical approach. Evidence, Inference and Enquiry. W. Twining, P. Dawid and M. Vasilaki, The British Academy.
[7] Gelman, A. (2008). "Objections to Bayesian statistics." Bayesian Analysis 3(3): 445-467.
[8] Gigerenzer, G. (2002). Reckoning with Risk: Learning to Live with Uncertainty. Penguin Books.
[9] Jensen, F. V. and T. Nielsen (2007). Bayesian Networks and Decision Graphs, Springer-Verlag
[10] Kahneman, D. (2012). Thinking, Fast and Slow. London, Penguin Books.
[11] Korb, K. B. and A. E. Nicholson (2004). Bayesian Artificial Intelligence, CRC Press.
[12] McGrayne, S. B. (2011). The Theory That Would Not Die, Yale University Press
[13] Neapolitan, R. E. (2004). Learning Bayesian Networks, Prentice Hall.
[14] O'Hagan, A., C. E. Buck, et al. (2006). Uncertain Judgements: Eliciting Experts' Probabilities. Chichester,
[15] Pearl J., Causality: Models, Reasoning and Inference, 2nd ed, CUP, 2009.
[16] Pearl, J., Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.
[17] Pourret, O., P. Naïm and B. Marcot, Eds. (2008). Bayesian Networks: A Practical Guide to Applications. Statistics in Practice. Chichester, John Wiley & Sons.
[18] Smith, J. (2010). Bayesian Decision Analysis: Principles and Practice. Cambridge University Press
[19] Spiegelhalter, D. J., K. R. Abrams and J. P. Myles (2004). Bayesian Approaches to Clinical Trials and Health-care Evaluation, John Wiley and Sons
[20] Tversky, A. and D. Kahneman (1974). "Judgement under Uncertainty: Heuristics and Biases." Science & Justice: 185: 1124-1131.

Section b: Norman Fenton Curriculum vitae (max 2 pages)

Details of qualifications, publications, projects, invited talks, professional memberships, external activities, prizes, blogs media appearances and consultancy can be found at www.eecs.qmul.ac.uk/~norman/

Professional Record

- Since March 2000 (70% employment): Professor of Computer Science and Director of the Risk and Information Management (RIM) Research Centre at ***Queen Mary University of London (Dept Electronic Eng and Comp Sci).*** My other current responsibilities include: School Management Team, Director of Examinations, running a double-unit module on software engineering, supervising 5 PhD students, supervising 13 project students per year (7 BSc and 6 MSc).
- Since Feb 1999: CEO of ***Agena Ltd*** (this is why my QM employment is part-time)
- Since March 2008: Affiliated Professor to the ***University of Haifa, Israel***.
- Feb 1989-March 2000: Professor at ***City University*** (Centre for Software Reliability)
- Sept 1984- Feb 1989: Reader and Director of CSSE at ***South Bank University***
- July 1988-Dec 1988: Seconded to Visiting Researcher at ***GMD, Bonn, Germany***
- Sept 1982-Sept 1984: PostDoctoral Research Fellow at ***Oxford University*** (Maths Institute)
- Sept 1981–Aug 1982: PostDoctoral Research Fellow at ***University College Dublin*** (Maths Dept)
- Since 2007: Acted as expert witness (in Bayesian analysis/decision-making and software assurance) in several major criminal and civil cases (see 10-year record)

Research Record Summary

- Published six books and 135 refereed research articles.
- PI on research projects worth £6.5 million since 1986
- Successfully supervised 14 PhD students, several of whom have continued on to academic posts.
- Pioneered novel approaches to measurement and risk assessment in a wide range of application domains for 28 years.
- My research in the 1980s and 1990s on software measurement transformed the use of metrics in software engineering by finally providing it with a rigorous framework grounded in measurement theory. My book "Software Metrics" (for which a third edition will be published in March 2013) has sold over 30,000 copies worldwide and has 4000 citations.
- Twice named as one of the world's top 15 software 'scholars' (by J Systems & Software) in the 1990s based on publications in the leading software journals.
- While at the Centre for Software Reliability, City University during the 1990s my research focused on the problems of assessing reliability in critical systems. This led to an embrace of Bayesian networks (BNs) that enabled us to develop causal models with expert knowledge supplementing data.
- In the late 1990's I pioneered a causal BN modelling approach to the problem of software defect prediction, and this original work[2] is among the most highly cited papers in computer science (668 citations). It formed the foundation for subsequent advanced modelling of software risk as described in subsequent papers (see 10-year record). The underlying philosophy has also led to a methodology for measuring and predicting resilience and vulnerability in complex systems such as in defence and banking by accounting for hard and soft factors.
- In 1999, I created a spin-off company - Agena - to exploit the many ways in which BNs could be applied; shortly afterward I left my full-time post at City to join Queen Mary on a 70% employment so that I could also spend time growing Agena. During 2001 and 2003 Agena received individual and institutional investments enabling it to employ a dedicated team to develop a software platform (AgenaRisk) for building and deploying BN-based solutions. During this development phase (2003-2008) Agena peaked at 11 employees. Agenarisk currently has 8000 users worldwide.
- When I joined Queen Mary in 2000 I formed the Risk Assessment and Decision and Analysis Research group (which later became RIM) with co-investigators Prof. Martin Neil and Dr William Marsh. The group has focused its research on addressing the problems of how to scale-up BN methods and theory to deal with large-scale real-world problems. We have developed breakthrough BN algorithms that, for example, for the first time enable modellers to use unconstrained continuous variables alongside discrete

[2] Fenton NE and Neil M, "A Critique of Software Defect Prediction Models", 25(5) IEEE Transactions on Software Engineering, 675-689,1999.

ones, without going through the difficult (and intrinsically inaccurate) process of static discretisation. Further details of my major research contributions in the period are in the 10-year record.

- Some of these research advances have been implemented in AgenaRisk.
- Between 2000-2009 much of my applications-driven research was focused on critical systems assessment (with companies normally not allowing results to be published). Examples include:
    - ***Defence/Security***: Worked with: *QinetiQ* on vehicle reliability and whole-life costs (our BN-based system "TRACS" has been routinely used to evaluate early prototype military vehicles to avoid costly testing); *DSTL* on dependability evaluation of land systems, sensor-based security architectures, and decision support for sensor fusion and threat assessment; Numerous US defence companies including *Raytheon*, *GE-Aviation*, *Boeing*, *General Dynamics* and *Lockheed Martin* use models and software originating from the research. Supply chain risk assessment has been an active area directed by recent initiatives from the *USAF*.
    - ***Finance:*** worked on operational risk assessment with banks in the UK, South Africa and Norway. *Royal Bank of Canada* use a model I developed for routine identification of risk in new and ongoing maintenance projects; worked with: *Milliman* (the biggest independent international actuarial firm) who adopted our approach for operational risk assessment; *VocaLink* on infrastructure vulnerability evaluation; *AON-Benfield*: comparing accuracy of competing models for predicting insurance losses arising from catastrophic flood events.
    - ***Transport***: worked with: *NATS* on air traffic management architecture risk modelling; various parts of the UK rail industry on safety risk modelling and *Bosch* to predict the impact of software defects in automobile control systems. In 2005-2006 I worked 11 months on a multi-million pound civil case (Westinghouse Rail Systems Ltd v Data Systems and Solutions Ltd, Reference HT 05 181, Technology and Construction Court) on software quality and risk assessment of safety critical software in the rail industry. My 220 page expert report played a prominent role in settling the case.
    - ***Telecoms/electronics***: Worked with: most of the key players (*Philips*, *Siemens*, *Ericsson*, *Orange*, *Motorola*, *Tellabs*) on models to support risk assessment of their software intensive systems, especially with respect to improved accuracy of defect prediction; with *Motorola* to more accurately monitor and control faults in hardware components (this work save Motorola $5m per project in 2007); Developed International Patent (Number WO 03/090466) for Improved TV Programme Selection (based on Bayesian Networks, Fuzzy Logic and an original approach to programme classification) and worked with NDS and Sky to implement a prototype on the Sky Box platform.
- Since 2008 my research has focused on improved decision-support in law and medicine. This work has created enormous media interest (I have been interviewed on TV and radio, and there have been articles describing the work in the Guardian, the New Scientist and Nature) and led to a demand for my services as expert witness in several major criminal and civil cases (see 10-year record). Since June 2011 my approach has been the focus of an international consortium (*Bayeslaw*)[3] of statisticians, lawyers and forensic scientists working to improve the use of statistics in court. In medicine I have been working with numerous prestigious specialist medical units to improve clinical performance (including the world renowned Trauma team at the Royal London, to develop improved decision support for the management of traumatically injured limbs). I also led an EPSRC digital economy consortium on Data Information and Analysis for Clinical Decision Making.
- I have been prominent in raising public awareness of probability and statistics, as demonstrated by my blogs and websites on these issues (I have also contributed to a number of TV documentaries); my interactive tutorial on BNs is consistently ranked number one in Google searches and receives thousands of hits per week. My new book (see 10 year record) will be the first to bring BNs to lay readers. I have published (in academic journals) several articles describing BN models for football prediction; the most recent is consistently 'beating the bookies' and the predictions are actually posted weekly on a popular website maintained by my PhD student (www.pi-football.com).
- To capitalise properly on the agenda described in this proposal I need to be freed from a substantial proportion of my current QM duties for a significant period of time, especially as much of the proposed work represents a change from my established work in software engineering and risk.

Current Funding ID

I have only one ongoing grant as PI "Improving Legal reasoning with Bayes" 01/10/2011 - 31/03/2015 which funds one PhD student's fees and scholarship. I am working on a proposal to the UK ESRC on causal modelling to improve energy monitoring devices, for which I will be requesting 5% of my time.

[3] https://sites.google.com/site/bayeslegal/

Section c: Norman Fenton 10-year track-record (max 2 pages)

Google Scholar states that I have 9737 citations from 100 cited Publications with an *h-Index* of 38.
All citations counts, including those below are taken from Google Scholar citations on 22 Oct 2012

1. Fenton NE, Marsh W, Neil M, Cates P, Forey S, Tailor T, "Making Resource Decisions for Software Projects", 26th International Conference on Software Engineering (ICSE 2004), May 2004, Edinburgh, United Kingdom. IEEE Computer Society 2004, ISBN 0-7695-2163-0, pp. 397-406 *Citation count 101*

   This paper, in the world's leading software engineering conference (8% paper acceptance rate) describes a BN model to help project managers trade-off resources used against outputs (delivered functionality, quality achieved) in complex software projects. The model was built in collaboration with senior project managers at Israel Aircraft Industries, Philips, and QinetiQ. The model has since undergone extensive validation and refinement by organisations outside the original consortium.
2. Fenton NE, Neil M, Hearty P, Marsh W, Marquez D, Krause P, Mishra R, "Predicting Software Defects in Varying Development Lifecycles using Bayesian Nets", Information & Software Technology, Vol 49, pp 32-43, Jan 2007 *Citation count 102*

   This paper, in a special issue celebrating the most cited papers in Software Engineering since 2000, describes work that followed the earlier Fenton and Neil most cited paper. The novelty is the use of a tailorable model that can be adapted to arbitrary software life-cycles. On 36 projects at Philips (Bangalore and Eindhoven) the model achieved predictive accuracy of 93%, a significant improvement on what can be achieved by other approaches with similar data. The modelling approach was incorporated into the AgenaRisk software and has been used by hundreds of researchers and companies, including some (e.g. Siemens) who published further validation results.
3. Fenton NE, Neil M, and Caballero JG, "Using Ranked nodes to model qualitative judgements in Bayesian Networks" IEEE TKDE 19(10), 1420-1432, Oct 2007 *Citation count 57*

   Provides a simple, but revolutionary method that enables modellers to capture common subjective relationships without having to specify complex probability tables. The approach, which is applicable to a large class of commonly occurring structures, was widely validated on real industrial problems. It was subsequently implemented in Version 5.0 of AgenaRisk (the 'ranked node' functionality). Hence, this work represented a major breakthrough in BN research and technology.
4. Fenton, N.E., Neil, M., Marsh, W., Hearty, P., Radlinski, L., and Krause, P., On the effectiveness of early life cycle defect prediction with Bayesian Nets. Empirical Software Engineering, 2008. 13: p. 499-537. 10.1007/s10664-008-9072-x *Citation count 38*

   Provides a comprehensive validation (in consumer electronics), with sensitivity analysis, of an approach that enables organisations to use expert judgement and limited data during early development and testing phases to predict subsequent operational defects. The data and models were also published in 2008 on the PROMISE website and have been used by numerous researchers since.
5. Neil, M., Tailor, M., Marquez, D., Fenton, N.E., and Hearty, P., Modelling dependable systems using hybrid Bayesian networks. Reliability Engineering and System Safety, 2008. 93(7): p. 933-939. http://dx.doi.org/10.1016/j.ress.2007.03.009 *Citation count 20*

   One of the biggest impediments to more widespread use of BNs was the lack of efficient algorithms and tools for handling BNs that contain continuous variables. This paper describes an algorithm that solves this problem, demonstrating that it can be used as an alternative to analytical or Monte Carlo methods with high precision. This work on dynamic discretisation was implemented in AgenaRisk 5.0, and is recognised as a major breakthrough for BN applications.
6. Fenton, N.E., Neil, M., and Marquez, D., Using Bayesian Networks to Predict Software Defects and Reliability. Proceedings of the Institution of Mechanical Engineers, Part O, Journal of Risk and Reliability, 2008. 222(O4): p. 701-712 **Citation count 20**

   This paper reviews our previous work on improved software defect prediction, but includes new theoretical and empirical work explaining how the crucial, and novel, role of dynamic discretisation leads to greater accuracy. The paper was 'highly commended by the Editor and Editorial Board of the Journal" and was nominated for the Professional Engineering Publishing Prize
7. Hearty, P., Fenton, N., Marquez, D., and Neil, M., Predicting Project Velocity in XP using a Learning Dynamic Bayesian Network Model. IEEE Trans Software Eng, 2009. 35(1): 124-137. **Citation count 25**

   Our previous work using BN models to improve software risk assessment required mature software engineering processes and data collection. This was fine for companies such as Philips and Motorola but not suitable for small companies or any extreme programming environment. The research described here represents a breakthrough in that it enables quantitative effort predictions and risk assessments without requiring any special metrics collection effort.

8. Fenton, N. and Neil, M. (2010). "Comparing risks of alternative medical diagnosis using Bayesian arguments." Journal of Biomedical Informatics, 43: 485-495, *Citation count 12*
   Explains an approach that enables complex Bayesian arguments to be presented in such a way that they can be understood by lay people. Was used here in a medical negligence case that was going nowhere because the lawyers and experts were unable to articulate quantitatively the different risks. Using our approach the claimants' lawyers and experts were finally able to make the case effectively. The approach formed the basis for subsequent novel research in BNs for legal reasoning used in several high profile criminal cases.
9. Fenton, N.E. and Neil, M. (2011), 'Avoiding Legal Fallacies in Practice Using Bayesian Networks', Australian Journal of Legal Philosophy 36, 114-151, 2011 ISSN 1440-4982 *Citation count 7*
   Describes our novel use of Bayesian reasoning applied to evidence in cases where we were involved as expert witnesses (notably Levi Bellfield's first murder trial). This work led to us being engaged to use the approach on several further high profile criminal cases (see Expert witness work below). Interest in the work resulted in the BayesLaw consortium described in the CV.
10. Fenton, N. E. (2011). "Science and law: Improve statistics in court." Nature 479: 36-37. *Citation count* 1
   This invited commentary article in the world's leading science publication summarised some of my recent work in using Bayes to improve legal arguments, while also describing the new challenges.

Other notable contributions in last 10 years:

- *Fenton, N.E. and M. Neil, Risk Assessment with Bayesian Networks*. Nov 2012, London: CRC Press, ISBN: 9781439809105, 28 October 2012 Foreword by Judea Pearl (2011 Turing Prize winner)
  This book was the culmination of several years of work. As confirmed by the statements of Judea Pearl and numerous other senior academics (see http://bayesianrisk.com/ and reviews on the publishers website[4]) it is the first book to present BNs to non-specialists. Without this book the proposed research in the fellowship application would be unachievable in the time-frame proposed.
- *Expert witness work*. I have been extensively involved in several high profile cases. Some were civil cases (e.g. Westinghouse Rail Systems Ltd v Data Systems and Solutions Ltd - described in the CV, and a medical negligence case against the NHS described in reference 8 above), but most have been criminal. These include:
  - R v Levi Bellfield (2007-2008) for the murder of 2 women and attempted murder of 3 others. My evidence identified probabilistic fallacies in the prosecution opening (the judge instructed prosecuting counsel not to repeat these statements in closing) and my 150-page report formed the core basis for the defence argument.
  - The retrial of R v Gary Dobson and David Norris (the Stephen Lawrence case) 2011-21012. I advised the defence counsel on probabilistic flaws in the presentation and analysis of the DNA evidence.
  - R v Kerbey (2012) Appeal. This was a youth convicted of involvement in the Croydon riots of the summer of 2012. The case was based on a flawed statistical argument that I exposed.
  - An ongoing case (R v LW) – and potentially the most significant ever for Bayes and DNA – where I was asked to review the extremely poorly presented DNA evidence that was the key part of the case against a man convicted of raping his half sister. Despite severe flaws in all of the other evidence, it was only my extensive Bayesian evaluation of the DNA evidence that finally led to the defendant being granted an appeal.
- As Director of Research for the Department of Computer Science between 2006-2009 my primary responsibility was to handle the crucial 2008 "REF" submission that determines the official UK national research ranking and the amount of government funding the Department receives over the subsequent 5 years. My submission resulted in Queen Mary being determined as the most improved CS department in the country (moving from 33rd out of 88 in the previous submission to 11th).
- 46 invited seminars or keynotes in last 10 years at international conferences, universities and major companies (full list on website). Also 10 conference programme committees since 2006.
- My online guide to good technical writing[5] (developed over a number years) has been adopted by dozens of companies and academic institutions worldwide.

[4] www.crcpress.com/product/isbn/9781439809105
[5] www.eecs.qmul.ac.uk/~norman/papers/good_writing/Technical%20writing.pdf