Ivan Jureta, PhD

RESEARCH OBJECTIVES

My principal research interest is in method engineering and method automation, focusing on the elicitation, modeling and analysis of knowledge that human experts apply in problem solving and decision making, the engineering of ontologies and processes capturing that knowledge, and the automation of the said processes. This interest falls within the various research fields concerned with the transfer, preservation and automation of knowledge.

Study of problem solving and decision making knowledge of human experts underlines numerous research problems and high-impact applications. For example, how do we elicit, represent and automate methods of human experts and make them available to millions in interactive format (which is not the case for web pages, books, scientific publications, and more generally, text, images and video)? How do we transform an expert's method into algorithms which can automatically deliver recommendations/advice at a scale that is elusive to the expert? How do we evaluate recommendation algorithms against methods of human experts (e.g., trading algorithms, fraud detection algorithms)? How do we evaluate the quality of an expert's problem solving and decision making independently from their reputation? How do we evaluate the quality of existing systems engineering methods, and how do we create new systematic methods based on expert engineer's knowledge? How do we create methods for incorporating into Business Intelligence systems the knowledge of expert decision makers, so that the systems can deliver advice?

This general research interest translates into the following research goals:

1) To construct, refine and empirically test novel conceptualizations, mathematical models and methods for the elicitation, modeling and analysis of knowledge that underlies problem solving and decision making methods, specifically methods applicable to settings characterized by unavailable/incomplete quantitative estimates of probability and/or utility and the availability of variously imprecise, vague, incomplete, conflicting, and unstable/changing qualitative decision information and advice from many stakeholders. Examples include methods applied to deliver complex services, such as an engineer applies when designing an information system, a sports coach when advising an athlete on how to train, an architect when designing an airport, and so on.
2) Apply conceptualizations, models and methods to method engineering in specific domains and empirically evaluate the relevance of the resulting methods through industry transfer, as I believe that the empirical evaluation of method engineering and method automation needs to be performed under industry-specific resource constraints.
3) Form an interdisciplinary research group focused on the elicitation, modeling and analysis of knowledge that underlies problem solving and decision making methods.
4) Create a graduate study program specializing computer science students for the elicitation, modeling and analysis of problem solving and decision making knowledge, preparing them thereby to lead the engineering of future decision-support and recommendation systems.

I believe that the success of computer systems which aim to automate problem solving and decision making methods of human experts depends to an important extent on how well the expert's knowledge is elicited, understood and modeled. Based on my research experience, I also believe that it is not relevant to search for universal conceptualizations, models and methods for knowledge elicitation, modeling and automation, but instead work on domain-specific conceptualizations, models and methods, and when possible, abstract from these the method engineering guidelines which systematically reappear as relevant across domains.

CURRENT ACHIEVEMENTS

My past and current research focused on methods for the elicitation, modeling and analysis of knowledge used in information system engineering and which may come from any stakeholder (buyer, user, engineer, etc.) of the information system. This knowledge is traditionally referred to as the *requirements of the system*, in software and requirements engineering research and practice.

I have been working on the so-called requirements problem: the essence of the problem is how to combine the various stakeholders' beliefs, desires, intentions, and evaluations about an organization's objectives and current information systems, with the system engineer's (i.e., advisor's) own expertise and experience in designing information systems, so as to jointly make design decisions that would change available information systems towards the achievement of the organizational objectives. In management science the problems conceptually closest to the requirements problem have been studied in decision analysis (i.e., applied decision theory). Below is a list of my main results. References are omitted in this section - please consult my publication list.

(1) Core ontology for Requirements Engineering and the Requirements Problem: I gave a precise and novel definition of the requirements problem. I pointed out the various limitations in the understanding of the requirements problem that has been dominant over the course of more than a decade, and have offered solutions to these limitations. I presented this in a paper entitled "Revisiting the Core Ontology and Problem in Requirements Engineering", at the 16th IEEE Int. Requirements Engineering Conf. The contribution was formally recognized at the conference, as we received the best paper award (the award is given to one paper per year). It was necessary to expand on these ideas to provide a more detailed account precisely in order to explain as best as feasible to the affected communities the consequences of the new ontology and definition of the requirements problem, which led to a journal paper entitled "A core ontology for requirements", published in Applied Ontology. My coauthors and I presented and argued for a new and core ontology for decision making in requirements engineering, precisely to avoid confusions over the meaning of terms, which in decision analysis literature at least are very real (cf., e.g., Ronald Howard's recent discussion, tellingly entitled "Speaking of Decisions: Precise Decision Language", in Decision Analysis (Vol.1, No.2) of the meaning of fundamental terms, such as subjective probability, confidence, uncertainty, and so on). In contrast to models of choice in decision analysis, our core ontology focuses less on (but it does not abstract from) quantitative utility and (subjective) probability estimates and instead emphasizes qualitative tools (in particular: conceptual analysis and informal and mathematically formal argumentation, the later modeled via non-monotonic and paraconsistent mathematical logics used with different kinds of order relations) for the modeling of the requirements decision problem. This change of emphasis is due to the well-known practical difficulties in the elicitation of quantitative utility and probability estimates, along with contentions of the philosophical foundations of the subjective probability concept, and the criticism of expected utility theory (and its variants) as a model of a rational decision maker.

(2) Techne: The new definition of the core ontology for requirements engineering and of the requirements problem led to research into formalisms for the modeling and analysis of knowledge identified in the core ontology for requirements and relevant for solving the new requirements problem. This led to Techne, a formal language for the modeling of the problem and solution space of the requirements problem, published at the 2010 IEEE Int. Conf. on Requirements Engineering (RE'10), and subsequently the work on the identification of solutions (published, respectively at the 2010 Int. Conf. on Conceptual Modeling, ER'10, and the 2011 Int. Conf. on Requirements Engineering, RE'11). The key departure made in Techne, compared to prior work on formalisms for requirements engineering, is that the logic in Techne is designed from scratch to fit the structure of our core ontology for requirements and of the requirements problem, rather than being designed by only using the ontology for requirements to organize formulas, e.g., of a liner temporal logic. Ongoing research on Techne consists of expanding the original Techne formalism into a family of formalisms, each of which is appropriate for the resolution of a particular class of requirements problems.

(3) Multi-party validation of requirements: An important problem in decision making for requirements engineering, which Techne does not cover, is how to determine if the various organizational stakeholders agree on the interpretation of the information elicited about the organization and its business objectives; it is only if they do agree that such information be input in the requirements problem formulation to establish the particular requirements problem that the organization should solve. I designed a versatile method for the analysis of discussions arising in joint decision making, i.e., discussions that the advisors have with organizational stakeholders in order to agree on a decision. It generalizes some of my previous results on this problem (my "Clear justification of modeling decisions for goal-oriented requirements engineering", Requirements Engineering Journal, Vol. 13, No. 2). The results appeared in the paper "Analysis of Multi-Party Agreement in Requirements Validation", presented at the 2009 Int. Conf. on Requirements Engineering (RE'09).

(4) Analysis and design of advice: I wrote the monograph Analysis and Design of Advice, published by Springer in 2011. The principal aim is to present and discuss the contributions mentioned above within a much broader problem of how advice can be conceptualized and analyzed, and what its role is in coordination systems (namely, in management science and economics). I have made the case that the modeling and analysis of requirements does amount to a particular case of advice-making and giving, and that the contributions we have made up to now are sufficient to warrant a synthesis and a more general discussion. It is in effect possible to restructure the require-

ments problem in order to define the "advisor's problem", be it a consultant, a recommendation system, or otherwise, then reuse and expand on the methods we have been exploring up to this point, and which we continue to refine and improve. Roughly speaking, the advisor's problem is a decision problem that involves the elicitation and reconciliation of knowledge coming from various stakeholders, and the design of recommendations on a course of action towards the satisfaction of specific objectives arising through knowledge elicitation. The approach is different from that of decision analysis in management science, as any interested reader will see from our publications cited above. One difference is, e.g., the reliance on argumentation models to deal with uncertainty, and the subsequent limited role that quantitative probability estimates play in the advisor's decision problem formulation and the search for its solution; this limits the usual need in decision analysis to come up with probability estimates, which often tend to be very difficult to rigorously come up with. A manager's role is very often one of a coordinator within the organization, so that she will necessarily often dispense advice to employees in order to ensure coordination. It follows that it is relevant to introduce students to the ideas on the analysis of advice that I developed in this research, and rendered more accessible in the manuscript. Together with Prof. Stéphane Faulkner, I created a 5 ECTS lecture for third year management students at the Louvain School of Management, at the University of Namur, entitled "Decision Making & Requirements Engineering". Parts of the manuscript serve as the lecture notes, along with compulsory and optional readings of classical papers in decision theory and decision analysis, and requirements engineering. After the positive feedback from the students who attended that lecture from September to December 2009, I taught it in 2011 as well.

(5) Other achievements: All of my publications are either directly or indirectly related to the definition of the core ontology for requirements, of the requirements problem, and of the conceptual and mathematical tools and methods for the resolution of the requirements decision problem. Such publications focus on complementary approaches to these three topics (e.g., my "Clear justification of modeling decisions for goal-oriented requirements engineering", Req. Eng. J., Vol. 13, No. 2), their discussion and criticism (e.g., "A More Expressive Softgoal Conceptualization for Quality Requirements Analysis", 25th Int. Conf. Conceptual Modeling), automation (e.g., "Continually Learning Optimal Allocations of Services to Tasks", IEEE Transact. Serv. Comput., Vol. 1, No. 3), and use in related fields (e.g., "A comprehensive quality model for service-oriented systems", Softw. Qual. J., Vol. 17, No. 1).

RELATIONSHIP BETWEEN CURRENT ACHIEVEMENTS AND RESEARCH OBJECTIVES

The following hypothesis relates my research on requirements engineering methods and my research objectives: Engineering a method from problem solving and decision making knowledge of experts requires the mobilization of similar principles and guidelines, and is based on similar theoretical foundations as the effort of making a method for requirements engineering. This is not to say that any problem solving or decision making method is the same as any requirements engineering method, but instead that one can approach the problem of making a problem solving and decision making method in the same way that one approaches the making of a method for requirements engineering.

In making both requirements engineering methods, and methods for problem solving and decision making, one is obliged to answer the following questions. Firstly, what are the categories of information that the expert problem solver or decision maker uses and what advice does she produce trough her method. I.e., what ontology of information the expert has (the Ontology question). Secondly, how the expert represents the information and advice (the Modeling question). Thirdly, how the expert analyzes information (the Reasoning question). And finally, how the expert transforms the information, produces and delivers advice (the Advice question). These same questions apply when trying to recreate the method of an engineer who focuses on determining the right system requirements, or of a sports coach who delivers advice to athletes, or of an adviser in foreign policy.

The hypothesis stated above is that answering these questions in different domains mobilizes common theoretical principles, drawing from ontology engineering, formal language design, and process engineering. I say drawing from, because it is today not clear how the contributions of the said fields can help in answering the Ontology, Modeling, Reasoning and Advice questions. My primary research objective is consequently to construct, refine and empirically test novel conceptualizations, mathematical models and methods for the elicitation, modeling and analysis of knowledge which answers, for a given problem solving and/or decision making method, the said Ontology, Modeling, Reasoning and Advice questions. Answers to these questions can subsequently be used to evaluate if and how the expert's method can be automated, and the advice it produces delivered at a large scale.

VISION FOR THE FUTURE

Consider a running coach who advises her athletes how to train. One may be able to learn about this coach's problem solving and decision making method by being trained or taught by that coach. If the coach published a book on the method, one may read and attempt to apply it. If the coach speaks publicly about her method, one can attend the lectures to get a glimpse of the principles, practices, rules, guidelines that the coach applies to produce advice.

Transferring and preserving knowledge in the described manner is common, but falls short in several important respects: the coach has a limited time to transfer her knowledge to others; the coach may not be thinking much about the structure, rules and principles of her method, about their consistency, completeness, and whether they are detailed enough; text, but also images and video, as formats for transferring knowledge are not interactive - if one believes having learned the coach's method from a book, one can hardly verify if the advice she subsequently gives does indeed satisfy the method; and so on.

Such problems of transferring, preserving and applying knowledge to generate advice are more prominent today, when there is an abundance of data. Data itself is not particularly interesting. If we assume, for simplicity, that interpreted data is information, information still may not be enough to inform action. Namely, turning information into advice requires applying a method to the information, so as to determine if and how that information is relevant for the decision problem one is facing, and what course of action can be recommended based on that information. Knowing that one had run very fast today and that one is tired is information, but not advice. A coach can produce advice based on this information through her problem solving and decision making method.

Today's Internet is data and information. Enabling the production of advice based on that information requires not only sophisticated and massive statistical or otherwise analysis of the data, but also ways to automate and deploy at a large scale the methods of human experts. Doing so in turn requires research into means for answering systematically the Ontology, Modeling, Reasoning and Advice questions for a wide range of domain-specific problem solving and decision making methods.

To move closer to my objectives, my research is organized along four dimensions:

1) Ontology question: Engineering of domain-specific ontologies for problem solving and decision making, consisting of the specialization and extension of contributions in ontology engineering to the common concepts and relations relevant in problem solving and decision making, such as the concepts of alternative, goal, and relations of refinement, conflict, preference;

2) Modeling question: Definition and evaluation of symbolic and visual syntaxes of formal languages for the representation of instances of the problem solving and decision making ontologies, and the study of if and how to map expressions in these languages to well-known logics;

3) Reasoning question: Definition and evaluation of desirable properties, such as domain-specific consistency and completeness, of models made using the formal languages, as well as the definition of what inconsistency means in a method and what should be done when inconsistency is detected in a model (e.g., if it should be eliminated or tolerated, and if the latter, then what exactly it means to tolerate it and until when).

4) Advice question: Definition and evaluation of formal properties of proof theory and semantics for formal languages used to represent instances of the problem solving and decision making ontology for a given method, to fit the expectations on what conclusions should be drawn from information to which the method is applied to generate advice.

I have worked in my current and past research on each of these topics specifically for requirements engineering methods. My research on method engineering for requirements is theoretically grounded and spans different areas of computer science and knowledge representation and reasoning. I look forward to approaching the Ontology, Modeling, Reasoning and Advice questions in problem solving and decision making beyond requirements engineering, and towards enabling large scale preservation, transfer and automation of knowledge on problem solving and decision making.