Patient Managed Decision-Support using Bayesian Networks (PAMBAYESIAN)

Track Record

The project is an interdisciplinary collaboration involving a diverse team of senior researchers in the School of Electronic Engineering and Computer Science (EECS) and clinical academics from the Barts and the London School of Medicine and Dentistry School (SMD). The team has evolved from previous relevant successful interdisciplinary projects in which members of EECS, and especially the Risk Information Management research group (RIM), have collaborated with multiple disciplines at SMD, the Royal London and St Bartholomew's hospitals, and other medical practitioners. These collaborative projects have resulted in the development of a series of intelligent causal models for improved medical decision-making. These models – based on Bayesian Networks (BNs) – include: strokes [10] and coagulopathy risk [23]; limb viability after traumatic event [24]; recidivism in violent psychiatric patients [6] and Warfarin therapy.

The PI, **Prof Norman Fenton** (EECS), is Professor and Director of the RIM group at Queen Mary, and a Director of Agena Ltd. The RIM Group (which includes co-PIs Prof Martin Neil and Dr William Marsh) has pioneered research and applications of causal BNs in a range of critical areas of intelligent risk assessment and decision analysis that involve the combination of expert judgment with data. The research includes both a) breakthrough algorithms that enable a much broader class of problems to be modelled accurately, e.g. [19][20], and b) a range of novel methods that enable non-specialist to develop and deploy realistic BN models, e.g. [9]. The RIM group is world-leading in its practical applications of Bayesian modelling and/or risk quantification which, in addition to medical decision support, includes such diverse areas as: operational risk in finance, security and defence, legal and forensic arguments, systems and design reliability, transport safety, software project risk, football prediction, simulation (using dynamic discretisation as an alternative to Monte Carlo), and cost benefit analysis.

Fenton has managed research projects to the value of over £8 million since 1986 as well as many major consulting projects through Agena. He has published 7 books and 240 refereed articles. Fenton and Neil's highly acclaimed book [11] was the first to bring BNs to a general audience. Fenton's google scholar entry (3 May 2016) shows 13,892 citations and h-index of 44. In April 2014 he was awarded one of the prestigious European Research Council Advanced Grants (value €1,572,562) to focus on the special problems of developing causal BNs in critical applications where there is minimal data. Fenton has been active in the media in raising public awareness of the importance of probability theory and Bayesian reasoning. In March 2015 he presented the award-winning BBC documentary Climate Change by Numbers. He has been an expert witness/consultant on probabilistic risk assessment in several high profile legal cases (including medical negligence) and in 2016 he will be leading a 6-month Programme on Probability and Statistics in Forensic Science at the Isaac Newton Institute for Mathematical Sciences, University of Cambridge. In addition to his research on risk assessment, Norman is renowned for his work in software engineering (including pioneering work on software metrics).

Prof Martin Neil (EECS), co-PI has led much of the BN algorithm development work of the RIM group, including the breakthrough dynamic discretisation algorithm [20]. He has been PI in research and consulting projects to the value of over £4 million and has published one book and 80 refereed articles. Before setting up Agena and joining academia Martin previously held senior positions with JP Morgan and Lloyds Register in the areas of software project governance and safety critical systems evaluation respectively.

Dr William Marsh (EECS, S&E), co-PI, is a Senior Lecturer and has played a major role in RIM's extensive previous work on decision support for medical applications, in trauma [24] and forensic psychiatry [6]. He has also pioneered the use of BNs in railway safety. Before joining QMUL he worked in software development and assessment, latterly at ERA Technology where he led software and safety assessment projects in avionics and railway engineering.

Prof Paul Curzon (EECS, S&E), co-PI is renowned for his work in both human computer interaction and public engagement. Curzon was co-Pi and management team member of the EPSRC programme grant CHI+MED on the safe medical device design [5]. This included work, for example, understanding the use in context of a new inpatient glucometer [12] and developing hazard analysis methods with user error central [17] as applied to infusion pumps. He cofounded

and runs the internationally reputed ‘Computer Science for Fun’ (cs4fn) public engagement project, currently sending >22,000 copies of its magazines on computer science research to UK school subscribers, and has personally delivered public engagement talks to over 10,000 students.

Dr Akram Alomainy (EECS), co-PI, is Senior Lecturer and Industry Strategy Coordinator for EECS. His expertise in the field of intelligent/cognitive wireless networks and wearable technologies from large to nano-scale has attracted substantial funding from national and international research councils. His current research portfolio is around £3.7m as PI and Co-I including EPSRC research grants and platform fund. He has authored and co-authored two books, five book chapters and more than 150 technical papers.

Dr Dylan Morrissey (SMD), co-PI, is NIHR/NEE Consultant Physiotherapist and Reader in Sport and Exercise Medicine and leads the musculoskeletal (MSK) case study. He currently holds a 5 year senior clinical lectureship, and has collaborated with Fenton and Marsh in developing early (MSK) causal models that will form the foundation for the MSK case study. He leads a PPI initiative and clinical trials in an MSK physiotherapy department with ~2500 referrals per month. Further, he co-leads a fully equipped human performance laboratory suitable for testing physical activity measurements with respect to motion capture and force measurement equipment. He has published over 70 papers and obtained more than £4m in research funding. Morrissey’s team also includes co-PIs Dr Frances Humby (SL and Honorary Consultant Rheumatologist) and Dr Victoria Tzortziou-Brown NIHR Clinical Lecturer, commissioning co-lead and local GP and – as a consultant – Orthopaedic Surgeon Mr Manoj Ramachandran, who is also Clinical Informatics Officer (CIO) and the Director of Informatics and Innovation for Surgery at Barts Health.

Dr Bobby Huda (SMD), co-PI leading the diabetes case study, is a Consultant Physician in Diabetes & Metabolism at Barts and the Royal London Hospital, and Honorary Senior Lecturer in the SMD. He leads a busy antenatal service at the Royal London which delivers 6000 women a year. His current research interests include type 1 diabetes and gestational diabetes. He has published over 30 per-reviewed articles and book chapters, and is currently PI on 3 clinical studies. Dr Huda will work as part of a team led by co-PI **Prof Graham Hitman** (Director of the Blizard Institute), and **Prof Shakila Thangaratinam** (Professor of Womens Medicine, QMUL and obstetrician to the diabetes pregnancy clinic). East London has a heavy burden of diabetes in pregnancy due to a young population with a large number of South Asian people and other ethnic groups; the Blizard Institute is leading centre for research in diabetes; e.g., Prof Hitman (H index 65 with >300 publications) has recently coordinated a €3m FP7 EU grant into the prevention of diabetes and obesity using pregnancy as a window of opportunity for both mother and child.

Dr David Collier (SMD), co-PI, is Joint Clinical Director of the William Harvey Research Institute Clinical Research Centre and leads the heart-monitoring case study and GP engagement. He is CI, PI and co-investigator on 30 current clinical trials, the MRC START methodology programme on recruitment methods for clinical trials, and the Barts NIHR cardiovascular Biomedical Research Unit. His team (with Mark Caulfield and Mel Lobo) hold many recruitment records for clinical trials in cardiovascular disease, and have influenced multiple clinical innovations and guidelines (ASCOT, ILLUMINATE, Pathway 1,2,3 MRC, ROX coupler). The HiLo project is putting ASCOT experience back into primary care and impacting QoF blood pressure and lipids in Tower Hamlets, City & Hackney and Enfield (9,200 recruited). Their Trials Connect project (with Paul Bowers-Isaacson FCIEA) is pioneering post-study patient training and linking 6th form students with trials patients and medical students. For the project David’s team will include – as a consultant – cardiologist **Prof Richard Schilling**, whose research interests focus on elucidating the mechanism and percutaneous treatments of complex cardiac arrhythmia in the human heart.

Prof Anita Patel (SMD), co-PI, is a Health Economist in the Blizard Institute. She has evaluated the cost-effectiveness of various health care treatments through applied multidisciplinary research. Her work has been influential in underpinning national and international clinical guidelines and service change. She has over 70 peer-reviewed publications and has been a named co-applicant/co-investigator on collaborative research contracts totalling £19m, including as lead economist for a £3.5m technology assessment centre (at King’s College London) funded by the National Institute for Health & Care Excellence to support its decision-making. She also undertook key economics-related roles in securing and/or delivering a further £17m of research.

Patient Managed Decision-Support using Bayesian Networks (PAMBAYESIAN)

1. Background and State-of-the-art

The ambition of this project is to create a new generation of easy-to-use medical decision support systems for direct patient use, with real-time monitoring for chronic conditions. This will increase patient independence and decrease reliance on direct consultation, allowing more autonomous care at home with a reduction in associated health care costs.

While substantial technology for remote monitoring exists (e.g [18]), it is not in widespread use or depends on expert data interpretation (e.g. [26]). A detailed evaluation of a telemonitoring system for COPD in Scotland (TELESCOT) highlighted two key enablers needed for wider adoption [21]: ***intelligent processing*** [3] of the data presented both to patients and clinicians, to avoid e.g. false alarms or excessive work; and ***user-centred interaction design*** [16], without which the technical systems do not lead to the expected benefits. Data processing needs to better simulate an expert's interpretation and understanding of what is really happening to a patient and then to support a patient's decision about how best to interact with the healthcare system. A novel aspect of our approach is to overcome these barriers in an inter-disciplinary and integrated way, providing decision-support for all stakeholders and considering, for example, predictable variations in the way patients interact with technology. Our vision for intelligent decision-support systems incorporates sensor data integrated alongside more standard clinical information sources, processed using ***causal models***, based on expert-built ***Bayesian Networks*** **(BNs)**[11].

BNs are graphical models representing causal or influential relationships between a set of variables, providing probabilities for each unknown variable. As described in the Track Record, the team has developed a series of intelligent causal models for improved medical decision-making based on BNs. The underlying models combine four different types of information:

1. Historical data relevant to the medical condition(s), its treatment and outcome
2. Published evidence and expert medical judgment, including causal relationships (since the relevant data is often limited and key factors may be unobservable)
3. Background patient-specific data
4. 'Live' data from patient examination and clinical tests

The first two data types form the basis of the pre-defined model (termed 'expert driven' because of data type 2), while the last two are the inputs for a specific patient. The models are probabilistic: when the patient-specific inputs are entered, Bayesian inference computes updated probabilities for variables of interest: e.g. diagnoses that the patient has a particular condition, may suffer a particular medical event in future, or respond to a treatment. The resulting 'decision support systems' have proven effective when compared to previous state-of-the-art methods, and also complement other medical applications of Bayesian reasoning which have focused more on parameter and machine learning rather than exploiting any rich causal structure provided by domain experts [2][13][15]. It is increasingly well understood how to build expert-driven BN decision-support models, and progress is being made with their introduction into routine clinical use.

In parallel with the BN developments it has become feasible to monitor patients continuously with wearable sensors, covering an increasing range of physiological data. Current wearable sensors are able to collect data that can be used to monitor movements of joints in the body, breathing, neurological, blood glucose and cardiovascular bio-signals. Devices are increasingly small and energy efficient, so better suited for continuous sensing. However, the use of real-time data in BNs for clinical decision making is less straightforward as it requires a time transformation from frequent sensor data to the relatively infrequent analysis of significant trends or events, with the introduction of summary variables (e.g. number of events). The use of Markov models, including BNs, for filtering continuous data streams is well known [4] and we can build on this to achieve a more integrated design process: a single model of causal mechanisms that can be abstracted into a coherent set of BNs at different levels. Most existing BN modelling approaches lack abstraction mechanisms; some structured approaches (such as Object Oriented BNs) exist, but do not provide the way to build a BN with several levels of abstraction. As a result, current techniques for designing a BN decision model are inadequate for our goals.

2. National importance

The proposed research maps directly to the EPSRC's Digital Health grand challenge 'Transforming Community Health and Care', to 'use real-time information to support self-management of health and to facilitate timely interventions'. Our contribution will be more 'intelligent' decision-support systems to work with the latest portable sensors so that patients can confidently manage chronic conditions with efficient support from healthcare professionals. The potential for technology to both improve health care and reduce its cost cannot be achieved by technologists working alone. Hence, the project is highly inter-disciplinary, with close co-operation with clinical researchers around case studies in ***three chronic medical conditions***, which are themselves of national importance and in which SMD has extensive ongoing research that can be exploited in this project. These conditions, which provide an ideal testbed, are:

Musculoskeletal problems (MSK): These present an increasing burden of care across all health care sectors, accounting for 30% of general practitioner time, with the ageing population bringing more long-term problems. We will focus on patients with inflammatory joint disease who require on-going monitoring and consultation to optimise care of fluctuating disease. New drugs have transformed the management of these patients but at huge cost. We will engage patients, commissioners, the primary care clinical effectiveness group, rheumatologists and GPs in creating and testing technological innovations for this environment.

Diabetes in Pregnancy: With diabetes reaching epidemic proportions, the evidence suggests that early treatment to maintain optimum glucose levels will reduce the future burden of complications. This is especially so for diabetes in pregnancy, which can affect both mother and child. This case study will examine how to help pregnant women with diabetes, in partnership with their health practitioners, to use clinical decision-systems in conjunction with self-testing of blood glucose to manage both lifestyle and appropriate pharmacotherapy.

Atrial Fibrillation (AF): Patients who have irregular heartbeat (AF) are at increased risk of stroke due to blood clots forming in the heart. Ablation of parts of the electrical circuit in the atria at the top of the heart can restore normal rhythm, but recurrence of AF means that blood thinning has to be continued. 24h ECG monitoring is now practical: the latest sub-matchbox sized devices provide comprehensive autonomic nervous system data, over a Bluetooth link. However, using this mass of data to predict who is at risk of AF requires the kind of complex reasoning made possible by the proposed work in the project.

The project will build on EPSRC funded research: CHI+MED (EP/G059063/1) on device safety (Curzon co-PI) and MATCH (EP/F063822/1) on health technology assessment and engage with the GetAMoveOn network (EP/N027299/1) and the Strathclyde CDT on medical devices. Research on Bayesian tracking (EP/K020153/1) complements our algorithms work. Exploiting new digital technology in healthcare is of national importance; for this the project is supported by UK medical technology companies (BeMoreDigital, Mediwise, Rescon, SMART Medical, uMotif) as well as IBM UK and Hasiba Medical, who will form the core of the project **Advisory Board**.

3. Academic impact

The project promises to transform the application of BNs, with methods and algorithms to allow more realistic models than are currently practical by overcoming known problems of poor methods of abstraction and difficulties combining judgment and data. Although many practitioners already use BNs for decision-support and probabilistic analysis, more widespread use is prohibited by limits on the size and complexity of models. The project will deliver the capability to build and execute far more complex models efficiently, and a framework for integrating models at different levels of abstraction. The human-computer interaction studies will provide academic impact for HCI researchers given the novel aspects of integrating state of the art decision-making and wearable sensor technology into the everyday lives of patients with chronic diseases. The project will deliver open source code so that academic researchers and users can use the new techniques without restriction.

The project's close collaboration with clinical researchers will also be transformative. Instead of a protracted process of technology developed independently of its users and then shown in trials to have a disappointing impact, e.g. [21], we will involve users from the start. We will set-up clinical focus groups around the case studies to ensure that technology solutions fit the actual problems.

Through this early engagement, we will develop new patterns of inter-disciplinary working relevant to other aspects of innovation in digital healthcare.

4. Research hypothesis and objectives

The research hypothesis is that, by incorporating intelligence, it is possible to develop patient self-monitoring decision support systems that are safe and easy-to-use, but most importantly allow much more effective interactions between patients and supervising clinicians. The project goal can be achieved by developing a new framework for distributed probabilistic decision-support systems. The framework will combine continuous patient data from local sensors with 'conventional' data (e.g. blood tests, imaging results) and allow both autonomous and collaborative decision-making between patients and clinicians. This requires us to overcome the complexity created by multiple system layers, each needing some intelligence:

Sensor layer: where intelligence is needed to extract clinically significant observations from the sensor data stream, specific to the patient's circumstances.

Monitoring and management layer: provides a patient interface, used to give advice and gather patient reports about their condition: intelligence is needed to tailor the information requests (as a clinician would) and to generate patient advice.

Diagnosis and intervention layer: provides an interface for hospital data and decisions of a supervising clinical team and avoiding unnecessary visits to a clinic or hospital. This layer needs intelligence to generate alerts and support decisions such as changing the parameters within which the other layers have autonomy.

The need to include intelligence, derived from expert knowledge of causal relationships, in the different layers so that the system works in an integrated way, gives rise to three major complexity challenges, which collectively define the research objectives:

AI complexity challenge. We will provide a new framework to elicit causal knowledge of the domain from experts and implement it in a group of co-operating probabilistic models. Time plays only a minor role in existing probabilistic decision support systems, which are called upon by the clinician on demand. In the proposed system, each layer runs at different rates. For example, the sensor layer is continuous but needs to behave like an experienced practitioner and detect significant changes. We expect to use sequential Bayesian updating to track underlying (and not directly observed) clinical states advised by the experts. Further, throughout the system, effective decision support depends on predictions conditioned on the possible interventions: while the theory for this is established more automation will be needed.

Interaction and safety complexity challenge: New technology changes the experience of healthcare for both providers and patients. We will develop an understanding of the contexts of both as the basis for appropriate interaction design. Patients and carers need clear, easy to understand and supportive guidance that fits their circumstances and their lives. The clinical team need new methods to 'interpret' and simplify model outputs, with easy to understand 'alerts', 'actions' etc., relevant to the condition. Interaction design must be safe, easy to use and fit healthcare workflows. Experienced clinicians do not interact with all patients in the same way but adjust the information exchanged to meet a specific patient's needs, so the system must too. We seek to understand and model the impediments to using self-monitoring devices across different application areas and contexts, and the reaction of different patients to varying amounts of information, then prototype and evaluate designs.

Systems integration complexity challenge. To exploit sensor technology in full, we will develop new techniques for efficient BN inference in small autonomous devices capable of real-time execution. The sensor level devices may be low-power micro-controllers with limited memory, required to do real-time model inputs and updates. The challenge is not so much a lack of available algorithmic approaches, but the lack of integration with the way BN models are developed using general-purpose packages. An integrated development process is needed that includes algorithms to perform the inference needed at the sensor layer. We also have to consider the reliability of devices and the safety of the system. Failures, arising either in individual devices or from the connections between them, may give rise to safety hazards with important conditions missed or incorrect advice given. Local intelligence provides the

opportunity to ensure safety by monitoring for potentially hazardous conditions.

While the research will be heavily informed by the selected medical case studies, the objectives are generic. To ensure both a clear route to impact, we will ensure that the innovations are not tied to specific chronic medical conditions or to current sensor technologies.

5. Programme and methodology

The plan is organised around six work packages (see also Work Plan attachment, which describes more detailed subtasks) with the three medical applications interwoven into each package. The MSK and diabetes applications (referred to as studies 1 and 2 respectively) will be used as initial test beds, developed concurrently with the methods in WP 1-4; after refinement of the methods the AF application (study 3) will subsequently be used to validate the framework's general applicability to decision support for chronic medical conditions.

WP 1 (The AI complexity challenge): This is composed from two tasks:

Knowledge elicitation framework: Starting from the concepts outlined above, including clinically significant observations, interventions and clinician alerts, we will develop a framework for modelling the medical knowledge needed to manage chronic conditions.

Interacting BNs: Rather than a single integrated BN, we will develop novel methods of connecting the BNs necessary for the different layers of the system into an integrated decision-support system. New algorithms will be developed and implemented for handling the analysis of interventions and also the use of counterfactuals [5].

WP 2 (The interaction and safety complexity challenge): This is broken into three tasks with patients and other stakeholders involved in up-front participatory design:

Understanding the Context of Use: Using in-context structured interviews, we will build an understanding of clinician and patient needs and investigate the extent of adaptation required in an effective interaction design. We will develop personas [7] for different patient groups and stakeholders, with linked contextually rich scenarios of use, as both general design tools and to directly drive our prototype design. The well-known difficulties of presenting probabilistic predictions are magnified for a diverse group and different approaches may be needed for different patients. We will investigate how to focus input requests so that patients understand the value to them. We will draw general lessons, relevant to applications beyond the case studies.

Use-Related Safety: We will apply techniques in safety engineering of identifying hazards and analysing their causes. It will draw on our work on CHI+MED, in collaboration with the US regulator, developing use-focussed hazard analysis techniques [8][17].

Interaction Design Prototyping and Validation: Interaction design prototypes will be developed with and for the different stakeholder groups, and also validated in interaction design experiments and role-play of scenarios with patients and professionals, including both those who participated in the activities and those who did not. We will seek to show that reliable evidence of usability can be generated before clinical trials are undertaken.

WP 3 (The systems integration complexity challenge): This addresses the challenge of a system composed of multiple devices, including some embedded devices (e.g. ARM Cortex-M) in the sensor layer. The tasks are:

BN algorithm development: Current BN inference algorithms [19][20] can handle large complex dynamic BN models that are hybrid in the sense of incorporating both discrete variables and arbitrary types of continuous variables needed in realistic medical models. A new generation of efficient approximation algorithms will be developed to allow relevant parts of such models to work given the limited computational resources of mobile or wearable devices.

Built-in Safety Monitoring: Since we have intelligence in multiple sub-systems, there is potential for each sub-system to monitor the others. The innovation is to build BNs that can detect the occurrence of hazardous failure conditions (identified in WP2) while running alongside the decision-support system itself, with redundancy between the different parts of the system.

Data and Sensors: For each application we will obtain data (subject to ethics and information governance) from multiple sources, including: sensors, patient reported outcome measures

(particularly for pain) and medical records, including trials. Details of the sensor technology to be used are described in Justification of Resources, but given very rapid developments in this area we will be flexible. We do not expect to be able to make this data fully coherent (e.g. by gathering sensor data from patients enrolled in trials, due to regulatory complexity) but will ensure enough coherence to show the potential performance, sufficient to support an application for a pilot trial, following the completion of the project. The new algorithm will be tested on wearable microcontroller-based solutions such as the Pebble smart watch (based on the ARM Cortex-M3 processor), with Bluetooth allowing connections to sensors. Alternatives, such as the newly released Microsoft band and the Apple watch that provide research development kits for sensor interfaces will also be evaluated.

***WP 4 (Prototype Systems for Validating Decision Making)*:** We will construct prototype decision systems for the case studies sufficient to allow us to compare recommendations with those of experts as a validation of the methods and algorithms from WP1 and WP3.

WP 5 (Medical applications): Whereas WPs 1-4 incorporate the medical applications as testbeds and validation of the methods, this work package focuses on the following tasks:

User Engagement: Early in the project, we will work with the case study stakeholders including GPs, physicians, surgeons, nurses and carers to model the current care pathways, including therapies, monitoring tools and disease progression. We will use participatory design techniques to explore scenarios for how the proposed technologies could modify existing care pathways. Early patient engagement will also be around the case studies via patient advisory groups formed from existing successful PPI groups, for example:

- There is a specific group for patients with inflammatory MSK conditions and a thriving general MSK PPI group who input to both service and research delivery from the design phase through to final dissemination.
- The North East London diabetes lay panel has been integral to the research pathway in diabetes for the last 10 years, advising on all aspects of studies including design, recruitment, development of lay material and dissemination.
- The NIHR cardiovascular Biomedical Research Unit (cvBRU) Public and Patient Advisory Group (PPAG) is well established under a patient-chair (Prof Vernon Trafford), routinely helping with patient information sheets and academic guidelines.

Later, we will also engage with wider patient groups such as Patient Opinion, UK eHealth Association and NHS Choices to explore the translation of our research into practice.

Regulatory Barriers: To ensure that the longer-term impact is not hampered by regulatory barriers (in addition to the safety aspects covered by the core research objectives) we have already engaged with key regulators such as the ethics committees, the Information Governance Alliance and the Caldicott Guardian and Digital Health lead for Barts Health.

Economic Impact: We will evaluate the potential economic impact of the proposed systems, using established early economic modelling techniques [1]. In consultation with healthcare professionals and patient groups, this will compare the current care pathway with an alternative that incorporates anticipated effects of remote monitoring based on plausible assumptions, with sensitivity analysis, about the performance of the decision algorithms. The analysis will support a business case to take applications of this type to a pilot study.

WP 6 (Management and Public Engagement): Fenton will have overall leadership, working with co-PIs to manage staff on each sub-package (see details in work plan and JoR). Co-ordination between technical and medical researchers will be managed using bi-monthly meetings, with frequent one-to-one interactions to bridge the gaps between disciplines. An **Advisory Group**, including the technology companies listed in Section 2, the EECS researchers, and medical staff, will meet twice a year to share progress and lessons learned and plan for impact.

Effective public engagement is essential if the public and patient groups are to understand the potential of technology and to play an informed role in shaping public opinion. In addition to three public workshops, we will produce an ‘insights booklet’ communicating the key results of the research, backed up by a website (also for the technical results). We will draw on our experience, e.g. the EPSRC CHI+MED project on safer medical devices (www.chi-med.ac.uk) and the tendon

research website (http://www.tendon.qmul.ac.uk/), designed to provide patients with advice as well as disseminating research. We will also produce a special issue of the cs4fn magazine.

6. References

[1] Annemans L., Genesté B, Jolain B. Early modelling for assessing health and economic outcomes of drug therapy. Value in Health 2000; 3: 427–434.

[2] Best, N., Ashby, D., Dunstan, F., Foreman, D., McIntosh, N., 2013. A Bayesian approach to complex clinical diagnoses: a case study in child abuse. J. R. Stat. Soc. Ser. A (Statistics Soc. 176, 53–96.

[3] Burton, C., Pinnock, H., McKinstry, B. Changes in telemonitored physiological variables and symptoms prior to exacerbations of chronic obstructive pulmonary disease. Journal of Telemedicine and Telecare, 2015, Vol. 21(1) 29–36

[4] Chen, Z. (2003) Bayesian filtering: From Kalman filters to particle filters, and beyond. Statistics, 182(1), 1-69

[5] Constantinou, A. C., Yet, B., Fenton, N., Neil, M., & Marsh, W. (2016). Value of Information analysis for interventional and counterfactual Bayesian networks in forensic medical sciences. Artificial Intelligence in Medicine 66, pp 41-52

[6] Constantinou, A., Freestone M., Marsh, W., Fenton, N. E., Coid, J. (2015) "Risk assessment and risk management of violent reoffending among prisoners", Expert Systems With Applications 42 (21), 7511-7529.

[7] Cooper, A. (2004) The Inmates Are Running the Asylum: Why High Tech Products Drive Us Crazy and How to Restore the Sanity, ISBN 0672326140, Pearson Higher Education

[8] Curzon, P. and Blandford, A. and Cox, A. and Thimbleby, H. (2015), "Safer Interactive Medical Device Design: Insights from the CHI+MED Project Proceedings of the 5th EAI International Conference on Wireless Mobile Communication and Healthcare, October, London, UK. ACM.

[9] Fenton NE, Neil M, and Caballero JG (2007), "Using Ranked nodes to model qualitative judgements in Bayesian Networks" IEEE TKDE 19(10), 1420-1432

[10] Fenton, N. and Neil, M. (2010). "Comparing risks of alternative medical diagnosis using Bayesian arguments." Journal of Biomedical Informatics, 43: 485-495

[11] Fenton, N.E. and M. Neil (2012), "Risk Assessment and Decision Analysis with Bayesian Networks". 2012, CRC Press

[12] Furniss, D., Masci, P., Curzon, P., Mayer, A., & Blandford, A. (2015). Exploring medical device design and use through layers of distributed cognition: How a glucometer is coupled with its context. J. of Biomedical Inf., 53, 330–341

[13] Greenland, S., 2006. Bayesian perspectives for epidemiological research: I. Foundations and basic methods. Int. J. Epidemiol. 35, 765–75.

[14] Haran Burrim, H (2013). Cardiac Pacing – Is Telemonitoring Now Essential? Arrhythm Electrophysiol Rev. 2013 Nov; 2(2): 95–98.

[15] Jackson, C.H., Best, N.G., Richardson, S., 2009. Bayesian graphical models for regression on multiple data sets with different variables. Biostatistics 10, 335–51.

[16] Jenny, U. *et al*. (2011) Piloting tele-monitoring in COPD: a mixed methods exploration of issues in design and implementation. Primary Care Respiratory Journal (2011); 21(1): 57-64.

[17] Masci, P and Curzon, P. and Thimbleby, H. (2015) "Early identification of software causes of use-related hazards in medical devices Proceedings of the 5th EAI International Conference on Wireless Mobile Communication and Healthcare, October, London, UK. ACM."

[18] Mushcab, H., Kernohan, W. G., Wallace, J., Martin, S. (2015) Web-Based Remote Monitoring Systems for Self-Managing Type 2 Diabetes: A Systematic Review. Diabetes Technology & Therapeutics. July 2015, 17(7): 498-509.

[19] Neil, M, Chen X, Fenton, N E, "Optimizing the Calculation of Conditional Probability Tables in Hybrid Bayesian Networks using Binary Factorization", IEEE Transactions on Knowledge and Data Engineering, 24(7), 1306 - 1312, 2012

[20] Neil, M., Tailor, M., & Marquez, D. (2007). Inference in hybrid Bayesian networks using dynamic discretization. Statistics and Computing, 17.

[21] Pinnock, H. *et al*. (2013) Effectiveness of telemonitoring integrated into existing clinical services on hospital admission for exacerbation of chronic obstructive pulmonary disease: researcher blind, multicentre, randomised controlled trial. BMJ 2013; 347:f6070

[22] Yet B, Perkins ZB, Rasmussen TE, Tai NR, and Marsh DWR (2014). “Combining Data and Meta-analysis to Build Bayesian Networks for Clinical Decision Support” Journal of Biomedical Informatics , 52, 373 – 385

[23] Yet B, Perkins ZB, Tai NR, and Marsh DWR (2014). “Explicit Evidence for Prognostic Bayesian Network Models” Studies in Health Technology and Informatics, 205, pp.53-57

[24] Yet, B., Perkins Z., Fenton, N.E., Tai, N., Marsh, W., (2014) "Not Just Data: A Method for Improving Prediction with Knowledge", Journal of Biomedical Informatics, Vol 48, 28-37