The Reasoner

Volume 2, Number 3
March 2008

www.thereasoner.org
ISSN 1757-0522

Contents

§1 Editorial

I was delighted when Jon asked me to be guest editor of the March edition of The Reasoner; this was several months ago, March has come around all too quickly and Jon is still waiting for copy! Several Reasoner editorials have started out with apologies and this one is no exception. My task was to find a suitable candidate to interview and since The Reasoner is an interdisciplinary journal, this provided an opportunity to explore some of the latent philosophical views of a researcher who does not publish in overtly philosophical journals. I interviewed a computer scientist who generally works in applied AI, an area close to my own heart. During the interview, the opinion was expressed that continuing research on reasoning is necessary in order to maintain progress on the design of intelligent machines. Where are these interdisciplinary researchers of the future to come from?

The graduate students I work with, expected to become part of the next generation of innovators, almost without exception, have no interest in philosophy; it is our job, as educators, to stimulate that interest. My own area, probabilistic reasoning, has been very much in the news recently, with meetings on Bayesian Nonparametric Regression (Isaac Newton Institute, Cambridge) and Combining Probability and Logic (Centre for Reasoning, University of Kent). By encouraging our students to attend meeting such as these, we can hope to produce better educated graduates, ready to push forward the current boundaries of research.

The undergraduate students I teach are generally more open to new ideas. The educational issues are broad and not just limited to my Department, but I believe that curriculum re-design is a serious issue for us all. The new interdisciplinary MA in Reasoning at the University of Kent UK leads the way by providing the background required for graduate research in all aspects of reasoning.

It became clear during the interview, that reasoning was considered to provide the theoretical foundation for all applied work. Indeed, that when progress ground to a halt, it was the logical and philosophical concepts involved that were questioned. It was refreshing to hear

that philosophy as well as logic, is considered important in this work.

It gives me great pleasure to introduce Professor Lakhmi C. Jain.

Dawn E. Holmes
Statistics and Applied Probability, University of California Santa Barbara

§2
Features

Interview with Lakhmi C. Jain

Brief Profile

Lakhmi C. Jain is Professor of Knowledge-Based Engineering in the Division of Information Technology, Engineering and the Environment, School of Electrical and Information Engineering, University of South Australia. His research interests include the creation of knowledge-based intelligent information and engineering systems incorporating neural nets, fuzzy systems, evolutionary computing, e-learning and multi-agent paradigms. Professor Jain is also interested in the application of knowledge-based intelligent information and engineering systems in various fields including aviation, business, management, decision science, defence, education, engineering, and science and health sciences.

Lakhmi is a prolific author and has been publishing in the most prestigious journals in his areas of interest since the early 1990's. He is the series editor for many of Springer's prestigious collections and co-edited eight volumes in 2007 alone! Lakhmi is the Founding Director, Knowledge-Based Intelligent Engineering Systems Centre at University of South Australia. In 1997 he became founding editor of the International Journal of Knowledge-Based Engineering Systems, which has since blossomed into a top journal.

Interview

Dawn Holmes: Lakhmi, was your doctorate in computer engineering? How did you first get into computer science and, in particular, reasoning under uncertainty as an area of research? When was your interest in knowledge-based systems first aroused?

Lakhmi C. Jain: I did my PhD in Electronic Engineering. I have always been interested in designing machines which can to some extent mimic human behavior. I began by automating the design of electronic systems by the use of readily available Expert System Shells. These early designs greatly influenced my thinking. I became very impressed by the power of the knowledge base and the inference engine of the expert systems. I soon realized the limitations of expert systems and of Artificial Intelligence in general. I started by thinking on how we could deal with the problem of uncertainty. I had always believed that reasoning must inevitably play a major role in the implementation of practical intelligent systems.

DH: Some of your work is theoretical, some of it is applied. For example, in 1996 you presented the paper 'A Fuzzy Wheelchair Controller' at the Fourth International Conference on Soft Computing. How important do you consider is an understanding of the philosophical concepts in general and of logic in particular to the understanding of artificial reasoning?

LJ: The theoretical foundation is undoubtedly extremely important but I consider it to be equally important to verify as far as possible any theoretical study results by the use of practical examples. In our present society we have many examples such as roads, transport, hospitals, schools all of which could possibly benefit by the application of an appropriate Artificial Intelligence during the design process and to the improvement of their performance during operation. I attempt to stimulate interest among designers, researchers and operators in these vital areas of our civilization. I have used the results obtained by Artificial Intelligence studies in areas as diverse as, Wheel chairs, Explosive detection for security purposes, and Breast Cancer Diagnosis for example. These studies were also educational for me as they served to make me better realize the importance of the logical and philosophical concepts involved.

DH: Do you consider that the Early Promise of Artificial Intelligence has been realized?

LJ: I am of the opinion that too much was expected from Artificial Intelligence in the early phases of its development which led to a degree of disillusionment with the new technology. I am very impressed by the vision of the many great researchers which include people such as John McCarthy, Marvin Minsky, Nils Nilsson, Daniel Bobrow and Bruce Buchanan. Their efforts are especially noteworthy when the limitations are considered. These limitations imposed by the early technology were formidable. They included the limitations of early computing, the limited computing power available and the limited availability of computers. The early promise of AI has now to a very great extent been realized.

DH: What is your long term research goal? Is your area of research growing?

LJ: My long term research goal has not changed yet. It is still the 'Design and Development of Intelligent Machines which can in a limited sense mimic human behavior'. There are many intelligent people who do wonders in the area of intelligent systems. It is a matter of creating enough interest among this group to pro-

vide the opportunities for the presentation of the new research results, discussion and publications.

Consequently, my colleagues and I initiated the following program.

- The Knowledge-Based Intelligent Engineering Systems Centre (KES) to provide applied research support to the Information, Defence and Health Industries. The KES research activities are mainly focused on modeling, analysis and design in the area of Intelligent Information Systems, Physiological Sciences Systems, Electronic Commerce and Service Engineering. The overall goal is to fuse contributions from researchers in diverse disciplines such as Engineering, Information Technology, Science, Philosophy, Cognitive Science and Psychology.

- The first International Conference on Knowledge-Based Intelligent Engineering and Information Systems (KES) was in 1997. The first three KES conferences were held in Adelaide. The KES conference has now attained full international status under the leadership of Professor Robert Howlett. The fourth conference was in the UK. The Fifth conference was in Japan, the Sixth conference in Italy, the Seventh in the UK, the Eighth in New Zealand, Ninth in Australia, Tenth in the UK, Eleventh in Italy in 2007. The Twelfth KES is scheduled to be held in Croatia in September 2008. This conference is now an established annual international event.

In addition, we launched several journals, including the International Journal of Knowledge-Based Intelligent Engineering Systems (KES) in 1997 and the International Journal of Hybrid Intelligent Systems in 2004, all under the auspices of the IOS Press, and several book series.

I believe that there is a need to expose all graduates to the computer science discipline, and especially the intelligent aspect of machines. It will be a wonderful achievement to convince our graduates that it is possible to design Intelligent Machines which can even at present mimic human behavior in a limited but impressive sense.

DH: Finally Lakhmi, in your view, what is an important open problem in knowledge-based systems?

LJ: Although I have put tremendous efforts towards achieving my goals I have not made sufficient progress towards realizing my dreams. I believe that there is still a need for ground breaking research on Reasoning and Learning.

Translating Utterances, Reporting Beliefs

Kripke's well known puzzle about belief can be portrayed as a puzzle about reasoning. We observe Pierre, a monolingual and untravelled Frenchman perusing in Paris many pictures of London and delightedly exclaiming 'Londres est jolie'. So we (readers of *The Reasoner*) accept:

1. Pierre believes that London is pretty.

Next, we observe Pierre move to London (not realising it to be the town he has been calling 'Londres'), pick up English through immersion, and, on the basis of what he sees, dejectedly and repeatedly mutter 'London is not pretty' (a verdict with which his new neighbours heartily concur). So we accept:

2. It is not the case that Pierre believes that London is pretty.

We reckon ourselves to be good reasoners, not liable to accepting flat-out contradictions. Yet here we are accepting 1. and 2.!

In response to this problem Cristian Constantinescu suggested that 'Londres', in Pierre's idiolect, does not mean 'London' and thus that we are wrong to translate his delighted exclamation as 'London is pretty' and hence wrong to use 1. to report his belief. I criticised this on the grounds that, as Kripke insists, Pierre uses 'Londres' in the same way as his compatriots, and 'London' in the same way as the English do. [I'll use '(C1)', '(G1)' and '(C2)' to label, respectively, Constantinescu's first piece, my reply and his rejoinder in *The Reasoner*, 1(4):8-9; 1(5):4-5; 1(7):5-6, 2007.]

There is, I believe, a perfectly straightforward solution to Kripke's puzzle that can be brought out by considering which of the following is the more accurate report of the belief that Pierre held while in Paris:

1a. Pierre believes that London (the parts he has seen in pictures) is pretty.

1b. Pierre believes that London (the whole city) is pretty.

According to me, then, even if I translate Pierre's Paris exclamation as 'London is pretty', this does not constrain me to use 1. to report his belief, because generally a person gives only rough verbal expression to his/her beliefs. Constantinescu, by contrast, sees the solution to Kripke's puzzle as turning on a principle about translation, namely that it must preserve cognitive content. Looking at his two examples of translation in (C2) will be instructive.

Discussing the first of these, Constantinescu says that 'we would normally have no qualms' about translating inside its quotation marks. The example is:

La ville qu'aujourd'hui nous appelons 'Londres' est située sur le site d'un campement Romain ancient, appelé 'Londinium'.

His translation is

The city that we nowadays call 'London' is situated on the site of an ancient Roman settlement called 'Londinium'.

But Constantinescu thinks that a good translation, apart from preserving cognitive content, must also preserve truth-value—and clearly his translation here may not, for the original is about the habits of French speakers ('nous') the proposed translation about the habits of English speakers ('we'). It is easy to imagine tokens of both sentences being used in some possible world in which English speakers don't use 'London', but call London 'The Smoke'. Indexicals pose a problem for translators—how, for example, should we translate 'Cette phrase est en anglais' so as to preserve truth-value?

Constantinescu's second example is designed to show that, for some non-quotational contexts, translating each word would make for a bad translation. His example (we imagine Pierre saying this on discovering that what he now calls 'London' is what he previously called 'Londres'):

Je viens de découvrir qu'on peut être simultanément à Londres et à London.

Constantinescu's point is that, if we translated that 'Londres' as 'London' our translation would, misleadingly, have 'Pierre foolishly rejoicing in the discovery of a trivial *a priori* truth', and so he recommends that, in translating, we leave the 'Londres' untouched. The first problem with this is that what we would end up with is not a translation but a mish-mash of English and French. (The original is too, but the speaker can get away with it because he is bilingual—which the person for whom we are providing the translation may not be.) Second, according to Constantinescu, the utterance in question 'is clearly a statement about objects, not names'. Well, the utterance is about one city under two different names. To see this, try translating it for the benefit of a monolingual Irish person. It would not do, for Constantinescu's reason, to render *both* the 'Londres' and the 'London' into Irish. The best translation (for help with which, I thank Arthur Keaveney) would be:

Tá mé díreach anois tar éis a fháil amach gur féidir le duine bheith ag an am chéana san áit ar a dtugtar 'Londres' na Fraincís agus san áit ar a dtugtar 'London' i mBéarla.

which translates back into English as:

I have just discovered that one can be at the same time in the place called 'Londres' in French and in the place called 'London' in English.

In this case, the translator has to look beyond the words of the speaker to the (probable) content of his utterance.

What this shows is what is known to all professional translators—that a *faithful* translation cannot always be *concise*, and the same is true *a fortiori* of a faithful report of someone's belief, for here what we are trying to convey is the content of that belief. Hence 1a. or 1b. in contrast to 1. If it is 1a. that faithfully reports Pierre's belief, then there is no inconsistency between that and

2a. It is not the case that Pierre believes that London (the parts he has seen in London) is pretty.

On the other hand, if it is 1b. that faithfully reports Pierre's belief, then, if his London-based belief is also about the whole city, we are correct to report that belief as

2b. Pierre believes that London (the whole city) is not pretty.

With 1b. and 2b., we are consistently reporting Pierre's beliefs, and these are inconsistent due to the fact, as I pointed out in (G1), that his inductive reasoning is *merde*.

Laurence Goldstein
Philosophy, Kent

How We Reason: A View from Psychology

Psychologists have studied reasoning for at least a century. But, for sixty years or so, they had no proper theory of what individuals are doing when they reason, or of the underlying mental processes, which are inaccessible to introspection. Computers made theorizing about reasoning feasible and respectable, and psychologists have developed several such theories, especially of deduction. One theory is that we are all equipped with formal rules of inference akin to a logic in a "natural deduction" formulation, see, e.g., Rips, L. (1994: *The Psychology of Proof*, Cambridge, MA: MIT Press). Reasoning on this account is a search for a proof leading from premises to conclusion. Another theory is that the probability calculus describes how we ought to reason and how in fact we do reason even deductively, see Oaksford, M., and Chater, N. (2001: The probabilistic approach to human reasoning, *Trends in Cognitive Sciences,* 5, 349–357). This theory describes well the results of certain psychological experiments, yet other experiments have shown that untrained individuals do distinguish between necessary conclusions and probable conclusions. Indeed, Louis Lee and Geoffrey Goodwin have shown in unpublished studies of Sudoku puzzles that naïve individuals soon realize that their solution depends, not on probabilities, but on valid deductions.

One difficulty for theories based on formal rules is that reasoning in daily life depends on the logical form, not of sentences in natural language, but of the propositions that they express. Hence, the use of formal rules depends on recovering the logical form of propositions.

For example, in the sentential calculus, an inference of this form is valid:

- If p then not q.
- q.
- Therefore, not p.

But, not surprisingly, individuals balk at this inference:

- If Jane played a game then she didn't play soccer.
- Jane did play soccer.
- Therefore, she didn't play a game.

They know that soccer is a game, and therefore that the conditional premise is consistent with only two possibilities, shown here on separate lines:

- Jane played a game. Jane didn't play soccer.
- Jane didn't play a game. Jane didn't play soccer.

The logical form of the first premise is therefore: (p or not p) and not q. Its recovery is a headache, and the general analysis of the logical form of propositions expressed in natural language is beyond any existing algorithm.

If my colleagues and I are correct, there is no need to recover logical forms and no need to search for proofs. Our theory postulates instead that individuals use the meanings of propositions and general knowledge to construct a set of mental models representing the possibilities consistent with the premises, see Johnson-Laird, P.N., and Byrne, R.M.J. (1991: *Deduction*, Hillsdale, NJ: Erlbaum) and Johnson-Laird, P.N. (2006: *How We Reason*, Oxford: Oxford University Press). If a conclusion holds in all these possibilities, then individuals infer that it is valid. And they are also able to show that an inference is invalid, not by searching in vain for its derivation, but by constructing a counterexample, i.e., a model of a possibility consistent with the premises but not with the conclusion.

One prediction of our theory is that the greater the number of models needed to make an inference, the harder it will be. For example, ask yourself what follows from these two disjunctions:

- Ann is in Atlanta or Ben is in Birmingham, or both.
- Ben is in Birmingham or Cate is in Clemington, or both.

Intelligent but untutored individuals usually overlook at least one of the five possibilities consistent with these premises, and their conclusions often describe only one of them.

A fundamental principle of the theory is that mental models represent what is true, but not what is false. Hence, given the first premise above, individuals enumerate the following three possibilities:

- 1. Ann is in Atlanta.
- 2. Ben is in Birmingham.
- 3. Ann is in Atlanta. Ben is in Birmingham.

where, for example, the falsity of first disjunct in the second possibility is not represented explicitly, see, e.g., Johnson-Laird, P.N., and Savary, F. (1999: Illusory inferences: A novel class of erroneous deductions, *Cognition*, 71, 191–229). This so-called principle of *truth* reduces the processing load on working memory, but, as we discovered from a computer program implementing the theory, it has a devastating effect on certain seemingly simple inferences. Consider this problem, for instance:

- Either Jane is kneeling by the fire and she is looking at the TV, or else Mark is standing at the window and he is peering into the garden.
- Jane is kneeling by the fire.

Does it follow that she is looking at the TV?

Most individuals say, "yes", see Walsh, C., and Johnson-Laird, P.N. (2004: Co-reference and reasoning. *Memory & Cognition*, 32, 96–106). Given the first premise, they think of two possibilities: in one, the first conjunction is true; and in the other, the second conjunction is true. They overlook that when the second conjunction is true, the first conjunction is false, and that one way in which it can be false is when only its first clause is true, i.e., Jane is kneeling by the fire but *not* looking at the TV. Hence, the correct answer to the question is: "no".

Invalid inferences of this sort are endemic, occurring in all domains of reasoning. Untutored individuals represent what is true (rather than what is false), what is possible (rather than impossible), what is permissible (rather than impermissible, except in the case of overt prohibitions), and what are instances of a concept (rather than non-instances). It is as though for them what is false etc. ceases to exist. Any theories, including those based on formal rules or on the probability calculus, that fail to predict this phenomenon have quite a bit of explanatory work to do. There is, of course, more to the model theory than I can describe in this outline—it extends to probabilistic reasoning, induction, and abduction. Likewise I am indebted to more colleagues than I can name here.

P.N. Johnson-Laird
Psychology, Princeton

Arithmetic and Logic Incompleteness: The Link

We have all often read that the incompleteness of second order logic is a consequence of Gödel's incompleteness result. How the former follows from the latter is, however, not so often explained. Here we offer an easy account as a nice example of meta-theoretical reasoning based on model theory.

In 1930 Kurt Gödel proved that first order logic (fol) is complete. More precisely, Gödel proved an equivalent to the statement that if q is a first order sentence which is a logical consequence of a set $\{p\}$ of first order sentences then q is derivable from $\{p\}$ by the rules of inference of fol. Fol can be presented as a natural deduction calculus, i.e. consisting solely of rules of inference. The soundness of fol guarantees that if the set $\{p\}$ consists of truths / logical truths, any q derivable by fol from $\{p\}$ will be true / logically true as well.

In 1931 Gödel proved that the system PA consisting of the Peano Axioms [Richard Dedekind (1888)-Giuseppe Peano (1889)] for arithmetic with rules of inference appropriate for the language, if omega-consistent, is incomplete. A system S is omega-consistent if and only if there is not a predicate P for which S proves every sentence of the form 'Pn' as well as the sentence 'There is an x which is not a P'. Rosser reduced the condition of omega-consistency to simple consistency in 1936. This implies that, if PA is consistent, there will be arithmetic truths, i.e. arithmetic sentences that are true under the standard interpretation, which PA will not prove. Furthermore, PA cannot be completed by any number of additional axioms if certain normal desiderata are obeyed. We shall henceforth take the condition of PA's consistency for granted.

A diversion: The most substantial axiom of PA is the axiom of mathematical induction which is not a first order sentence. It is expressible as the second order sentence: '$(P)\{[P0\&(x)(Px \Rightarrow P(x+1))] \Rightarrow (x)Px\}$'; where '$x$' ranges over numbers and 'P' ranges over their properties. It states that for any property P, if it is true that 0 has it and it is also true that if any number x has it so does its successor $x + 1$, then so does every number. We may mimic the axiom of mathematical induction by presenting it as a first order axiom schema, thus: '$[F0\&(x)(Fx \Rightarrow F(x + 1))] \Rightarrow (x)Fx$'; where '$F$' is a schematic letter or place holder for a formula expressing a property. The axiom schema allows us to state an infinite number of induction axioms. These axioms are instances of what is known as the "full" axiom of mathematical induction. Nevertheless, all those instances are, even taken together, weaker than the single "full" axiom.

The Peano system PA1 (which is often referred to as just "PA") consists of the Peano Axioms except for the full axiom of induction which is replaced by the induction schema. Since PA1 consists of first order sentences its logic is fol. Hence every logical consequence of the axioms of PA1 is a theorem of PA1. Gödel however proved that there is a true arithmetic sentence which is not a logical consequence of the axioms of PA1. Hence there is a true arithmetic sentence which is not provable in PA1, that is, PA1 is incomplete. Historically he proved incompleteness for PA with the full axiom of induction. But then it holds for the weaker PA1 as well. Thus there will be some arithmetic sentence q (true under the standard interpretation) such that the axioms of PA1 will be consistent with q as well as with its (standardly false) negation $\sim q$. In other words there will be a PA1 model satisfying the (standardly true) sentence q and a PA1 model satisfying the (standardly false) sentence $\sim q$. Of course, q will not keep in this model the meaning it has in the standard one.

The Peano system PA2 consists of the Peano Axioms with the full induction axiom. Since the induction axiom requires a second order language, the required logic for PA2 is Second Order Logic (sol). We can think of sol (similarly to fol) as a set of inference rules; we can then ask whether sol is also complete i.e., whether sol will allow us to derive all logical consequences of any set of second order sentences. The answer is 'no' and it follows from Gödel's incompleteness theorem and a theorem by Dedekind. Dedekind proved in 1888 that all models of PA2 are isomorphic, i.e. structurally identical to the set of natural numbers. Therefore all arithmetical truths are satisfied in all models of PA2. Hence all truths of arithmetic are logical consequences of the axioms of PA2. That is, there can be no model of PA2 in which an arithmetic sentence q is true, and yet another in which q is false. This is a consequence of PA2's possessing the stronger "full" induction axiom. Gödel however showed that there is an arithmetical truth q which is not provable in PA2. Hence the logic of PA2, i.e., sol, is incomplete.

Laureano Luna
IES Doctor Francisco Marín, Siles, Jaén, Spain

Alex Blum
Philosophy, Bar-Ilan University

§3 News

Formal Models of Norm Change, 29–30 November

Last November the University of Luxembourg hosted a workshop on Formal Models of Norm Change. I evaluate the contributions by grouping them together and putting the groups in a, I hope, self-explanatory order.

Norm Creation

In 'A Logic for Social Norms Resulting from Conflicting Group Preferences' Jan Broersen, Rosja Mastop and Paolo Turrini argued that to understand the reason for why a norm emerges or is useful for a society, one has to understand in what sense norms are related to the social preferences and abilities of coalitions of agents. This relation is investigated by defining an operator for social rational choice, and by considering situations where what is rational for sub-groups may conflict with what is rational for the whole group. The presence of reachable outcomes that are optimal for the whole group gives rise to a social norm saying that sub-groups should not pursue their own best interest if that conflicts with what is best for the group.

Formal Theories of Change

Six contributions to the workshop dealt directly with the problem of how to model (norm) change.

The first contribution we discuss is 'The Statics of Rule Update' by Alexander Bochman. A central idea here is that in as far norms can be represented by defeasible rules, the problem of updating rules with new rules is a problem of 'priority' between defeasible rules. This reduces the problem of the dynamics of norms to a static one (hence the title), viz. the problem of finding a correct semantics for prioritized defeasible conditional theories. Many proposals for intuitive semantics were discussed.

Like the contribution discussed above, 'From AGM to Input / Output Contraction' by Gabriella Pigozzi and Leon van der Torre was about about dynamics of rules. The idea is to investigate to what extent Makinson's contraction postulates (AGM) for propositional theories apply to conditional theories such as Input / Output logic. The main conclusion so far is that the 'recovery' postulate of AGM can only be made to work for specific Input / Output logics, showing that contraction can easily be 'too strong' when applied to rules.

Fenrong Liu's contribution 'Exploring Dynamics in Preference' demonstrated the correspondence between two perspectives on change: change as modeled by relational changes in dynamic modal formalisms, and change as priority change in first-order theories on which objects are preferred over other objects. In an elegant way, this result connects two lines of formal research on change she explored in her PhD thesis.

'Legal Modifications in Temporal Defeasible Logic' by Antonino Rotolo investigates the dynamics of normative (legal) systems in a temporal defeasible logic framework. Legal norms are parametrized by a number of different temporal dimensions like time of validity, time of references, etc. Meta-rules are used to implement the idea that legal systems themselves should specify how and under what conditions they should change.

Rosja Mastop contributed 'The logic of prescriptions: free choice permission and normative gaps'. This contribution can be seen as a coalition logic account of Veltman's update semantics as applied to normative concepts. Natural accounts of free choice permission and normative gaps were given. In update semantics the dynamics is not so much in the norms themselves as in the way that given static norms can be understood, that is, as updates on sets of actions / choices agents are permitted or obliged to do. But, of course this mechanism can also be directly interpreted as a method to update permissions and obligations in the light of a new obligation.

Norm Acceptance

Another five contributions to the workshop can be classified as dealing with the dynamics of how agents cope with existing norms. So, the dynamics is not in the norms, but in their acceptance. For instance, in 'Abstracts Normative System Games' by Thomas Ågotnes, Wiebe van der Hoek and Mike Wooldridge, the central idea is to see the decision of accepting or rejecting a norm as a game played with other agents being subject to the same norm. A normative system is modeled here as a set of illegal state transitions. Furthermore, each agent is assumed to have a prioritized list of goals represented as formulae of Computation Tree Logic (CTL). Game theoretic properties and the computational complexity of related decision problems were discussed.

In 'Autonomous Multi-Agent Systems that Change: a Framework and Some Challenges' by Luca Tummolini and Emiliano Lorini (joint work with Cristiano Castelfranchi, Dominique Longin and Benoit Gaudou) the logic AL (Acceptance Logic) is proposed in which agents reason differently depending on which group they consider themselves part of. This enables one to characterize normative facts (i.e., obligations, prohibitions and permissions stemming from social conventions) and institutional facts. One of the issues studied is how normative dynamics interacts with a process of group acceptance revision.

In 'Some Problems with Regulations' by Laurence Cholvy a logic for the merging of possible conflicting regulations was proposed. Furthermore, it was discussed how the proposed logic might deal with 'incompleteness' of regulations and how this relates to the dynamic perspective.

In 'Obligation Change with Flexible Beliefs: when obligations are not always mandatory', Clia da Costa Pereira focused in on the influence of belief change on the acceptance of norms. The idea is that since beliefs come in grades, also acceptance of norms comes in grades. The set of accepted obligations of an agent is

modeled as a set of 'weighted' elements. For goal generation, three different ways for comparing fuzzy sets of obligations were presented.

Finally, in Souhila Kaci's contribution 'Merging Rules' a survey of existing merging operators developed in belief merging was presented, and it was discussed to what extent these operators can be adapted to operators for rule merging. One of the suggestions put forward is that we can use the stratification of a defeasible rule base into a prioritized rule base as the basis for a merging algorithm.

Abstracts and slides are available at http://icr.uni.lu/normchange07/.

Jan Broersen
Intelligent Systems, Utrecht

Second Indian Winter School on Logic, 14–26 January

In continuation with the logic events that were held in the past few years in India, the Second Indian Winter School on Logic was held at the Indian Institute of Technology (IIT), Kanpur during January 14-26, 2008. It was organized under the aegis of the Association for Logic in India (ALI) and the Association for Symbolic Logic (ASL); and was funded by the Research I Foundation, Department of Computer Science.

As the name suggests, this was the second in the series of Indian Winter Schools on Logic to be held biennially, with the first one held at IIT, Bombay in 2006. The effort initiated with the organizing of the First Indian Conference On Logic and its Relationship with other Disciplines (ICL) at IIT, Bombay in 2005. It was followed by the second ICL at IIT, Bombay in 2007 and the International Conference on Logic, Navya-Nyaya and Applications at Jadavpur University, Kolkata in 2007. ICL, too, constitutes a biennial series of conferences, with the third one to be held in January, 2009. All these events have paved the way for a grand coming together of the logic community in India, which was already existing for so many years in certain parts of India, not to mention the ancient traditions of Indian Logic. ALI, a product of this attempt, was formed in 2007 with this winter school, its founding event.

Logicians from all over the world as well as from different parts of India participated in this two-week school, resulting in a rich flow of ideas, blossoming into a huge network of interactions. This year saw a large participation of students from all over India, mostly coming from the disciplines of computer science, mathematics and philosophy. A variety of logic courses were taught over the fortnight starting at the introductory level and then moving onto the recent advances in each of these areas.

Mai Gehrke, Ramon Jansana and Alessandra Palmigiano gave a thorough introduction on algebraic logic, moving to relational semantics via cannonical extensions and Sahlqvist theory. Gregory Wheeler gave an illuminating overview on probabilistic logic. In the philosophical logic section, Chhanda Chakraborti spoke about ontologies, whereas Ranjan Mukhopadhyay discussed various philosophical issues about truth. Jean-Yves Beziau introduced the ideas of universal logic. An interesting intermingling of bisimulation, markov processes and logic was exemplified in Prakash Panangaden's lectures. Rohit Parikh put forward his enlightening views over a variety of topics including games, finite and infinite dialogues, and human and animal rationality. Eric Pacuit gave a thorough introduction to foundations of game theory. Agatha Walczak-Typke, Benedikt Löwe and S.M. Srivastava delved into the realm of set theory to a great depth. Various ways of dealing with uncertainties in knowledge representation were discussed in minute details by M.K. Chakraborty and Didier Dubois.

In an open session meant for the participating delegates, there was a great opportunity for them to present some problems that they are currently interested in. At the end of the day the participants went back greatly enriched.

On the penultimate day, a discussion convened by R. Ramanujam, convener of the steering committee of ALI about the logistics of the third ICL and the various ways in which efforts could be taken to further build up a stronger logic community in India was held, followed by a cultural presentation by the speakers and the participants. In conclusion thanks are due to the co-ordinators Mohua Banerjee and Anil Seth of IIT, Kanpur, who left no stone unturned in making the school a huge success.

Sujata Ghosh
Indian Statistical Institute, Kolkata

Logic of change, change of logic

The Prague International Colloquium is held once a year, and each edition is dedicated to a topic of interest to both logicians and philosophers. It brings together leading experts on the subject, as well as active researchers in the domain, to discuss current issues and challenges. This year's colloquium, Logic of change, change of logic, organized by Ondrej Majer, Michal Pelis and myself, will take place between the 10^{th} and the 14^{th} September. The theme is attitude change.

Tools from logic and mathematics have played a central role in models of human beliefs, of human desires and preferences and indeed the actions which are based on them. However, beliefs, preferences and perhaps even desires change. Thus the development, which

has been greatly accelerated in recent times, of extensions of the logical and mathematical techniques to account for the problems of change. However, as different paradigms (AGM theory and dynamic logic in the 'logic' camp, Bayesian update and Jeffrey conditionalisation in the probability camp, to take just a few examples of theories of belief change) jostle to impose themselves, it is perhaps the moment to take a step back and ask: what do we want from a theory of attitude change?

This question—as philosophical and methodological as it is technical—is at the heart of this year's colloquium. The aim is to bring together specialists working on the problem of attitude change, from a wide range of paradigms, to present and discuss their views on the objectives for theories of change. The ambition is to identify the main issues for theories of change, and clarify the major positions one could hold concerning the project of understanding or modelling attitude change. Such reflection is essential for the future development of the domain; anyone interesting in contributing to and benefiting from the discussion is very welcome to join us.

Brian Hill
Groupe HEC, Jouy-en-Josas, France

Calls for Papers

HYBRID LOGIC: Special Issue of the Journal of Logic, Language and Information, deadline 1 March.

MACHINE LEARNING IN SPACE: Special Issue of the Machine Learning Journal, deadline 31 March.

MULTIPLE SIMULTANEOUS HYPOTHESIS TESTING: Special Issue of the Journal of Machine Learning Research, deadline 31 March.

INFORMATION FUSION: Information Fusion in Public Health Informatics and Surveillance, special issue of Information Fusion, deadline 30 May.

CONDITIONALS AND RANKING FUNCTIONS: Special issue of Erkenntnis, franz.huber@uni-konstanz.de, deadline 31 May.

CAUSALITY AND PROBABILITY IN THE SCIENCES

Deadline 1 July

PROBABILISTIC MODELS FOR IMAGE UNDERSTANDING: Special Issue of the International Journal of Computer Vision, deadline 21 July.

DEPENDENCE ISSUES IN KNOWLEDGE-BASED SYSTEMS: Special Issue of International Journal of Approximate Reasoning, deadline 15 September.

§4
INTRODUCING ...

In this section we introduce a selection of key terms, texts and authors connected with reasoning. Entries will be collected in a volume *Key Terms in Logic*, to be published by Continuum. If you would like to contribute, please click here for more information. If you have feedback concerning any of the items printed here, please email thereasoner@kent.ac.uk with your comments.

Thomas Bayes

Thomas Bayes (ca. 1702-1761) was a British Presbyterian minister, theologian and logician. His 'Essay Towards Solving a Problem in the Doctrine of Chances' (1764)—published posthumously by his friend Richard Price in the 'Philosophical Transactions of the Royal Society of London'—contains the statement of a special case of what is known today as Bayes' theorem. In the essay, Bayes deals with the chance of events in connection to pre-existing circumstances and after the occurrence of particular events—which he termed 'prior odds' (or probability) and 'posterior odds', respectively. Bayes proposes that evidence confirms the likelihood of a hypothesis only to the degree that the evidence would be more probable with the assumption of the hypothesis than without it. Although his name is nowadays connected to a number of interpretations of probability that share the idea that probabilities express degrees of beliefs rather than frequencies, it remains unclear whether Bayes himself would have endorsed such an understanding or, instead, put more emphasis on observable entities and events.

Matteo Morganti
IHPST, Paris

A priori / A posteriori

A proposition is knowable a priori if one can know that it is true without appeal to experience. In order to know that bachelors are unmarried men I do not have to interview various bachelors; I just have to understand the terms 'bachelor' and 'unmarried man'. In contrast, a proposition is knowable a posteriori if it can be known on the basis of experience. That sugar is sweet is knowable a posteriori because I can come to know this by tasting it.

Dan O'Brien
Philosophy, Warwick & Birmingham

§5
Letters

Dear Reasoners,

In the last Reasoner Ed Brandon acknowdged the significance of Fred Sommers' work on the logic underlying our ordinary reasoning. Yet he went on to claim that Sommers' plus / minus algorithm is unfit for revealing our ordinary 'inclination to accept denying the antecedent or affirming the consequent'. However, a closer inspection of that algorithm shows just how easily antecedent denial (-p+q, -p hence -q) can be confused with modus tollens (-p+q, -q hence -p), and how easily consequent affirmation (-p+q, +q hence +p) can be confused with modus ponens (-p+q, +p hence +q). Sommers and I have presented the entire system in detail in An Invitation to Formal Reasoning, Aldershot: Ashgate, 2000.

George Englebretsen
Philosophy, Bishop's University

§6
Events

March

Relativism and Rational Reflection: 10th Annual Pitt–CMU Graduate Student Philosophy Conference, University of Pittsburgh, 1 March.

Artificial General Intelligence: The First Conference on Artificial General Intelligence, Memphis, Tennessee, 1–3 March.

Science and Pseudoscience: University of Birmingham, UK, 15 March.

Russell: Proof Theory meets Type Theory, Swansea, 15–16 March.

Constraint-Sac: Track on Constraint Solving and Programming, at the 23rd Annual ACM Symposium on Applied Computing, Fortaleza, Brazil 16–20 March.

Causation 1500-2000: King's Manor, University of York, 25–27 March.

UnClog: International Workshop on Interval / Probabilistic Uncertainty and Non-Classical Logics, Ishikawa, Japan, 25–28 March.

AITA: Architectures for Intelligent Theory-Based Agents, Stanford University, 26–28 March.

Naturalism, Normativity, and the Space of Reasons: University College Dublin, 28–29 March.

April

AISB: Artificial Intelligence and Simulation of Behaviour, Aberdeen, 1–4 April.

Subjective Bayesian Methods: Department of Probability and Statistics, University of Sheffield, 2 April.

LSIR: Logic and the Simulation of Interaction and Reasoning, Aberdeen, 3–4 April.

RelMiCS10-AKA5: 10th International Conference on Relational Methods in Computer Science & 5th International Conference on Applications of Kleene Algebra, Frauenwörth, Germany, 7–11 April.

Reduction and the Special Sciences: Tilburg Center for Logic and Philosophy of Science, 10–12 April.

Theoretical Frameworks: Theoretical Frameworks and Empirical Underdetermination Workshop, University of D"usseldorf, 10–12 April.

FLOPS: Ninth International Symposium on Functional and Logic Programming, Ise, Japan, 14–16 April.

Workshop: XVIII Inter-University Workshop on Philosophy and Cognitive Science, Madrid, luis.fernandez@filos.ucm.es, 22–24 April.

Practical Rationality: Intentionality, Normativity and Reflexivity, University of Navarra, 23–25 April.

Non-Classical Logics: From Foundations to Applications, Centro di Ricerca Matematica Ennio de Giorgi, Pisa, Italy, 24–26 April.

SDM: 8th Siam International Conference on Data Mining, Hyatt Regency Hotel, Atlanta, Georgia, USA, 24–26 April.

Transcendental Philosophy: Transcendental Philosophy and Naturalism in Maths and Logic, London, 25 April.

May

SBIES: Seminar on Bayesian Inference in Econometrics and Statistics, University of Chicago Graduate School of Business Gleacher Center, 2–3 May.

Pragmatism and Naturalism: Workshop, Tilburg Center for Logic and Philosophy of Science, 7–9 May.

SIG16: 3rd Biennial Meeting of the EARLI-Special Interest Group 16—Metacognition, Ioannina, Greece, 8–10 May.

CLE, EBL & SLALM: 30th Anniversary of the Centre for Logic, Epistemology and the History of Science (CLE), UNICAMP, 15th Brazilian Logic Conference, and 14th Latin-American Symposium on Mathematical Logic, Paraty, Brazil, 11–17 May.

ArgMAS: Fifth International Workshop on Argumentation in Multi-Agent Systems, Estoril, Portugal, 12–13 May.

Interval Probability: Workshop on Principles and Methods of Statistical Inference with Interval Probability, Durham, 12–16 May.

DL: 21st International Workshop on Description Logics, Dresden, 13–16 May.

FEW: Fifth Annual Formal Epistemology Workshop, Madison, Wisconsin, 14–18 May.

UR: Special Track on Uncertain Reasoning, 21st International Florida Artificial Intelligence Research Society Conference, Coconut Grove, Florida, 15–17 May.

AI Planning and Scheduling: A Special Track at the 21st International FLAIRS Conference, Coconut Grove, Florida, 15–17 May.

RSKT: Rough Sets and Knowledge Technology, Chengdu, 17–19 May.

ManyVal: Applications of Topological Dualities to Measure Theory in Algebraic Many-Valued Logic, Milan, 19–21 May.

NAFIPS: North American Fuzzy Information Processing Society Annual Conference, Rockefeller University, New York, 19–22 May.

ISMIS: The Seventeenth International Symposium on Methodologies for Intelligent Systems, York University, Toronto, Canada, 20–23 May.

WCB: Workshop on Constraint Based Methods for Bioinformatics, Paris, 22 May.

COMMA: Second International Conference on Computational Models of Argument, Toulouse, 28–30 May.

AI: 21st Canadian Conference on Artificial Intelligence, Windsor, Ontario, 28–30 May.

Expressions of Analogy: Faculty of Social and Human Sciences, New University of Lisbon, 29–31 May.

June

ARea: International Workshop on Advancing Reasoning on the Web: Scalability and Commonsense, Tenerife, 1 June.

WCCI: IEEE World Congress on Computational Intelligence, Hong Kong, 1–6 June.

Meta-Analysis: Synthesis and Appraisal of Multiple Sources of Empirical Evidence, Statistical and Applied Mathematical Sciences Institute, North Carolina, 2–13 June.

CSHPS: Canadian Society for History and Philosophy of Science, University of British Columbia, Vancouver, 3–5 June.

CiE: Computability in Europe 2008: Logic and Theory of Algorithms, University of Athens, Athens, 15–20 June.

IIS: Intelligent Information Systems, Zakopane, Poland, 16–18 June.

DM: SIAM Conference on Discrete Mathematics, University of Vermont, Burlington, Vermont, 16–19 June.

Logica: Hejnice, Czech Republic, 16–20 June.

IEA-AIE: 21st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, Wroclaw, Poland, 18–20 June.

HOPOS: Seventh Congress of the International Society for the History of Philosophy of Science, Vancouver, Canada, 18–21 June.

HDM: Multivariate statistical modelling and high dimensional data mining, Kayseri, Turkey, 19–23 June.

EPISTEME: Law and Evidence, Dartmouth College, 20–21 June.

IPMU: Information Processing and Management of Uncertainty in Knowledge-Based Systems, Malaga, Spain, 22–27 June.

Med: 16th Mediterranean Conference on Control and Automation, Ajaccio, Corsica, 25–27 June.

ESPP: European Society for Philosophy and Psychology, Utrecht, 26–28 June.

Philosophy of Probability: Graduate Conference, London School of Economics, 27–28 June.

DGL: Second Workshop in Decisions, Games and Logic, Institute for Logic, Language and Computation, Amsterdam, 30 June – 2 July.

EWRL: European Workshop on Reinforcement Learning, INRIA, Lille, 30 June – 3 July.

July

WoLLIC: 15th Workshop on Logic, Language, Information and Computation, Edinburgh, 1–4 July.

LOFT: 8th Conference on Logic and the Foundations of Game and Decision Theory, 3–5 July.

Logic Colloquium: Bern, Switzerland, 3–8 July.

ICML: International Conference on Machine Learning, Helsinki, 5–9 July.

SMT: 6th International Workshop on Satisfiability Modulo Theories, Princeton, 7–8 July.

Computation and Cognitive Science: King's College, Cambridge, 7–8 July.

Negation and Denial: Philosophy Centre, University of Lisbon, 7–8 July.

CAV: 20th International Conference on Computer Aided Verification, Princeton, 7–14 July.

Induction: Historical and Contemporary Approaches, 5th Ghentian Conference in the Philosophy of Science, Centre for Logic and Philosophy of Science, Ghent, 8–10 July.

Bayesian Modelling: 6th Bayesian Modelling Applications Workshop, Helsinki, 9 July.

UAI: Uncertainty in Artificial Intelligence, Helsinki, 9–12 July.

COLT: Conference on Learning Theory, Helsinki, 9–12 July.

Classical Logic and Computation: Reykjavik, 13 July.

WCP4: Fourth World Congress of Paraconsistency, Melbourne, 1–31 July.

BPR: The 1st International Workshop on Bit-Precise Reasoning, Princeton, 14 July.

ITSL: Information Theory and Statistical Learning, Las Vegas, 14–15 July.

IKE: International Conference on Information and Knowledge Engineering, Las Vegas, 14–17 July.

DMIN: International Conference on Data Mining, Las Vegas, 14–17 July.

NorMAS: 3rd International Workshop on Normative Multiagent Systems, Luxembourg, 15–16 July.

DEON: 9th International Conference on Deontic Logic in Computer Science, Luxembourg, 15–18 July.

NCPW: 11th Neural Computation and Psychology Workshop, Oxford, 16–18 July.

Proof Theory: Workshop on Logic, Foundational Research, and Metamathematics II, WWU Institute for Mathematical Logic, Münster, 18–19 July.

MoChArt: Fifth Workshop on Model Checking and Artificial Intelligence, Patras, Greece, 21–22 July.

WIGSK: Inference methods based on graphical structures of knowledge, Patras, Greece, 21–22 July.

ISBA: 9th World Meeting, International Society for Bayesian Analysis, Hamilton Island, Australia, 21–25 July.

Model Selection: Current Trends and Challenges in Model Selection and Related Areas, University of Vienna, 24–26 July.

ESARM: Workshop on Empirically Successful Automated Reasoning for Mathematics, Birmingham, UK, 26 July – 2 August.

First Formal Epistemology Festival: Conditionals and Ranking Functions, Konstanz, 28–30 July.

August

Conference: Language, Communication and Cognition, University of Brighton, 4–7 August, Brighton, UK.

ESSLLI: European Summer School in Logic, Language and Information, Freie und Hansestadt Hamburg, Germany, 5–15 August.

IJCAR: The 4th International Joint Conference on Automated Reasoning, Sydney, 10–15 August.

ICT: The Sixth International Conference on Thinking, San Servolo, Venice, 21–23 August.

Compstat: International Conference on Computational Statistics, Porto, Portugal, 24–29 August.

LSFA: Third Workshop on Logical and Semantic Frameworks, with Applications, Salvador, Bahia, Brazil, 26 August.

September

IVA: The Eighth International Conference on Intelligent Virtual Agents, Tokyo, 1–3 September.

10th Asian Logic Conference: Kobe University, Japan, 1–6 September.

COMSOC: 2nd International Workshop on Computational Social Choice, Liverpool, 3–5 September.

KES: 12th International Conference on Knowledge-Based and Intelligent Information & Engineering Systems, Zagreb, 3–5 September.

ICANN: 18th International Conference on Artificial Neural Networks, Prague, 3–6 September.

BLC: British Logic Colloquium, Nottingham, 4–6 September.

SMPS: Soft Methods for Probability and Statistics, 4th International Conference, Toulouse, 8–10 September.

AiML: Advances in Modal Logic, LORIA, Nancy, France, 9–12 September.

Causality and Probability in the Sciences
University of Kent, Canterbury UK, 10–12 September

Colloquium Logicum: he biennial meeting of the German Society for Mathematical Logic, Technische Universitaet Darmstadt, 10–12 September.

Logic of Change, Change of Logic: Prague, 10–14 September.

ICAPS: International Conference on Automated Planning and Scheduling, Sydney, 14–18 September.

CSL: 17th Annual Conference of the European Association for Computer Science Logic, Bertinoro, Italy, 15–20 September.

PGM: The fourth European Workshop on Probabilistic Graphical Models, Aalborg, Denmark, 16–19 September.

HAIS: 3rd International Workshop on Hybrid Artificial Intelligence Systems, Burgos, Spain, 24–26 September.

October

SETN: 5th Hellenic Conference on Artificial Intelligence, Syros, Greece, 2–4 October.

Reason, Activism, and Change: University of Windsor, 3–5 October.

MICAI: 7th Mexican International Conference on Artificial Intelligence, Mexico City, 27–31 October.

MDAI: Modeling Decisions for Artificial Intelligence, Barcelona, 30–31 October.

December

ICLP: 24th International Conference on Logic Programming, Udine, Italy, 9–13 December.

§7 Jobs

Artificial Intelligence and Assistive Technology: Research Associate, University of Dundee, deadline 3 March.

IHPST, Paris: Postdoctoral Fellowship, History and philosophy of logic / history and philosophy of science, deadline 1 April.

§8
Courses and Studentships

Courses

MA in Reasoning

An interdisciplinary programme at the University of Kent, Canterbury, UK. Core modules on logical, causal, probabilistic, scientific and mathematical reasoning and further modules from Philosophy, Psychology, Computing, Statistics and Law.

MLSS: 10th Machine Learning Summer School, Kioloa Coastal Campus, Australian National University, 3–14 March.

Algorithmic Decision Theory: MCDA, Data Mining and Rough Sets, Doctoral School, Troina, Italy, 11–16 April.

EASSS: 10th European Agent Systems Summer School, New University of Lisbon, 5–9 May.

Logic School: State University of Campinas, Brazil, 7–9 May.

Logic and Formal Epistemology: Summer school for undergraduates, Department of Philosophy, Carnegie Mellon University, Pittsburg, 9–27 June.

SIPTA: 3rd SIPTA School on Imprecise Probabilities, Montpellier, 2–8 July.

Probabilistic Causality: Central European University, Budapest, 21 July–1 August.

ESSLLI: European Summer School in Logic, Language and Information, Hamburg, 4–15 August.

Mathematics, Algorithms, and Proofs: Summer School, Abdus Salam International Centre for Theoretical Physics, Trieste, 11–29 August.

Causality Study Fortnight

University of Kent, Canterbury UK, 8–19 September

Studentships

Leeds Philosophy: 2-3 postgraduate studentships in philosophy and history and philosophy of science, deadline 1 March.

Kent Computing: 5 PhD Studentships in theoretical computer science, contact computer-science@kent.ac.uk.

Acknowledgements

The Reasoner is a development of the progicnet academic network supported by the Leverhulme Trust.