

Belief change in rational agents: Perspectives from Artificial Intelligence, Philosophy, and Economics

Dagstuhl seminar 05321, August 8-12, 2005

Executive summary

1 Motivation

1.1 Introduction

The area of belief change studies how a rational agent may maintain its beliefs when obtaining or perceiving new information about the environment. This new information could include properties of the actual world, occurrences of events, and, in the case of multiple agents, actions performed by other agents, as well as the beliefs and preferences of other agents. Not surprisingly, this area has been of interest to researchers in different communities.

The initial research in belief change came from the philosophical community, wherein belief change was studied generally from a normative point of view (that is, providing axiomatic foundations about how rational agents should behave with respect to the information flux). Subsequently, computer scientists, especially in the artificial intelligence (AI) and the database (DB) communities, have been building on these results. Belief change, as studied by computer scientists, not only pays attention to behavioural properties characterising evolving databases or knowledge bases, but must also address computational issues such as how to represent beliefs states in a concise way and how to efficiently compute the revision of a belief state. More recently, the economics and game theory community, in particular the emerging field of cognitive economics, has become active in belief change research, adopting a normative point of view, like philosophers, but paying more attention to the “cognitive plausibility” or “fitness” of the belief change operators.

The goal of the seminar was to bring together researchers from these areas, allowing for the identification and addressing of problems of common interest in

this area, as well as providing a means to explore ways in which one area may contribute to another.

1.2 Goals and Content of the Seminar

As mentioned, the area of belief change can be regarded as originating in the philosophical logic community. This work provided abstract, formal, and precise specifications of desirable properties for belief change operators, as well as the identification of distinct types of change. However, this research says nothing about specific implementable operators nor computational issues – issues of fundamental importance to computer scientists. Researchers in artificial intelligence and computer science have followed up on these latter issues, as well as developed other specific operators (addressing e.g. sensor fusion and belief base merging) and examined their complexity characteristics. In artificial intelligence, the relatively recent emergence of the field of cognitive robotics, which is concerned with endowing artificial agents with cognitive functions that involve reasoning, for example, about goals, actions, the states of other agents, collaboration and negotiation, etc., has given impetus to the development of computational operators for belief change and the identification of issues arising from concrete, evolving sets of knowledge. As well, more recently, economists have been using work in belief revision, and applying it to notions of mistaken and changing beliefs among interacting and negotiating agents. Such work is also of obvious interest to researchers in artificial intelligence.

Previously, there has been limited interaction among these communities. Clearly there are deep problems of common interest, and results in one area will contribute to another. We have already mentioned that research in economics has made use of the work from the philosophical community, and that such results will be of use to researchers in AI. As well, contributions may also flow back from economics to research in the foundations of belief revision: For example, recently it has been suggested that that economic principles (dealing with choice, preferences, and utility) may provide a more appropriate foundation for belief change. Computational issues raised and addressed by researchers in computer science and AI will be of use to economists addressing related problems; as well such work can contribute to the other areas by further elucidating the abstract area of belief change, as well as providing implementations and identifying philosophically-interesting “pragmatic” or “practical” problems.

Thus we saw the Dagstuhl seminar as providing a forum where researchers in three broad areas (philosophy and logic, artificial intelligence and computer science, and economics and game theory) could address highly related (in some cases, the same) problems, in which work in one area could benefit research in another. Hence for the Dagstuhl seminar, we felt that there would be valuable interactions and contributions that would be anticipated by bringing people together in these areas.

2 The participants

The seminar brought together 37 researchers from the following countries:

Europe Austria (1), France (8), Germany (8), Israel (1), Italy (1), Luxemburg (1), The Netherlands (1), UK (2);

Americas Brazil (1), Canada (2), USA (3), Venezuela (1);

Asia & Pacific Australia (3), Japan (1), New Zealand (2)

Five of the participants were PhD students. Slightly more than one half of the participants come from computer science (mainly Artificial Intelligence), the other ones coming from philosophy, mathematics (formal logic) and economics. The different backgrounds of the researchers resulted in stimulating discussions on various issues.

3 The program

We organized the program so as to have rather homogeneous sessions, each day being focussed on one or two particular issues.

The focus of the first day was on *philosophical and logical foundations of belief change*. The talks of Hans van Ditmarsch and Wiebe van der Hoek together gave an overview of the role of dynamic epistemic logic for representing belief revision (including the multiagent case). Then we had an invited talk by Isaac Levi, who gave deep philosophical insights on various kinds of degrees of belief. The talks by Alexander Bochman and by Olivier Schulte addressed foundational issues of belief expansion, revision and contraction, while the talks by Ken Satoh and Alexander Nittka both concerned meta-level reasoning about rational agents performing belief revision.

The second day was dedicated to *belief merging*. The talks by Paolo Liberatore and by Torsten Schaub introduced new approaches to belief merging: Liberatore argued that inconsistency between knowledge bases is due to mistakes and therefore proposes to merge belief bases by minimizing the set of these mistakes; Schaub's talk proposed two new specific belief merging operators: a first one consisting in computing a maximal set of formulas comprising the joint knowledge of the knowledge bases, and a second consisting in projecting all belief bases onto a specified belief base. The joint talks by Gabriella Pigozzi and Daniel Eckert gave an overview of judgement aggregation, and the role that belief merging can play in solving this problem. These talks on judgement aggregation also discussed the connections between belief merging and social choice. Such connections were also central to the talks by Pierre Marquis and by Ramon Pino-Pérez, who both addressed the strategical aspects (manipulability) of belief merging. Lastly, the talk by Sébastien Konieczny addressed connections

between belief merging and some game-theoretical issues.

On the third day, we had only four talks, on a variety of subjects (the afternoon was left free so as to give participants some time for discussions). The talk by Bernard Walliser addressed the value of information of communication in multi-agent belief revision settings. The talk by Giacomo Bonanno defined a new epistemic logic based on three modal operators representing initial beliefs, information, and revised beliefs. The talk by Jan Broersen addressed some comparisons between epistemic and deontic changes. Finally, the talk by Didier Dubois was concerned with belief revision and abduction in a numerical setting, where uncertainty is modelled by a family of probabilities.

On the fourth day there was a focus on *iterated belief change*. The joint talks by Michael Thielscher and Yi Jin presented the axiomatic foundations and the computational properties of a framework for action and (iterated) belief revision. The talk by Abhaya Nayak addressed iterated expansion and/or contraction. The talk by Hans Rott considered prioritized belief bases encoding compactly an agent's belief state, and showed how the shifting of priorities in such bases could be used for representing a large variety of methods for changing belief states. The talk by Jonathan Bennaïm focussed on distance-based revision (a subclass of AGM revision operators), for which he gave an impossibility result. The talk by Guillaume Aucher considered the introduction of actions and probabilities into the belief revision process. The talk by Thomas Meyer discussed different approaches to resolving inconsistencies in ontologies, specified in description logic, suggesting that belief revision has a role to play in a new and promising application field: the semantic web. Lastly, the day was concluded by a panel on the new insights in belief change.

During the last half-day we had two talks on action and belief change: one by Steven Shapiro on belief change with acting and sensing in the situation calculus, and one by Andreas Herzig, on belief update. We also had a talk by Aditya Ghose on outcome-driven revision, a talk by Emil Weydert on revision by conditionals and a talk by Eduardo Fermé, who presented new maps between different kind of base contraction functions. The session ended up by a general brainstorming discussion.

4 Conclusion

We found the workshop successful, especially on the following two achievements: first, the seminar made participants aware of a commonality of interests across different disciplines; second, it suggested new directions for research that will probably be taken up by researchers in the next couple of years.

Where is the field going? We can mention at least two emerging issues:

- the field is broadening both with respect to theoretical underpinnings, and

so begining in incorporate notions from game theory and social choice theory. As well, it is broadening wrt application areas, moving beyond traditional areas in AI and database systems, to including areas in description logics, the semantic web, and in economics.

- As well, there is an emerging focus on epistemic notions having to do with communicating, negotiating, competing, and collaborating agents in belief change. Dynamic epistemic logic seems to have an important role to play here.

Moreover, it looks like belief merging and iterated belief revision are still hot topics and will remain so for the next few years.

For the future, we plan first a followup workshop, and therein we will further encourage interdisciplinary work.

The organizers:

James Delgrande
Jérôme Lang
Hans Rott
Jean-Marc Tallon