The CREDO stack: theory to practice in cognitive systems engineering

John Fox, Department of Engineering Science, University of Oxford

Abstract

CREDO is a framework for understanding and designing cognitive systems. It has evolved through a long program of empirical studies and engineering of practical AI systems in medicine, and it is now being used for an increasing range of applications. This short retrospective review summarises the main theoretical and technical results that underpin CREDO and outlines future research directions. [1]

Introduction

CREDO is the outcome of an extended programme of research on reasoning, decision-making and planning in medicine. The programme has drawn on a range of disciplines including empirical decision research and mathematical decision theory; knowledge and software engineering; logic and logic programming, and autonomous agent theory and multi-agent systems.

The work has taken medicine as a specific research focus but the goal has been to develop a generic theoretical foundation for understanding high level cognition and practical technologies for developing cognitive agents (Fox, 2013; 2014).

These varied perspectives have come together in a flexible framework called the CREDO stack (figure 1). This is a set of inter-related layers of concepts and techniques for modeling high level cognition in complex, knowledge rich domains. The stack includes an "end to end" set of development tools for building and deploying cognitive systems of diverse kinds.

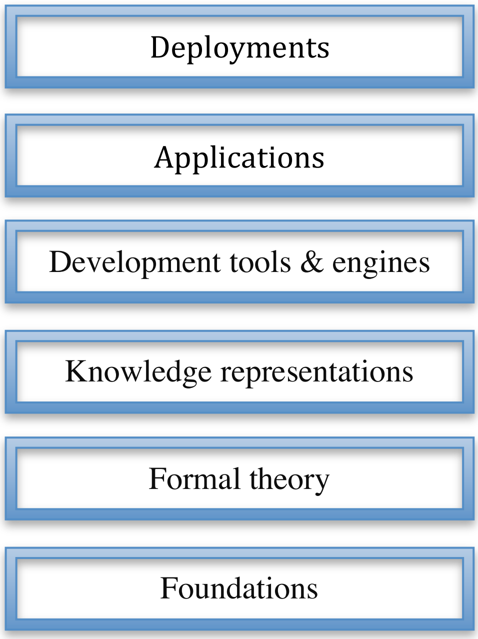

Figure 1: The CREDO stack

Medicine has always been fertile soil for research in cognitive science and systems engineering because of its complexity and the range of challenges that it throws up for theoreticians as well as developers. It has been a focus for research in cognitive science (psychology, AI), decision theory (normative and heuristic models), computer science (e.g. Bayesian networks, non-classical logics), and knowledge engineering (e.g. expert systems, ontology representation).

Medicine is also challenging because of the complexity and diversity of clinical tasks; the high levels of uncertainty that are ubiquitous

[1] **Acknowledgements** Too many people have made significant contributions to the CREDO research programme to mention them all here. However I would particularly like to acknowledge Rick Cooper (Birkbeck, London), Subrata Das (Machine Analytics, Boston), Paul Krause (Surrey University) and Simon Parsons (Liverpool University), and my present collaborators who have created so much of the current generation of technology: Ioannis Chronakis, David Glasspool & Vivek Patkar (Deontics Ltd.).

in clinical practice; the need to address multiple goals and constraints simultaneously and adapt to unexpected situations. A major challenge is the vast amount of background knowledge that may need to be called upon at any time in the clinical setting. Finally, an understanding of clinical expertise requires a unified view of cognitive capabilities which are individually major research areas in their own right, from reasoning and problem-solving to decision-making and planning, communication and learning.

I think that the main theoretical contributions of our work are a novel approach to reasoning under uncertainty and about evidence (*argumentation*), decision making (*symbolic decision theory*) and a general architecture for autonomous cognitive agents (*the domino architecture*). A key practical result is PRO*forma*, a declarative but executable specification language for modeling decisions, plans and other cognitive tasks. Technical details can be found in Fox and Das (2000).

A brief history of the CREDO stack

Foundations

The origins of the CREDO program were in laboratory studies of human decision-making and computer simulations of memory, reasoning and decision-making (e.g. Fox, 1980; Cooper and Fox, 1997). The studies suggested that human decision-making is at least as well described by a qualitative account as by more traditional quantitative decision models.

The early work made much use of *if…then…* production rule techniques to model medical reasoning and decision making, and we had some practical success with PROPS, a hybrid tool that combined data-driven production systems with goal-directed logic programming. PROPS evolved into X*i*Plus, an expert system package that was quite successful in commercial applications but proved to be less effective in medicine.

Rule-based systems still attract considerable interest for clinical decision-support systems though many have noted that they have significant limitations in practical use (e.g., Musen 1998; 2013).

A key question about rule-based and other knowledge-based expert systems was how to define sound theoretical foundations for symbolic decision methods since, despite their practical successes, the *ad hoc* nature of heuristic methods in many systems seemed to compare unfavourably with "rational" mathematical decision models.

While orthodox mathematical models offer a normative basis for making decisions they show little of the versatility and flexibility of human decision-making. In particular they say little about when and how decisions are to be initiated and what the options are ("framing the decision"), when domain-specific knowledge is relevant, or why, or how it may be used flexibly (and sometimes creatively).

To address these issues we developed a model for decision-making based on classical first-order logic with some extensions that subsumed four distinct patterns of reasoning:

- Reasoning about beliefs (e.g. about a clinical situation) and introducing goals;
- Generating candidates or options to resolve goals on the basis that there must be a valid argument for each candidate;
- Constructing sets of arguments for and against each option;
- Aggregating arguments pro/con each option to establish a preference order.

The scheme was dubbed *Symbolic Decision Theory* (SDT: Fox et al, 1990a; Fox and Krause 1991; Huang et al, 1993). Applications were implemented using logic programming techniques (Fox et al, 1990b) and we found SDT to be a versatile and extensible framework for implementing decision support systems in the medical domain. However it was an informal theory whose properties and limitations were unclear. The next step was to put these informal ideas on a sound formal footing.

Formalizing SDT

The primary step was to formalize the everyday concept of argumentation, a natural mode of reasoning in the presence of uncertainty, and reasoning about arguments. Argumentation can be formalized as a labeled deduction system and straightforward extensions to classical deductive axioms. LA, a Logic of Argument, is thought to have been the first argumentation model for decision making with a clear axiomatic treatment (Fox & Krause, 1992; Fox, Krause & Elvang-Горannson, 1993; Krause et al, 1995).

Another driver for formalising SDT was the goal of understanding decision-making in autonomous agents, a major research interest in AI, and particularly for agents that might be used in safety critical applications. The domino agent architecture shown in figure 2 (Das et al, 1997; Fox and Das, 2000) extended the 4 step decision model to include planning and plan execution.

The complete model brings together concepts of beliefs, goals, decision options, arguments, constraints and commitments within a framework of modal propositional temporal logics (Das et al 1997; Fox and Das, 2000).

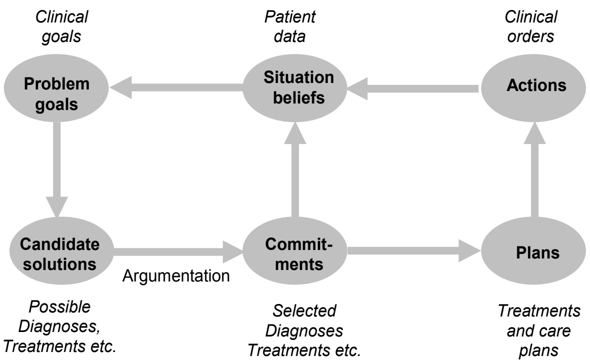

Figure 2: the domino agent architecture

The model has connections with the well-known BDI agent framework; a domino agent can reason about its current situation (beliefs), desires (goals) and intentions (plans) as events occur and situations change, but it can also generate multiple (tentative) beliefs, goals and plans, argue the merits of the alternatives, and decide which to commit to on the basis of its confidence and priorities.

The PROforma language

The domino model led to an extended logic programming language called R^2L (Das et al, 1997; Fox and Das, 2000). However we also wanted to be able to take advantage of techniques from other programming paradigms like object-oriented programming, and from software engineering, such as modularity of design, component reusability, software verification and testing etc.

R^2L evolved into the PRO*forma* language (Fox, 1997; Fox and Das, 2000) which we believe combines strengths of logic programming and agent-programming in an object-oriented framework (Fox et al, 2003). The key change in PRO*forma* was that the network of *inference processes* (arrows) in the domino model was reified into a small set of *objects* that formed an upper ontology of tasks: decisions, plans, actions and enquiries as shown in figure 3.

The central concepts of the PRO*forma* language are *tasks* (decisions, plans, actions and enquiries). These four classes are subclasses of a general concept called a "keystone" from which they inherit some general attributes including the *goal* of the task, *triggering* events, *pre-conditions* and *post-conditions* and other properties.

Each type of task also has class-specific attributes. These are used to specify the details of the method for carrying out a decision, plan or other task.

- *Decisions* are modeled in terms of a set of candidates or *options*, logical *schemas* for constructing arguments for and against options, and rule schemas for selecting or recommending one or more of the options in light of arguments.

- *Plans* are containers for sets of tasks that can be enacted sequentially or concurrently, and may (recursively) have sub-plans. Plans can also have termination and abort conditions.
- *Actions* and *enquiries* are tasks whose role is to communicate with the external environment, interacting with a user or another software component.

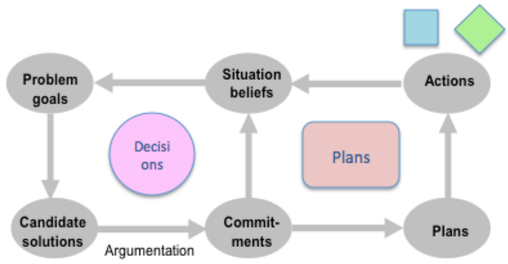

Figure 3: Reification of the domino model into "tasks". Arrows (logics) on the left of the domino are reified into a generic decision task, and the arrows (logics) on the right are reified as plans.

The syntax and semantics of PRO*forma* have been published as an open standard (Sutton and Fox, 2003) and it is the modeling language for several commercial application development tools (e.g. Arezzo®, Tallis[2]).

Applications

The earliest PRO*forma* applications were decision support systems for physicians in making routine medication decisions (Walton et al, 1997), genetic risk assessments (Emery et al, 1999; Coulson et al, 2001) and making decisions about whether patients required urgent referral to cancer specialists. Later developments focused on specialist care, particularly cancer care, including radiological image interpretation (Taylor et al, 1998); treatment of leukemia in children (Bury et al, 2005) and diagnosis and treatment of breast cancer (Patkar et al, 2006; Patkar et al, 2012). All trials to date have demonstrated measurable benefits in supporting better clinical decision making, and in some cases dramatic improvements over unaided clinician's judgments.

PRO*forma* applications in routine use include a national program for supporting decision making for a variety of common conditions (e.g. management of children with Asthma[3]) by general practitioners in New Zealand. The scalability and economic benefits of the technology are demonstrated by the patient triage "symptom checker" service run by the UK NHS[4].

A more complex application of the CREDO stack in oncology supports decision-making by multidisciplinary teams (MDTs). An early version of the system for assisting in the treatment of breast cancer provided recommendations and supporting argumentation for diagnosis, imaging, staging, treatment, prognosis and trial eligibility decisions. It has been shown to make recommendations that comply with best practice significantly more frequently than the MDT alone (Patkar et al, 2012).

Deployment

The last layer of the CREDO stack is a collection of services for deploying PRO*forma* applications. Although the primary purpose of the language was to formalize decisions and other tasks in a machine interpretable format, it was also anticipated that naturalistic task ontology and argument-based decision procedure might facilitate design of intuitive user interfaces.

The simplest arrangement is to map each class of task to a particular style of web page. For example a *decision* is visualized as a single page with a set of options, each of which can be expanded to show the arguments for and against each option, and each argument can be selected to review the logical structure and backing evidence for it. An *enquiry* maps to a page that provides an

[2] ; www.openclinical.org/gmm_proforma.html

[3] http://www.bestpractice.net.nz/news_art_childhoodAsthma_nov2010.php

[4] Reported at www.infermed.com (case study at http://www.infermed.com/en/Clinical-Decision-Support/Case-Studies/NHSDirect.aspx accessed at 3/12/2013).

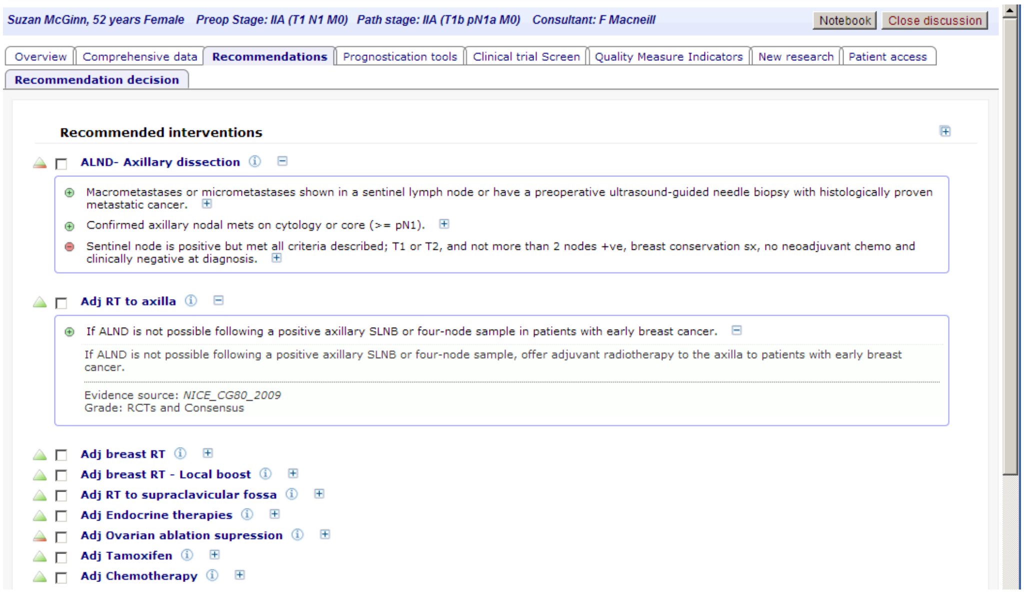

Figure 4: Set of "recommendations" in a CREDO breast cancer application. All possible treatment options have been evaluated in the light of what is believed about the patient and all the arguments for and against each option. The list shows the recommended subset of chemotherapy, radiotherapy, surgery and other interventions. Two options have been expanded to show arguments for (green dot) and against (red dot). Each argument is linked to the "backing" evidence that is the basis of the argument. The tabs at the top of the screen provide access other services.

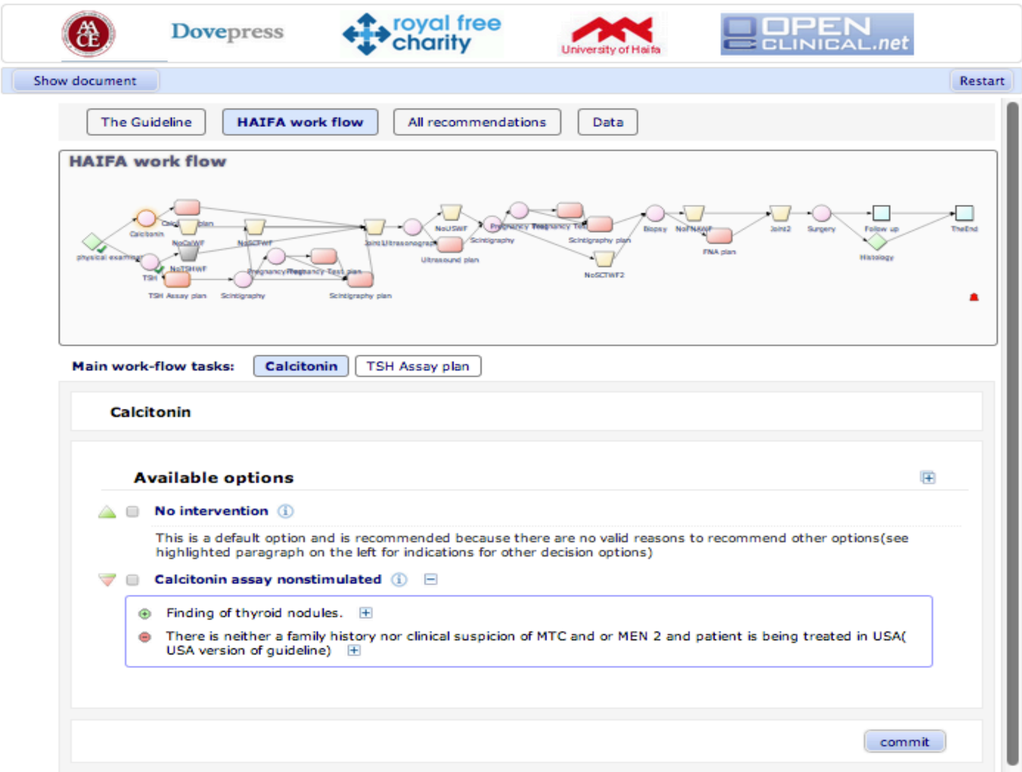

Figure 5: A "care pathway" for the diagnosis and treatment of patients presenting with thyroid nodules (developed for the American Association of Clinical Endocrinologists (AACE). The top panel is a visualization of a PRO*forma* plan showing the clinical pathway using the standard PRO*forma* task iconography.

electronic form for data entry, and an *action* might map to a page that presents an alert or a message to the user using a variety of widget types. Finally a *plan* can be visualized as a collection of tasks, where a workflow is visualized as a pathway of enquiries, decisions, plans and actions organized as an executable network (see figures 5 and 6).

The detailed presentation style of any task UI can be customized as required but the basic task ontology and process modeling framework have proved to be an effective basis for automatically generating interaction dialogues for many different kinds of decision support services. PRO*forma* was used process representation language for a spoken language interface to decision support systems (Beveridge and Fox, 2006) - see demo at http://www.openclinical.org/dm_homey.html#.

Conclusions and future aims

CREDO appears to be an effective framework for understanding human expertise and for engineering cognitive systems. A lasting practical success has been the PRO*forma* language for specifying and implementing autonomous agents. The language has proved to be very versatile and is an important element of many systems that are being routinely used in a variety of clinical settings. The task models used in PRO*forma* are "clinician friendly" in two senses: the task ontology facilitates the design and implementation of intuitive and flexible user interfaces, and clinicians can themselves design and test PRO*forma* knowledge bases using the CREDO toolset.

The software components of the stack include knowledge authoring and testing tools, execution and test engines, and web deployment software (servers and clients). The software is constantly evolving and being refined to improve application delivery capabilities but the underlying theory and language have remained largely unchanged since they were originally developed.

Reflections on the CREDO stack

Despite these results, however, our 20+ years of experience building software and applications have made us aware of ways in which all the elements of the stack could be improved. For example argumentation theory has become mainstream in AI and computer science and there are significant results to be exploited. Knowledge representation techniques (notably ontology modelling and description logics) and the theory and design of autonomous agents and multi-agent systems have also advanced considerably.

Experience building diverse medical applications has also suggested ways of improving the stack, particularly the PRO*forma* decision model and some changes and extensions to the task ontology would be desirable. A general area offering both theoretical and technical advances is in understanding the (large) family of computational architectures that could instantiate the domino agent model.

From CREDO to ARIA

ARIA represents our aspirations to begin a new research programme whose aim is to build on the CREDO experience and to explore a number of new directions which we believe will advance our understanding of cognitive agents and ability to design them.

Agent models in AI commonly adopt an explicit representation of agents' *mental states* such as beliefs, desires and intentions. CREDO uses an extended set including *situation states* (beliefs, goals, plans, arguments, preferences, commitments) and *task states* (pending, in_progress, supended, terminated, aborted). An important line of enquiry in ARIA would be to investigate a richer model and explicit representation of the meaning, context, history and rationale for such states (e.g. the reasons for and expected consequences of the agent's current beliefs and plans).

This enriched state model will allow an ARIA agent to be able to *reflect* upon what it believes *and why* at all times, and what it intends to do *and why*. We believe that this meta-cognitive ability is a major source of human intelligence and adaptability, yet to our knowledge there is little theory and few tools for building and deploying reflective agents.

We aim to systematically investigate variants of the ARIA approach by developing a new family of execution engines, with meta-cognitive capabilities built into the architecture. The project will also carry out a systematic study of meta-cognitive computation in which cognitive states are modeled as ontologically constrained meta-types. A number of recent papers have begun to explore this possibility (Fox et al, 2007; Fox, Cooper and Glasspool, 2012; Fox, 2014 (in press)).

Given the length of time that was required to get to the current level of capability in the CREDO programme, however, it seems likely that the scale of the challenge and the range of issues and expertise that will be needed to make major progress on the ARIA programme is beyond the capabilities and resources of a single group. My colleagues and I are therefore interested in meeting colleagues who share our interest in sound theories of cognitive systems and developing practical technologies for intelligent agents to explore opportunities for collaboration.

References

Beveridge M, Fox J “Automatic generation of spoken dialogue from medical plans and ontologies” *Journal of Biomedical Informatics,* Volume 39 Issue 5, October 2006 Pages 482 – 499

Bury J, Hurt C, Roy A, Cheesman L, Bradburn M, Cross S, Fox J, Saha V. LISA: a web-based decision-support system for trial management of childhood acute lymphoblastic leukaemia. *Br J Haematol.* 2005 Jun;129(6):746-54.

Cooper R, Fox J Learning to make decisions under uncertainty: The contribution of qualitative reasoning *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, 1997

Coulson A, Glasspool DW, Fox J, Emery J. “RAGs: A novel approach to computerised genetic risk assessment and decision support from pedigrees” *Methods of Information in Medicine* 2001; 40; 315-322

Das S, Fox J, Elsdon D, Hammond P “A flexible architecture for a general intelligent agent*”, Journal of Experimental and Theoretical Artificial Intelligence*, 9, 407-440, 1997.

Das S “Symbolic Argumentation for Decision Making under Uncertainty” in *7th International Conference on Information Fusion (FUSION),* 2005

Elvang-Goransson M, Krause P and Fox J "Acceptability of arguments as logical uncertainty" in M Clarke, R Kruse and S Moral (eds) Symbolic and Quantitative approaches to reasoning and uncertainty *Proc. European Conference ECSQUARU 93*, Berlin: Springer Verlag, 79-84.

Emery J, Walton R, Murphy M, Austoker J, Yudkin P, Chapman C, Coulson A, Glasspool DW, Fox J, “Computer support for interpreting family histories of breast and ovarian cancer in primary care: comparative study with simulated cases” *British Medical Journal*, 321, 28-32, 1st July, 2000.

Fox, John Making decisions under the influence of memory. *Psychological Review,* Vol 87(2), Mar 1980, 190-211

Fox, J., Myers, C. D., Greaves, M. F., & Pegram, S. “Knowledge acquisition for expert systems: Experience In leukaemla diagnosis” *Methods of Information in Medicine,* (1985) 2 65-72.

Fox J, Clark D, Glowinski A and O'Neil M "Using predicate logic to integrate qualitative reasoning and classical decision theory" *IEEE Trans. on Systems, Man and Cybernetics*, 20 (2), 347-357 (1990a).

Fox J, Glowinski A, Gordon C, Hajnal S, O'Neil M Logic engineering for knowledge engineering: design and implementation of the Oxford System of Medicine, *Artificial Intelligence in Medicine,* Volume 2, Issue 6, December 1990, 323–339

Fox J, Krause P J, Symbolic decision theory and autonomous systems Proceeding UAI'91 *Proceedings of the Seventh conference on Uncertainty in Artificial Intelligence* Pages 103-110 San Francisco: Morgan Kaufmann Inc. 1991)

Fox J, Krause P, Qualitative frameworks for decision support: lessons from medicine *The Knowledge Engineering Review* Volume 7 Issue 01, March 1992, pp 19-33

Fox J, Krause P J and Elvang-Goransson M "Argumentation as a general framework for uncertain reasoning", *Proc. Int. Conf. on Uncertainty in Artificial Intelligence,* Washington, 1993

Fox J "Dissemination of Medical Knowledge: the PRO*forma* approach" *Artificial Intelligence in Medicine 1997*

Fox J and Das S, *Safe and Sound: Artificial Intelligence in Hazardous Applications*, AAAI/MIT Press, 2000

Fox J, Patkar V, Thomson R, "Decision support for health care: the PROforma evidence base" *Informatics in Primary Care*, 14, 2006.

Fox J, Beveridge M, Glasspool D "Understanding intelligent agents: analysis and synthesis" *AI Communications* Volume 16 Issue 3, Pages 139 - 152, August 2003

Fox J, Glasspool D, Modgil S "A canonical agent model for healthcare applications" *IEEE Intelligent Systems: Special Issue on Intelligent Agents in Healthcare*, 21-28, Nov/Dec 2006.

Fox J "Arguing about the evidence" in *Evidence, Inference and Enquiry* (Editors: Philip Dawid, William Twining, Mimi Vasilaki) `The British Academy (2011)

Fox J, Cooper R, Glasspool D W "A canonical theory of dynamic decision-making" *Frontiers in Psychology*, 02 April 2013

Fox J "Credo for decision science and engineering" *Axioms* 2014 (accepted subject to revision).

Huang J., Fox J., Gordon C., and Jackson-Smale A., (1993) "Symbolic decision support in Medical Care", *Artificial Intelligence in Medicine* 5, pp 415-430

Krause, P., Ambler, S., Elvang-Goransson, M. and Fox, J. (1995) A logic of argumentation for reasoning under uncertainty, Computational Intelligence, 11, 113{131

Musen M "Modern architectures for intelligent systems: reusable ontologies and problem-solving methods. *Proc AMIA Symp*. 1998:46-52.

Musen M: The knowledge acquisition workshops: A remarkable convergence of ideas. *International Journal of Human-Computer Studies*, 71 (2), February 2013, 195–199

Patkar V, Hurt C, Steele F, Love S, Purushotham A, Williams M, Thomson R and Fox J "Evidence-based guidelines and decision support services: a discussion and evaluation in triple assessment of suspected breast cancer *Br J Cancer*. 2006 December 4; 95(11): 1490–1496

Patkar V, Acosta D, Davidson T, Jones A, Fox J, Keshtgar M "Using computerised decision support to improve compliance of cancer multidisciplinary meetings with evidence-based guidance *BMJ Open* 2012; 2.

Sutton D and Fox J, "Syntax and semantics of the PRO*forma* guideline modeling language", *JAMIA* 2003

Taylor P, Fox J Pokropek A. The development and evaluation of CADMIUM: a prototype system to assist in the interpretation of mammograms, *Medical Image Analysis*, Volume 3, Issue 4, 1999, pp. 321-337.

Tural C et al. Clinical utility of HIV-1 genotyping and expert advice: the Havana Trial. *AIDS*. 2002; 16:209-218.

Walton RT, Gierl C, Yudkin P, Mistry H, Vessey MP, Fox J. Evaluation of computer support for prescribing (CAPSULE) using simulated cases. *British Medical Journal* 1997 Sep 27;315(7111):791-5.