THE REASONER

Volume 6, Number 1
January 2012

www.thereasoner.org

ISSN 1757-0522

Contents

Features

News

What's Hot in . . .

Events

Courses and Programmes

Jobs and Studentships

Editorial

On the occasion of the recent 50th anniversary of artificial intelligence, I interviewed two prominent figures of AI, Professor Jacques Pitrat and Professor Aaron Sloman, inviting them to make a personal evaluation of where AI stands right now and a personal declaration on its evolution, so as to crystallize a

point in history of science on which future generations might reflect. Our long and delightful discussions revolved loosely around a number of themes, among which some related to the evolution of AI goals, its misses and its prospects are summarized in the interview excerpts presented in this issue.

I want to point out a number of conclusions drawn from comparing pthe scientists opinions. Firstly, both scientists firmly insist on what may be called "the non-linear nature of scientific progress". The developments in AI are attributed less to an objective and expected order of discovery of things (problems, methods, approaches) and more to external influences coming from a number of sources: the structure of the research environment, that induces lack of sufficient research time and fragmentation of problems, the tremendous growth of computational power, the priorities of the funding agencies etc. According to Prof. Pitrat, it's an art to select one's own problem.

Secondly, both scientists share a taste for large-scale projects: Prof. Pitrat declares this to be the only road to AI, while Prof. Sloman calls for coordinated efforts on putting all the pieces together in a working system whose behaviour would compare to that of a little child. They both claim that the most important questions arise only in complex enough situations, when many parts or functions of a system will be simultaneously recruited, be it rule-based systems or animal-mimicking mechanisms or whatever.

Thirdly, despite both scientists having a well-known "symbolic past", that is to say, recognition during the so-called symbolic or good old-fashioned AI days, their ideas are deeply rooted in a broad evolutionary and developmental thinking, hardly visible in the presented excerpts. Both of them believe that the large systems that they preach have to be built incrementally and in an evolutionary way and successive versions or partial designs will present novel and increasingly complex problems that will allow more intricate solutions to be found. Their approach is also largely developmental, both explicitly because they claim so, as well as tacitly through the ideas of self-observation, self-reorganization and internalization, and through their favorite references, to Karmiloff-Smith (Prof. Sloman) or to a special issue on child development (Prof. Pitrat). Is this symptomatic of the pervasive evo/devo thinking of nowadays or is it the distilled outcome of over forty years of research each? It's hard to know.

A final important belief that both scientists share is the need for interdisciplinarity: Prof. Sloman insists on the importance of a broad educational curriculum for prospective AI researchers, while Prof. Pitrat denies that AI is pure computer science business and expects major future breakthroughs to come from other unrelated domains, such as neuroscience. The non-linear development of AI so far gives room to think that new roads will be explored in the future, far away from those taken in the past.

ELPIDA TZAFESTAS
University of Athens

Interview with Jacques Pitrat and Aaron Sloman

Jacques Pitrat is an Emeritus Research Director in University Paris VI, France, and has been working in Artificial Intelligence since 1958. He is well-known worldwide for his work on meta-knowledge and meta-reasoning and author of a number books, the latest one entitled: "Artificial Beings: The conscience of a conscious machine" (Wiley, 2009).

Aaron Sloman is an Honorary Professor of Artificial Intelligence and Cognitive Science in the University of Birmingham, UK. He has been involved with mainstream AI work as the "best way to do philosophy" and is well known for his "Cognition and Affect" lifetime research project.

Both have been very influential and key figures in establishing the artificial intelligence communities in their respective countries.

1. Where do you think AI stands right now?

2. Do you think the original AI goals have changed? Evolved? Abandoned? Where are we now?

3. Which are in your opinion the most important steps and achievements in the history of AI so far? And which are the biggest misses so far?

4. What you think are the most promising directions/ideas for the future?

Jacques Pitrat: Firstly, I think that the original AI goals have been completely abandoned. I think actually AI doesn't stand for Artificial Intelligence, it stands for Advanced Informatics. Advanced Informatics isn't bad, people are right in doing this, but they are wrong in only doing this. So why did things change? That's because AI is extremely hard, and I am convinced no simple algorithm can solve everything, the fundamental concepts are still to discover, despite advances sofar.

I am also convinced that we should experiment with large systems, which is a very ambitious task, because by putting such systems to work, things are discovered or clarified and give us new ideas. I believe that's the only approach to *really* artificially intelligent systems. Unfortunately, if you wish to build a large system but not wish to devote half of your time to it, it's zero. For instance, when I return from holiday I need quite some time to re-adapt myself to the 10000 rules of my system. So, if people can only work on it for a couple of hours per week, then it's zero. Nowadays, most researchers cannot afford putting half of their time to such a project, except for PhD students, who are not experienced enough and are only there for a short time—enough to learn to do things, but not enough to do AI. For me, these worldwide developments are disastrous. Most researchers, as soon as they defend their PhD, start giving lectures

and courses, doing administrative tasks, raising funding for their research groups and so on. Result: it is materially impossible for them to do what they should do.

Another serious problem concerns publications: researchers are asked to publish at a mad rate, maybe a dozen papers per year. Consequently, people lose their time writing papers, reviewing papers, which is redundant, etc. If they decided to build a really big system, they would only publish one paper every other year and they would be soon thrown out of academia and be unable to continue doing research, unless their dad is called Rothschild ... Thus the structure of the research environment practically forbids doing research.

Not to mention that, when the progress of a big system like this can only be found in what it does internally and not in its results, then most people not working in exactly the same domain would be incapable of understanding.

This is mainly why the goals have changed. In the beginning, AI was relatively easy to do, because everything had to be done from scratch. People would create small-scale but extremely interesting systems: the original systems by Newell and Simon were small and one could easily work on them in parallel with other things. But nowadays, the easy things are gone and what's left to do, it's the very difficult ones. That's why I work for 25 years on my AAIS with the goal to bootstrap AI, because I am convinced that artificial intelligence can be achieved, but we are not intelligent enough to achieve it! So, to achieve AI, we have to seek the help of AI itself. Sadly, this takes a long time and it's an art to choose what's worth doing first. So I am building an artificial AI scientist knowing that I will never reach my goal ...

Simon has extensively studied problem solving and has identified two kinds of techniques. One is the means-ends analysis, where there's a goal which differs from the current situation but which can be progressively reached by subdividing differences. The interesting thing with this approach is that the user can see that he's near the goal and often perceive his distance from it. On the contrary, in problem solving by representation switching, like in the case of the mutilated checkerboard (C. Kaplan and H. Simon 1990: In search of insight, *Cognitive Psychology*, 22: 374–419), one is unable to see how close to the solution he is and he may even realize the problem was solved only a posteriori. So, I think AI is a problem of this type, one that should be solved by representation switching, which is a kind of meta reasoning.

The original goals have also changed because people are abusing the term AI: knowing that AI is very hard, they are doing something else, which is not bad but irrelevant. I compare AI to a church whose priests have lost their faith: the AI people don't really believe in it anymore. These very intelligent people demonstrate lots of theorems (AI journal is full of them), but no theorem has been demonstrated by the machine itself. It's like a student's parents who are doing their child's homework in its place, which results in excellent marks, but it doesn't solve the problem. Researchers don't do service to AI by struggling for programs of outstanding performance, like contemporary chess programs, because there's no AI in there. These systems use methods developed in the 60's, but the machines are now much more powerful than in the past and the amount of built-in knowledge, such as large opening libraries, is tremendous. This is why in human-machine matches, it is often disallowed to the program to use these libraries, otherwise they couldn't compete!

The most important developments in AI remain those of the early days: theoretical logics, GPS, etc., From 1965 on, most papers have been either too theoretical or based on combinatorics, which may be objectively excellent, but highly insufficient. I didn't fall in love with too many papers, except for Lenat's papers, where many ingenious ideas may be found, and Laurière's work.

Still, there are important fundamental problems that have to be solved first, and specifically I emphasize the importance of working at the meta level, because this is what distinguishes essentially man from animal: the possibility to observe oneself, to reflect on what one is doing, to know what one has done etc. This is why I am very pessimistic about the future of AI, because only the people who have nothing to lose, like me, can work on such large-scale systems.

Aaron Sloman: I'll start with the negative point. I think a huge amount has been achieved in the last sixty years, but there are a number of things that are going wrong. One is that the theory now is very deeply fragmented, for example in the 80's, there was a major AI conference like IJCAI, and that would be the place where most of the researchers would want to go and present their latest results, whereas now they will mostly want to go to the specialised conferences. As a result of this fragmentation there are lots of subgroups of researchers that all agree on a set of problems, often defined by benchmark tests, and developments are achieved which show results that look like great improvements on some kind of test, but they have nothing to do with putting the pieces together in a working system, like a robot that can do things that a young child could do. So, regarding engineering there are many impressive applications, but regarding the understanding of the science of how to put together an intelligent system, the theory is very fragmented and that's seriously holding up progress.

One of the consequences is that the education of people in AI tends to be very inadequate, so people will learn a lot of things in a particular subfield and not learn much about anything else. Another thing that has happened in AI and also to some extent in some other disciplines is that there have been fashions and factions, that's to say that people often—just because they are competing for funding—like to be able to claim that what other people are doing is no good and they've got the new answer to the old problems that's the only good answer. So people won't even look at what others are doing because they have been convinced that the others' approaches are worthless.

But I think that there is some hope, in particular the European initiatives in the last few years of trying to set up funding for projects that put the pieces together. The cognitive systems and cognitive robotics projects can help to overcome that, although it may take a long time for that to feed into educational systems, so that we get young people come into universities and get a broad education in the whole field. I think the problems are very difficult and I don't know whether we've got enough people who are of high enough quality and broad enough level of education to solve these problems, but at least something is being attempted.

One way I see of trying to go forward reflects my own interests in AI as a way of
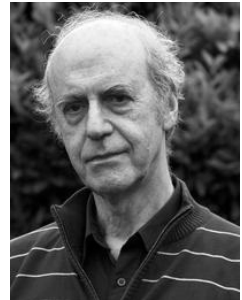
understanding natural intelligence. I want to understand what biological evolution produced and I think it useful to use AI as a kind of probe. So, one question that can lead in many different directions is, what were the advances or the developments in information processing that biological evolution produced and why were they needed and in what ways were they inadequate? It seems to me that between the earliest microorganisms and things like humans and elephants, there were very many transitions in designs of organisms which were provoked by transitions in the problems as the environments changed and as organisms became more complex. Complex behavior requires the ability to perceive and to take in information about the processes that are going on or could go on in the environment. And what can be done with that information about processes would depend partly on what kind of body the organism has, as well as on what's in the environment.

So we need to look in great detail at what animals do and what the problems are and what the advantages are of different kinds of designs and if the designs are good, the kinds of constraints they have etc. We have nothing like a deep, systematic overview of those problems and the different kinds of possible solutions because most of work in AI has been very piecemeal, and there is no attempt to produce a systematic overview of all this. I suggest to begin to develop such an overview by trying to look for the variety of the forms that evolution produced, analyzing it as engineers and then trying to model and mimick them.

I also think there's a lot that's going on when the animals are not behaving: that ability, not just to perceive and respond immediately, to behave promptly in a specific situation, but to think about what you might do if you wanted to assemble for instance a shelter from the materials that are available somewhere, that requires an ability to take in a deep and possibly broad collection of possibilities for change and for rearrangement of relationships and structures. Humans are not the only animals that can do that, others can do it in varying degrees and very little is understood about that process. One of the reasons is that it's actually very hard to study. You see the behavior but you don't see the thinking and the reasoning behind, and therefore you don't know that you have to replicate that as well.

Annette Karmiloff-Smith introduces the notion that humans and some other animals, but only very few robots, can acquire what she calls "behavioral mastery" in a domain (A. Karmiloff-Smith 1995: *Beyond modularity: A developmental perspective on cognitive science*, MIT Press). Once you develop some behavioral mastery, then you reorganize what you've learnt to do in a way that gives new capabilities, human language learning being the most obvious example. Our brains, after having mastered some collection of capabilities, can reorganize their knowledge into a new form that goes beyond what has currently been learnt. And that ability to transcend examples is not present in most of the currently popular forms of learning that people in AI are investigating, for instance neural nets. And I think that that's something about human and animal minds that biological evolution produced, that we need to understand and do something about.

Frege's Puzzle from a Model-Based Point of View

Frege's puzzle about propositional attitude reports can be presented in terms of Superman comics. See for example, Thomas McKay, Michael Nelson (2010: *Propositional Attitude Reports*, *Stanford Encyclopedia of Philosophy*):

> At the beginning, Lois Lane does not realize that Clark Kent and Superman are the same person, and she concludes from her observations that *Superman is strong*, but *Clark Kent is not strong*. Thus, it is true that *Lois believes that Superman is strong*, and that *Lois does not believe that Clark Kent is strong*. But since Clark and Superman are the same person, *Clark Kent = Superman* is true as well.

Now, is the rule $F(x) \& x = y \rightarrow F(y)$ valid as a general logical principle? If it is, then, by applying it to true sentences, *Lois does not believe that Clark Kent is strong*, and *Clark Kent = Superman*, we should obtain a true sentence: *Lois does not believe that Superman is strong*. However, the sentence *Lois believes that Superman is strong* is true as well, which is a contradiction. Thus, as a general logical principle, $F(x) \& x = y \rightarrow F(y)$ seems to be wrong.

This kind of disorder has caused more than a century of controversy. Let's try one more approach to solving the puzzle.

The model-based approach used below can be traced to Marvin Minsky (1965: *Matter, Mind and Models, Proceedings of IFIP Congress 65*, 1: 45–49). In my (2009: *Towards Model-Based Model of Cognition*, *The Reasoner* 3(6): 5–6) I presented this "robotic ontology" as follows:

> In my head, I have a *world model* (an incomplete one, incoherent, inconsistent, containing all my knowledge, beliefs, etc.). In this model, other persons are believed to have their own world models. Thus, my world model may contain "models of models", for example, a simplified model of your world model.

But, despite the possible inconsistency of my world model, I don't wish to admit contradictions like Frege's puzzle into it.

How does Frege's puzzle look from this point of view? At the beginning, Lois' world model includes the axiom *Clark Kent ≠ Superman*. Thus, in Lois' world model, her conclusions that *Superman is strong*, but *Clark Kent is not strong* do not contradict each other. But, as a reader of the Superman comics, I know from the very beginning that Clark and Superman are the same person. Hence, in *my* world model, *Clark Kent is strong*, but Lois believes the opposite. At the end of story, Lois is forced to *change* her world model axioms, and Clark becomes strong in her model, too. No puzzle here!

What could have caused the "puzzlification" of the situation?

The statements *Superman is strong* and *Clark Kent is not strong* belong to Lois' initial world model. In this model, *Superman ≠ Clark Kent*. Of course, Lois will not try replacing Superman with Clark Kent in these statements.

The statements *Lois believes that "Superman is strong"*, *Lois does not believe that "Clark Kent is strong"*, *Lois believes that "Superman ≠ Clark Kent"*, and *Superman =*

Clark Kent belong to the world model of the reader, but the parts of the statements in quotes refer to Lois' initial world model. Of course, the reader will not try replacing Superman with Clark Kent in the statement parts referring to Lois' world model.

Thus, one can run into puzzles only by *confusion* of different world models.

A formal model of the situation can be presented as follows. Let's imagine that all sentences we are interested in belong to some *uninterpreted* formal language plus some suitable system of logic. The world model of some person X is represented by a set of axioms, which allows one to derive all sentences that X believes in. Let's denote this axiom set by *WorldModel[X]*. If our logic includes the principle *Ex contradictione sequitur quodlibet*, then we must assume that WorldModel[X] doesn't contain *known* contradictions. The situation of Frege's puzzle is represented as follows:

$$\vdash P[Y1] \,\&\, Y1=Y2 \rightarrow P[Y2];$$

$$WorldModel[X] \vdash P[Y1] \,\&\, Y1 \neq Y2 \,\&\, \neg P[Y2].$$

Of course, no puzzle here!

The triviality of this solution is due to the purely *syntactical* character of the approach. Namely, let's regard world models *not* as "models of the world" with the world itself as their unique "reference". Let's consider world models simply as the way that people are thinking and talking about the world. When trying to understand their utterances, let's analyse what people *are thinking to be true*, and not what is true "in fact".

People are comparing and coordinating their world models. But no "independent jury" can be established for comparing of two models M1 and M2, or for comparing of some model M3 with its target system S3 in the world. Speaking strictly, I only can compare things that are contained in my world model: compare *my models* of the models M1 and M2, or compare *my model* of M3 with *my model* of S3.

As demonstrated above, under this approach, at least some of the puzzles disappear …

A similar formulation is attributed to Niels Bohr: "There is no quantum world. There is only an abstract quantum physical description. It is wrong to think the task of physics is to find out how nature *is*. Physics concerns what we can *say* about nature."— quoted after Aage Petersen (1963: The Philosophy of Niels Bohr. *Bulletin of the Atomic Scientists*, XIX(7): 8–14).

Every utterance comes from the world model of the speaker. More generally, every sentence comes from some kind of world model. It may be the world model of a (real or imagined) person, the world model represented in a novel, movie, scientific book, virtual reality, etc. In principle, even smaller informational units (stories, poems, newspaper articles, jokes, mathematical proofs, video clips, dreams, hallucinations, etc.) may introduce their own "partial world models" as small additions to "bigger" world models (regarded as background knowledge). Sometimes, sentences contain references to other world models. Trying to understand such sentences, we should identify, and keep separated, the world models involved.

Karlis Podnieks
Computer Science, University of Latvia

Games and the Reason-Knowledge Principle

John Hawthorne and Jason Stanley (2008: "Knowledge and Action", *Journal of Philosophy* 105(10), 571–90) defend what they call the "Reason-Knowledge Principle".

> Where one's choice is *p*-dependent, it is appropriate to treat the proposition that *p* as a reason for acting iff you know that *p*. (578)

There have been many attempts in the literature to show that this leads to implausible *actions*. As Jonathan Ichikawa (2012, "Knowledge Norms and Acting Well") shows, most of these attempts rest on further, and arguably false, assumptions about the connection between reasons and action. Relatedly, most of those responses concern the role of knowledge and reasons in decision-making. I'll argue that we can formulate a sharper problem for the principle if we focus on game-playing, and say exactly which extra assumptions we are making.

The Reason-Knowledge Principle should have the following implications, at least for cases where S's aim is to produce the best outcome.

(1) If S knows that φ and ψ will produce the same outcome, and S must choose φ or ψ, then it is rationally permissible for S to choose ψ.
(2) If S knows that φ and ψ will produce the same outcome if p, and φ will produce a better outcome if $\neg p$, then it is rationally permissible for S to choose ψ iff she knows p.

The point of (1) is that S can use her knowledge that φ and ψ will produce the same outcome to justify making an arbitrary choice between φ and ψ. And the point of (2) is that the Reason-Knowledge Principle suggests only knowledge that p could justify ignoring the fact that ψ does worse than φ if $\neg p$.

Define a **symmetric** game as having these features:

- The game is purely co-operative; each player gets the same payoffs;
- Each player knows nothing about the other save that it is common knowledge the players are rational, and hence know what each other's rational requirements are;
- Each player has the same moves available; and,
- The payoffs are a function of just which moves are made, not of who makes them.

Assume A and B are playing a symmetric game, and it is common knowledge which symmetric game they are playing. Then the following premise seems hard to dispute:

(3) It is rationally required for A to play φ iff A knows B will play φ.

What makes (3) so compelling is that we can derive it from (4), (5) and (6).

(4) A knows that B will play φ iff A knows that any rational player will play φ.
(5) If A knows any rational player will play φ, then A is rationally required to play φ.
(6) If A is rationally required to play φ, then A knows that any rational player will play φ.

We get (4) from the fact that *A* knows nothing about *B* save that she is rational. We get (5) by the factivity of knowledge. And we get (6) by the requirement that the players are rational, and hence know what rationality requires of each player. And these three together entail (3). So (3) is true, and (1) and (2) are entailed by the Reason-Knowledge Principle. Unfortunately, (1), (2) and (3) are inconsistent, as we'll now show.

Informally, in this game *A* and *B* must each play either a green or red card. I will capitalise *A*'s moves, i.e., *A* can play GREEN or RED, and italicise *B*'s moves, i.e., *B* can play *green* or *red*. If two green cards, or one green card and one red card are played, each player gets $1. If two red cards are played, each gets nothing. Each cares just about their own wealth, so getting $1 is worth 1 util. All of this is common knowledge. More formally, here is the game table, with *A* on the row and *B* on the column.

	green	*red*
GREEN	1, 1	1, 1
RED	1, 1	0, 0

Assume *A* knows *B* will play *green*. By (3), it is rationally required that *A* plays GREEN. But *A* can use this knowledge of *B* to deduce that GREEN and RED have the same payoff. So by (1), it is rationally permissible to play RED. Contradiction.

Now assume *A* does not know *B* will play *green*. By (3), it is not a rational requirement that *A* plays GREEN. But *A* knows that GREEN does better than RED unless *B* plays *green*. And since she doesn't know *B* plays *green*, by (2), she's required to play GREEN. Contradiction.

So either assuming that *A* does or does not know that it is rationally required for *B* to play *green* leads to a contradiction given (1), (2) and (3). So these three premises are inconsistent. Since (3) is true, that means (1) or (2) is false. And since the Reason-Knowledge principle entails those two premises, one of which is false, the Reason-Knowledge Principle is false.

I'm not entirely sure which of (1) and (2) is false; both of them do feel plausible. I suspect the problem is (1). Assume *A* deduces from premises she believes that rational players will play a green card. Perhaps she agrees with Robert Stalnaker (1998, "Belief revision in games: forward and backward induction", *Mathematical Social Sciences*, 36(1), 31–56) that rationality requires avoiding weakly dominated options. Then she knows it doesn't matter to her outcome whether she plays GREEN or RED; she will get $1 either way. But if she plays RED, she is incoherent; she is doing something she thinks no rational player does. And perhaps this incoherence is a bad thing in itself. Niko Kolodny (2005, "Why be Rational?", *Mind*, 114(455), 509–563) argues that incoherence is not bad in itself; Jacob Ross (2012, "Rationality, Normativity, and Commitment", *Oxford Studies in Metaethics*, 7, forthcoming) argues that it is. The suggestion that (1) is the false premise favours Ross's view over Kolodny's. But this conclusion is very speculative; the main thing I wanted to note was the problem this game raises for the Reason-Knowledge Principle.

BRIAN WEATHERSON
University of Michigan &
Arché, University of St Andrews

History and Philosophy of Computing, 7–10 November

The International Conference on *History and Philosophy of Computing* (Ghent University, Belgium), collected for the first time in a single venue philosophers, logicians, historians of computing, mathematicians, computer scientists. Aiming at raising historical awareness of the evolution of computing and soliciting a philosophical insight into its fundamental problems, topics included: history of computation; philosophical, foundational and practical issues of computability in logic, mathematics, computer science and other sciences. The programme included 29 contributed papers (selected out of 52 submitted) and 7 invited speakers:

> William Aspray, (University of Texas, US): "Three Topics in the History of Computing". An overview on the development of the concepts of computability, information business, information society and the use of history in their study.

> Martin Davis (Courant Institute, NYU): "Universality is Ubiquitous". A discussion of the foundational work on computability underpinning the development of all-purpose computers.

> Fairouz Kamareddine (Heriot-Watt University, Edinburgh): "From the foundation of mathematics to the birth of computation". A reconstruction of basic ideas of computer science in the light of principles of mathematics and logic, in particular the theory of types.

> Sybille Krämer (Freie Universität Berlin, Germany): "Mathematicizing power, formalization, and the diagrammatical mind or: What does 'computation' mean?". A philosophical understanding of the notion of computation as characterized by symbolism, visualization and diagrammatics.

> Giovanni Sambin (Università di Padova, Italy): "Computability without Turing Machines". The development of intuitionistic and predicative topology conceived as including an abstract theory of computation without Turing machines.

> Raymond Turner (University of Essex, UK): "Towards a Philosophy of Computing Science". An exploration of issues related to the semantics of programming languages and specifications.

> Stephen Wolfram (Wolfram Research): "Making the World Computable". An historical and theoretical overview of mechanization of mathematical knowledge in the Mathematica and Wolfram Alpha projects.

The programme of contributed talks covered all relevant aspects of computing sciences. Foundational issues were tackled with an overview of computable (non-standard) analysis in Turing and today (Gherardi, Sanders) and with a refinement of computability for models of complexity (Dean). The historical approach was largely influential, with

recounters of the development of machines and software (D'Udekem-Gevers, Henriksson, Mounier-Kuhn), including those that never made to the market (Shilov & Kitov, Bondecka-Krzykowska) and the related retrospective conceptual analysis of algorithms from the current viewpoint (Bullynck, Durnova, Numerico). The relation between computing and the sciences touched theory choice and simulation (De Langhe, Brand); AI and its epistemology (Nickel, Bach); information and complexity theories (D'Alfonso, Camardi); bio-computing and the medical sciences (Moore & Kirby, Parolini). A significant number of contributions were dedicated to computer science: its evolution and philosophical significance in modern infrastructures (Geske, Gobbo & Benini); the role of computer experiments in the philosophy of science (Schiaffonati & Verdicchio); the meaning and the semantics of programming (Mascella, Hernandez-Quiroz). Abstracts of the talks can be found at www.computing-conference.ugent.be.

HAPOC has revealed a grey area of important interactions among the different fields connected by computing and we believe the philosophical and historical approach will prove methodologically crucial. A next smaller event will be http://www.computing-conference.ugent.be/hapop12.

LIESBETH DE MOL
GIUSEPPE PRIMIERO
Centre for Logic and Philosophy of Science,
Universiteit Gent

Workshop on Semantics, Pragmatics and Rhetoric, 9–11 November

The SPR-11 was organized by the Institute for Logic, Cognition, Language and Information (ILCLI) of the University of the Basque Country and it was held on November 9–11 in the city of Donostia. The workshop included philosophers, linguists and psychologists, with a total of 9 plenary lectures and 42 contributed papers, organized in parallel sessions. The schedule of the workshop was tight but the high quality of both the plenary and the contributed papers made of it an intense and stimulating experience.

The topics discussed were diverse, but at least a few can be singled out as "persistent" throughout the workshop. The first was reference. Starting from the opening lecture, given by Howard Wettstein (Riverside) on *The myth of reference determination*; continuing with other plenary talks, like those of Genoveva Martí (Barcelona and ICREA): *Reference without having in mind* and Eros Corazza (ILCLI and Carleton): *Revisiting Mill(ianism) from a cognitive perspective*; and finishing with the closing lecture, by John Perry (Stanford and Riverside): *The strange case of Ivan and Donostia*. The second persistent topic was De-se attitudes. This topic was present both in the plenary lectures, with Ofra Magidor (Oxford): *The myth of the de se*, and in the contributed papers, with, to mention but two, Yan Huang's (Auckland) presentation: *Marking of de se attribution and neo-Gricean pragmatics: Logophoric expressions in west African languages and long-distance reflexives in east, south, and southeast Asian languages* and Marie Guillot's (Jean Nicod): *Semantic relativism and the epistemology of de se thoughts*.

Finally, a third prominent topic was, roughly, expressions-type meanings and con-

text. Among the plenary speakers, Jeff Pelletier (Alberta) presented a paper entitled *The syntactic, semantic, and pragmatic properties of the count/mass distinction in English and what this says about syntactic/semantic/pragmatic distinctions generally* and Robyn Carston (UCL and CSMN) talked about *Lexical meaning and concept communicated*. Other topics included incomplete predicates and the desirability of relativizing them, with Maite Ezcurdia's lecture (UNAM): *Incomplete predicates*; and the role of salient meaning, as opposed to literal one, in communication, with Istvan Kecskes' (SUNY): *Why do we say what we say the way we say it: Salience in language production*. Worth noticing in this respect were also two contributions on the analysis of Basque particles, made by Joana Garmendia and by Larraitz Zubeldia.

One of the factors that deserve to be pointed out here was the high quality of the contributed papers, too many and too varied to be mentioned in this short note; and of the discussions during the whole workshop. Perhaps, one of the most debated topics during the workshop were some of the ideas regarding the role of pragmatic analysis (mostly of the speaker's intentions) for reference and communication (particularly interesting were the debates concerning some ideas contained in Korta & Perry's new book *Critical Pragmatics*). In this line were the papers presented by Michael O'Rourke, Carlo Penco, Eduardo García-Ramírez, Robin Jeshion and Wayne Davis, to name just a few of the authors of the contributed papers.

MARÍA PONTE
ILCLI

Explanation, Causality, and Unification, 11–12 November

The workshop was held at the Heinrich–Heine–University Düsseldorf, Germany. It took place as part of the DFG–funded research unit "Causation | Laws | Dispositions | Explanation" (FOR 1063). Its focus were the three concepts mentioned in the workshop's title and the diverse connections among them, as well as related issues within Philosophy of Science. There were 14 slots with 40 minutes talk and 20 minutes discussion each.

The workshop started with Michael Baumgartner (Konstanz) demonstrating that there is an empirically equivalent deterministic common cause structure for every causal chain, but not vice versa, and that there is thus no formal reason to prefer deterministic causal chains over common cause structures. Next was Andreas Hüttemann (Cologne) who argued for a dispositional account of causality. Matti Sintonen (Helsinki) presented a paper in which he showed that the term 'cause' in causal explanations may have several—and very different from each other—meanings. The Friday afternoon session started with Gerhard Schurz (Düsseldorf) presenting a general theory of causality based on causal graph theory that can explain diverse probabilistic properties of certain systems and does have, as a whole theory, empirical consequences. The next presentation was given by Alexander Gebharter (Düsseldorf). He applied the theory presented by Gerhard Schurz and sketched some ideas of how it can be used to capture the concept of a mechanism. Theo Kuipers (Groningen) gave reasons for causal laws to be nomic laws that satisfy an updated set of logico-empirical conditions of adequacy for causal lawhood. The last talk on Friday was given by Kevin Kelly (Pittsburgh). He presented a

topological theory of empirical simplicity and an argument according to which, in light of said theory, Occam's razor is as close to deduction as an inductive truth-finder can possibly get.

The Saturday morning session was opened by Erik Weber's (Ghent) paper on how a mechanistic explanation of laws can be unifying by distinguishing between analogical and theoretical unification. Next was Victor Gijsbers (Leiden) who argued for the controversial thesis that there is understanding by classification but without explanation. Margaret Morrison (Toronto) presented a paper in which she argued that micro–phenomena can be unified by means of the Renormalization Group, a mathematical framework to investigate changes of physical systems at different scales. The Saturday afternoon session was opened by Henk de Regt's (Amsterdam) paper on why Wesley Salmon's complementary thesis is deemed to fail. Lorenzo Casini (Kent) argued for an inferentialist approach for interpreting causal claims in computational economics. His talk was followed by Jon Williamson (Kent) giving reasons for why phenomena should be explained by pointing to their underlying mechanisms. In the light of this hypothesis, he evaluated which one of several accounts of causality is closest to explanation in the sciences. Stathis Psillos (Athens) gave the last talk of the workshop and closed the Saturday afternoon session with a presentation in which he developed a lightweight metaphysics for the Regularity View of Causation.

The workshop led to new insights in issues surrounding its main topics as well as to a rich exchange among and between speakers and guests and thus, also to many ideas for new projects. A publication of the proceedings is planned for 2012.

Gerhard Schurz
Alexander Gebharter
Philosophy, Heinrich-Heine-University Düsseldorf

Workshop in Social Epistemology, 8–9 December

Jointly sponsored by the Einar Hansen Research Funds and the Lund Philosophy Department, some 25 speakers and participants witnessed a total of nine presentations on formal and informal work relating to the four key-phenomena: belief polarization, information cascading, pluralistic ignorance, and echo chambers.

Drawing on agent-based modeling and assuming, amongst others, the states 'belief', 'disbelief', 'withhold belief', and confidence values, Bert Baumgärtner (Davis/Amsterdam) analyzed echo chambers as Nash equilibria where an agent's confidence value fails to decrease upon meeting with the beliefs of other agents.

Chiara Lisciandra, in joint work with Ryan Muldoon, Stephan Hartmann, and Soroush Rafiee Rad (Tilburg), analyzed pluralistic ignorance as a mismatch between agents' intrinsic properties to act and their publically displayed behavior, providing a Bayesian account where "trendsetters" and "conformists" assign different priors to privately and publically displayed beliefs.

Hans van Ditmarsch (Sevilla) presented a multi-agent doxastic logic which sees standard axioms for truthful public announcement "mirrored" by dual axioms for lying and truth-telling, captures the case of bluffing, and—*modulo* intentional aspects—serves

to trace, in particular, the informational consequences of lying.

Drawing on several well-documented cases satisfying nearly the entire range of key-phenomena Jonathan Robson (Nottingham) suggested that—pace writers such as Kant—autonomous aesthetic judgment may count as the exception among humans.

Merel Lefevere, in joint work with Eric Schliesser (Gent), argued that active pluralism in the sciences—i.e., not just minding, but seeking alternative perspectives and problem-solving approaches—may be one of the ways in which a more responsible scientific community can be achieved, especially in cases where policy makers rely on scientists' advice.

Using plausibility models and through variation of information that previous agents' actions convey, Rasmus K. Rendsvig (Roskilde) showed that particular distributions of private signals in a "narrow-minded" (vs. an "open minded") population do (not) suffice to trigger informational cascades.

Rogier de Lange (Helsinki/Tilburg) offered a model which serves to explain the occurrence of the four key-phenomena, particularly the divergences between intrinsic qualities and community belief, through reference to an optimal interplay between competition and coordination.

Ruhalah Ramzani (Tehran) outlined a view according to which belief polarization arises from a conflict between two epistemic norms—the truth and the consistency norm—, holding that polarization is not irrational, and in turn avoidable, given these two are "balanced."

In joint and ongoing work with Wlodek Rabinowicz (Lund), Stephan Hartmann (Tilburg) revealed preliminary results of computationally studying measures for preference aggregation in a value context, then compared their differential reliability in approximating the truth.

Principal investigators are Erik J. Olsson (Lund) and Vincent F. Hendricks (Copenhagen). Presentations are available here. 2012 workshops remain to be announced.

<div align="right">

Frank Zenker
Lund University, Sweden

</div>

Calls for Papers

Structure of Scientific Revolutions: 50 Years On: special issue of *Topoi*, deadline 15 January.

Imprecision in Statistical Data Analysis: special issue of *Computational Statistics & Data Analysis*, deadline 30 January.

Input & Output Analysis for Simulation: special issue of the *Journal of Simulation*, deadline 1 March.

Formal and Intentional Semantics: special issue of *The Monist*, deadline 30 April.

The Mind-Body Problem in Cognitive Neuroscience: special issue of *Philosophia Scientiæ*, deadline 1 May.

Inforgs and the Infosphere: Themes from Luciano Floridi's Philosophy of Artificial Intelligence: special issue of *The Journal of Experimental & Theoretical Artificial Intelligence*, 1 July.

: special issue of *Teorema*, deadline 15 September.

WHAT'S HOT IN . . .

. . . Uncertain Reasoning

Every now and then, a theorem gets its biography. Sharon Bertsch McGrayne (2011: "The Theory That Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted Down Russian Submarines, and Emerged Triumphant from Two Centuries of Controversy", Yale University Press) is by far one of the most engaging I've ever come across. A one hour presentation of "The Theory That Would Not Die" given by the author at Carnegie Mellon University in October 2011 is available on CMU's Youtube channel.

The book is packed with interesting stories, yet I have no hesitation in choosing my favourite one. It appears in Chapter 3 and describes how Bayes' theorem had a place in the infamous *Affaire Dreyfus*. McGrayne doesn't give the references, but the question has indeed attracted some technical interest, e.g., D. Kaye (2007: "Revisiting Dreyfus : A More Complete Account of a Trial by Mathematics". Minnesota Law Review, 167(91), 825–835) and L. Rollet (2010: "Des mathématiciens dans l'affaire Dreyfus? Autoforgerie, bertillonnage et calcul des probabilités", Images des Mathématiques, CNRS).

In 1894, Alfred Dreyfus, a high official of the French artillery, is accused of having attempted to pass military secrets to the Germans. An Alsatian of Jewish background, Dreyfus is convicted for treason and sentenced to life imprisonment at Devil's island. To many, Dreyfus's conviction is a blatant case of antisemitism and in little time the public opinion splits between *dreyfusards* and an *antidreyfusards*. On 13 January 1898, the progressive paper *L'Aurore* publishes an open letter by novelist Émile Zola which will go down in history as the act of birth of the "intellectuals'" involvement in public life. The resonance of Zola's *J'accuse* is huge, and it certainly contributes to igniting the complicated chain of events which led to Dreyfus's complete public rehabilitation in 1906.

Intellectuals' first-person involvement in the Dreyfus affair is quite well known, and the enormous popularity of Zola probably led subsequent commentators to identifying "intellectuals" with humanists. Yet McGrayne reminds us that we shouldn't underestimate the role played by mathematics, and probability in particular, in the 1899 Rennes court-martial, the turning point of the Dreyfus trial. On that occasion, Alphonse Bertillon, father of anthropometry and then in charge of the Bureau of Identification in the Paris Police Department, concocted a pseudo-statistical analysis of the *bordereau*— the memoir containing classified information which Dreyfus allegedly intended to pass on to the Germans—to argue that it was indeed Dreyfus's own writing. Yet Bertillon's argument was proved to contain "a colossal mistake" by Henri Poincaré, perhaps the greatest French mathematician of the time. Poincaré showed the court-martial that in his derivation of the "high probability of coincidence" Bertillon had neglected the likelihood of those very coincidences. After illustrating the mathematical details of Bertillon's mistake, Poincaré commented as follows:

> Je ne sais si l'accusé sera condamné, mais s'il l'est, ce sera sur d'autres preuves. Il est impossible qu'une pareille argumentation fasse quelque impression sur des hommes sans parti pris et qui ont reçu une éducation scientifique solide.

The "solid scientific education" described by Poincaré was certainly available to the Rennes court-martial, which was made of military officials who had studied at the École Polytechnique under Joseph Bertrand. His textbook contained a detailed discussion of Bayes' theorem.

This is perhaps the most striking aspect of the story: what judges agreed to be part of a solid scientific education in 1899, does no longer appear to be so. The latest in a long list of cases in which Bayes' theorem has not been considered admissible in court was reported by The Guardian on 2 October 2011.

Hykel Hosni
Scuola Normale Superiore, Pisa

Events

January

ISAIM: 12th International Symposium on Artificial Intelligence and Mathematics, Fort Lauderdale, Florida, 9–11 January.

University of Miami Graduate Student Conference in Epistemology: Miami, FL, 12–14 January.

MAMLS: Mid-Atlantic Mathematical Logic Seminar, Florida, 13–15 January.

Perspectivalism Workshop: Ghent, 19–20 January.

MathLog: 5th Annual Cambridge Graduate Conference on the Philosophy of Logic and Mathematics, Cambridge, 21–22 January.

GPUs: Computational Statistics, University of Warwick, 25 January.

Vagueness in Language, Reasoning and Cognition: Amsterdam, 27–28 January.

Workshop on Bayesian Approaches to Handling Missing Data: Imperial College London, 30 January.

February

Colombian Conference on Logic, Epistemology, and Philosophy of Science: Bogota, Colombia, 8–10 February.

Conference on Computer Science & Computational Mathematics: Melaka, Malaysia, 9–10 February.

Perspectives on Structuralism: Center for Advanced Studies (CAS) and Munich Center for Mathematical Philosophy (MCMP), LMU Munich, Germany, 16–18 February.

ICIIN: International Conference on Intelligent Information and Networks, Hong Kong, 17–18 February.

ICICA: International Conference on Information and Computer Applications, Hong Kong, 17–18 February.

ICCMS: 4th International Conference on Computer Modeling and Simulation, Hong Kong, 17–18 February.

ICDC: International Conference on Digital Convergence, India, 18–19 February.

Theoretical Computer Science: Auckland, New Zealand, 21–24 February.

The Epistemology of Modality Workshop: Cologne, 23–24 February.

ICICN: International Conference on Information and Computer Networks, Singapore, 26 February.

March

FoIKS: 7th International Symposium on Foundations of Information and Knowledge Systems, Kiel, Germany, 5–9 March.

LATA: 6th International Conference on Language and Automata Theory and Applications, La Coruña, Spain, 5–9 March.

Dispositions, Causes, Modality Workshop: Cologne, 7–9 March.

Graduate Conference in Philosophy of Science: Erasmus University Rotterdam, 8–9 March.

Nothing but the Truth: Vienna Forum for Analytic Philosophy, University of Vienna, 9–11 March.

ICMLC: 4th International Conference on Machine Learning and Computing, Hong Kong, 10–12 March.

LPAR: 18th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Merida, Venezuela, 11–15 March.

Axiomatic vs Semantic Truth: Munich, 14–16 March.

&HPS4: Integrated History and Philosophy of Science, Department of Philosophy and History of Science, University of Athens, 15–18 March.

Empirical Philosophy of Science. Qualitative Methods: Sandbjerg, Denmark, 21–23 March.

Workshop on Philosophical and Formal Theories of Truth: Amsterdam, 23–25 March.

Pragmatism, Law, and Language: University of Idaho, 23–25 March.

LABCII: Logical Approaches to Barriers in Complexity II, Newton Institute, Cambridge, UK, 26–30 March.

CIFEr: Computational Intelligence for Financial Engineering and Economics, New York, 29–30 March.

New Science, New Risks: University of Pittsburgh, 30–31 March.

DICE: 3rd Workshop on Developments in Implicit Complexity, Tallinn, Estonia, 31 March–1 April.

April

YSM: Young Statisticians' Meeting, Cambridge, 2–3 April.

SBP: International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction, University of Maryland, 3–5 April.

Mind, Method and Morality: Pittsburgh, 6–7 April.

CNCS: International Conference on Computer Networks and Communication Systems, Malaysia, 7–8 April.

EMCSR: European Meetings on Cybernetics and Systems Research, Vienna, 10–13 April.

TIME FOR CAUSALITY: Workshop on Causal Inference and Dynamic Decisions in Longitudinal Studies, Bristol, 10–13 April.

evoSTOC: Evolutionary Algorithms in Stochastic and Dynamic Environments, Malaga, Spain, 11-13 April.

PHDS IN LOGIC IV: Ghent, 12–13 April.

NORTHWESTERN/NOTRE DAME GRADUATE EPISTEMOLOGY CONFERENCE: Northwestern University, Evanston, IL, 13–14 April.

BMC2012: Workshop on Turing's Legacy in Mathematics and Computer Science, University of Kent, 16–19 April.

CONFRONTING INTRACTABILITY IN STATISTICAL INFERENCE: University of Bristol, 16–19 April.

COLLECTIVE INTELLIGENCE: MIT, Cambridge, MA, 18–20 April.

BEING FREE, DOING FREE: Freedom Between Theoretical and Practical Philosophy, University of Freiburg, Germany, 19–21 April.

PSYCHOLOGY, EMOTION, AND THE HUMAN SCIENCES: University of Windsor, Windsor, Ontario Canada, 20–21 April.

MAICS: 23rd Midwest Artificial Intelligence and Cognitive Science Conference, Ohio, 21–22 April.

AISTATS: 15th International Conference on Artificial Intelligence and Statistics, La Palma, Canary Islands, 21–23 April.

THE PROGRESS OF SCIENCE: Tilburg Center for Logic and Philosophy of Science, 25–27 April.

SDM: 12th SIAM International Conference on Data Mining, Anaheim, California, USA, 26–28 April.

MAY

SOPHA: Société de philosophie analytique, Paris, 4–6 May.

ICDDM: International Conference on Database and Data Mining, Chengdu, China, 5–6 May.

BELIEF FUNCTIONS: Compiégne, France, 9–11 May.

NATURALISM AND NORMATIVITY IN THE SOCIAL SCIENCES: University of Hradec Králové, Czech Republic, 10–12 May.

CASI: 32nd Conference on Applied Statistics Ireland, 16–18 May.

GAMES, GAME THEORY AND GAME SEMANTICS: 8th International Symposium of Cognition, Logic and Communication, Riga, Latvia, 18–20 May.

SLACRR: St. Louis Annual Conference on Reasons and Rationality, 20–22 May.

IPDPS: 26th IEEE International Parallel and Distributed Processing Symposium, Shanghai, China, 21–25 May.

JDS: 44th Journées de Statistique, Brussels, 21–25 May.

UR: Uncertain Reasoning, Special Track at FLAIRS-25, Marco Island, Florida, USA, 23–25 May.

SSHAP: Mind, Language and Cognition, McMaster University, Canada, 24–26 May.

PHILMILCOG: 10th Annual Graduate Conference on the Philosophy of Mind, Language and Cognitive Science, University of Western Ontario, 24–26 May.

THE AIMS OF INQUIRY AND COGNITION: Edinburgh Epistemology Research Group, University of Edinburgh, 25–26 May.

CSAE: IEEE International Conference on Computer Science and Automation Engineering, Zhangjiajie, China, 25–27 May.

ICKD: 2012 International Conference on Knowledge Discovery, Indonesia, 26–27 May.

AI2012: Canadian Conference on Artificial Intelligence, 28–30 May.

RTA: 23rd International Conference on Rewriting Techniques and Applications, Japan, 28 May–2 June.

FEW: 9th Annual Formal Epistemology Workshop, Munich, 29 May–1 June.

ICCC12: Third International Conference on Computational Creativity, Dublin, 30 May–1 June.

STOCHMOD: 4th meeting of the EURO Working Group on Stochastic Modeling, Ecole Centrale Paris, 30 May–1 June.

HUMAN COMPLEXITY: The University of North Carolina, Charlotte, 30 May–1 June.

CAMBRIDGE PRAGMATISM: a Research Workshop, Cambridge, UK, 31 May–1 June.

JUNE

ADVANCES IN PHILOSOPHICAL LOGIC: Ruhr University Bochum, 3–5 June.

FEW: Formal Epistemology Week, Konstanz, 4–6 June.

AAMAS: 11th International Conference on Autonomous Agents and Multiagent Systems, Valencia, Spain, 4–8 June.

MFPS: 28th Conference on the Mathematical Foundations of Programming Semantics, University of Bath, 6–9 June.

MINDS, BODIES, AND PROBLEMS: Bilkent University, Ankara, 7–8 June.

EDINBURGH EPISTEMOLOGY GRADUATE CONFERENCE: University of Edinburgh, 8–9 June.

FOUNDATIONS OF LOGICAL CONSEQUENCE: University St Andrews, 8–10 June.

NORDSTAT: 24th Nordic Conference in Mathematical Statistics, Northern Sweden, 10–14 June.

MS5: Conference on Models and Simulations, Helsinki, 14–16 June.

CSAM: Classification Society Annual Meeting, Carnegie Mellon University, Pittsburgh, PA, 14–16 June.

BASIC KNOWLEDGE: Conference on the A Priori, Aberdeen, 16–17 June.

SAT: International Conference on Theory and Applications of Satisfiability Testing, Trento, Italy, 17–20 June.

LOFT: Tenth Conference on Logic and the Foundations of Game and Decision Theory, Sevilla, Spain, 18–20 June.

DM: Discrete Mathematics, Dalhousie University, Halifax, Nova Scotia, Canada, 18–21 June.

LOGICA: Hejnice, northern Bohemia, 18–22 June.

CɪE: Computability in Europe, University of Cambridge, Cambridge, 18–23 June.

SISSM: Scientific Meeting of the Italian Statistical Society, Rome, Italy, 20–22 June.

Pʜɪʟᴏsᴏᴘʜɪᴄᴀʟ Iɴsɪɢʜᴛs: Senate House, University of London, 21–23 June.

MBR12: Model-Based Reasoning in Science and Technology, Sestri Levante, Italy, 21–23 June.

HOPOS: Halifax, Nova Scotia, Canada, 21–24 June.

CCA: 9th International Conference on Computability and Complexity in Analysis, Cambridge, UK, 24–27 June.

COLT: 25th Annual Conference on Learning Theory, Edinburgh, 25–27 June.

MPC: 11th International Conference on Mathematics of Program Construction, Madrid, Spain, 25–27 June.

Aʀᴛɪғɪᴄɪᴀʟ Iɴᴛᴇʟʟɪɢᴇɴᴄᴇ ᴀɴᴅ Sᴏғᴛ Cᴏᴍᴘᴜᴛɪɴɢ: Naples, Italy, 25–27 June.

VᴀNɪM: Values and Norms in Modeling, Eindhoven, The Netherlands, 25–27 June.

LICS: 27th ACM/IEEE Symposium on Logic In Computer Science, Dubrovnik, Croatia, 25–28 June.

Sǫᴜᴀʀᴇ ᴏғ Oᴘᴘᴏsɪᴛɪᴏɴ: American University of Beirut, 26–29 June.

ICML: 29th International Conference on Machine Learning, University of Edinburgh, 26 June–1 July.

IJCAR: 6th International Joint Conference on Automated Reasoning, Manchester, UK, 26 June–1 July.

DGL12: Sixth Workshop in Decisions, Games & Logic, LMU Munich, 28–30 June.

Eᴠᴏʟᴜᴛɪᴏɴ ᴀɴᴅ Fᴜɴᴄᴛɪᴏɴ ᴏғ Cᴏɴsᴄɪᴏᴜsɴᴇss: Summer School in Cognitive Science 2012, Montreal, Canada, 30 June–9 July.

Jᴜʟʏ

Uɴᴄᴇʀᴛᴀɪɴᴛʏ ɪɴ Cᴏᴍᴘᴜᴛᴇʀ Mᴏᴅᴇʟs: Sheffield, UK, 2–4 July.

AISB/IACAP: Birmingham, UK, 2–6 July.

LASR: 31st Leeds Annual Statistical Research Workshop, University of Leeds, 3–5 July.

IIBM: 5th International Workshop on Intelligent Informatics in Biology and Medicine, Palermo, Italy, 4–6 July.

Hɪsᴛᴏʀʏ ᴀɴᴅ Pʜɪʟᴏsᴏᴘʜʏ ᴏғ Pʀᴏɢʀᴀᴍᴍɪɴɢ: Ghent University, 5–6 July.

CAV: 24th International Conference on Computer Aided Verification, Berkeley, 7–13 July.

ISSCSS: International Summer School in Cognitive Sciences and Semantics, Latvia, 8–18 July.

IPMU: 14th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Catania, Italy, 9–13 July.

ICALP: 39th International Colloquium on Automata, Languages and Programming, University of Warwick, 9–13 July.

Fᴏᴜɴᴅᴀᴛɪᴏɴs ᴏғ Mᴀᴛʜᴇᴍᴀᴛɪᴄs: University of Cambridge, 10–12 July.

DEON: 11th International Conference on Deontic Logic in Computer Science, University of Bergen, Norway, 16–18 July.

SIPTAss: Society for Imprecise Probability: Theories and Applications Summer School, Pescara, Italy, 16–20 July.

Paradox and Logical Revision: LMU, Munich, 23–25 July.

FOIS: 7th International Conference on Formal Ontologies in Information Systems, Graz, Austria, 24–27 July.

Courses and Programmes

Courses

LI: Logic and Interactions, Winter School and Workshops, CIRM, Luminy, Marseille, France, 30 January–2 March.

ESSLLI: 24th European Summer School in Logic, Language and Information, Opole, Poland, 6–17 August.

Programmes

APhil: MA/PhD in Analytic Philosophy, University of Barcelona.

Doctoral Programme in Philosophy: Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.

HPSM: MA in the History and Philosophy of Science and Medicine, Durham University.

LoPhiSC: Master in Logic, Philosophy of Science & Epistemology, Pantheon-Sorbonne University (Paris 1) and Paris-Sorbonne University (Paris 4).

Master Programme: in Artificial Intelligence, Radboud University Nijmegen, the Netherlands.

Master Programme: Philosophy and Economics, Institute of Philosophy, University of Bayreuth.

Master Programme: Philosophy of Science, Technology and Society, Enschede, the Netherlands.

MA in Cognitive Science: School of Politics, International Studies and Philosophy, Queen's University Belfast.

MA in Logic and the Philosophy of Mathematics: Department of Philosophy, University of Bristol.

MA in Logic and Theory of Science: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.

MA in Metaphysics, Language, and Mind: Department of Philosophy, University of Liverpool.

MA in Mind, Brain and Learning: Westminster Institute of Education, Oxford Brookes University.

MA in Philosophy: by research, Tilburg University.

MA in Philosophy of Biological and Cognitive Sciences: Department of Philosophy, University of Bristol.

MA in Rhetoric: School of Journalism, Media and Communication, University of Central Lancashire.

MA programmes: in Philosophy of Language and Linguistics, and Philosophy of Mind and Psychology, University of Birmingham.

MRes in Cognitive Science and Humanities: Language, Communication and Organization: Institute for Logic, Cognition, Language, and Information, University of the Basque Country, Donostia, San Sebastian.

MRes in Methods and Practices of Philosophical Research: Northern Institute of Philosophy, University of Aberdeen.

MSc in Applied Statistics: Department of Economics, Mathematics and Statistics, Birkbeck, University of London.

MSc in Applied Statistics and Datamining: School of Mathematics and Statistics, University of St Andrews.

MSc in Artificial Intelligence: Faculty of Engineering, University of Leeds.

MA in Reasoning

An interdisciplinary programme at the
University of Kent, Canterbury, UK.
Core modules provided by Philosophy and further modules from Psychology,
Computing, Statistics, Social Policy, Law, Biosciences and History.

MSc in Cognitive & Decision Sciences: Psychology, University College London.

MSc in Cognitive Science: University of Osnabrück, Germany.

MSc in Cognitive Psychology/Neuropsychology: School of Psychology, University of Kent.

MSc in Logic: Institute for Logic, Language and Computation, University of Amsterdam.

MSc in Mathematical Logic and the Theory of Computation: Mathematics, University of Manchester.

MSc in Mind, Language & Embodied Cognition: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.

MSc in Philosophy of Science, Technology and Society: University of Twente, The Netherlands.

MRes in Cognitive Science and Humanities: Language, Communication and Organization: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastian).

Open Mind: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.

PhD School: in Statistics, Padua University.

Jobs and Studentships

Jobs

POST-DOC POSITIONS: in Robot Learning and Reinforcement Learning, Intelligent Autonomous Systems Group, Darmstadt University of Technology / Technische Universitaet Darmstadt, Germany, to be filled asap.

POST-DOC POSITION: in the area of developmental robotics and robot learning, INRIA, Bordeaux, until filled.

TWO POST-DOC POSITIONS: in Machine Learning, in the project "Composing Learning for Artificial Cognitive Systems", INRIA Lille, until filled.

POST-DOC POSITION: in Machine Learning, University of Massachusetts, until filled.

POST-DOC POSITION: in Machine Learning, SUNY at Buffalo, until filled.

POST-DOC POSITION: in Philosophy of Mind, psychology, Neuroscience, and Computing, Department of Philosophy, University of Missouri-St. Louis, until filled.

POST-DOC POSITION: in Advanced Bayesian Computation for Cross-Disciplinary Research, Warwick Systems Biology Centre, until filled.

POST-DOC POSITION: in Probabilistic Graphical Models, Intelligent Systems Laboratory at Rensselaer Polytechnic Institute, New York, until filled.

POST-DOC POSITIONS: in all areas of speech and language processing at the Human Language Technology Center of Excellence at Johns Hopkins University, until filled.

POST-DOC POSITION: on the project "Explanatory Reasoning: Normative and Empirical Considerations," Tilburg Center for Logic and Philosophy of Science, until filled.

POST-DOC POSITION: applying probabilistic modelling and Bayesian statistical inference to problems in computational systems biology, University of Sheffield, deadline 9 January.

POST-DOC POSITION: in the Statistics Laboratory, University of Cambridge, deadline 15 January.

PROFESSOR AND TIER I CANADA RESEARCH CHAIR: in Epistemology and Metaphysics, Department of Philosophy, University of Alberta, deadline 15 January.

LECTURER: in Philosophy, AOS: Philosophy of Logic, Mathematics, and Formal Epistemology, University of St. Andrews, deadline 16 January.

POST-DOC POSITION: in Mathematical Logic at the Department of Mathematics, Stockholm University, deadline 23 January.

PROFESSOR: Department of Communication and Information Sciences, Tilburg Center for Cognition and Communication, deadline 1 February.

3-YEAR POST-DOC

To work on the relationship between Bayesian epistemology and inductive logic.
Philosophy, University of Kent, deadline 15 February

POST-DOC POSITION: in the History and Philosophy of Science, University of Pittsburgh, deadline 15 February.

PROFESSOR: of Statistics, Queen Mary, University of London, deadline 25 February.

FELLOWSHIP: in Statistics, University of Edinburgh, deadline 29 February.

Studentships

PHD POSITIONS: in Robot Learning and Reinforcement Learning, Intelligent Autonomous Systems Group, Darmstadt University of Technology / Technische Universitaet Darmstadt, Germany, to be filled asap.

THREE DOCTORAL TRAINING GRANTS: School of Computing, Faculty of Engineering, University of Leeds, until filled.

ONE DOCTORAL RESEARCHER POSITION AND ONE STUDENT RESEARCH ASSISTANT: in uncertain reasoning in the intersection of philosophy, psychology and cognitive science, Munich Center for Mathematical Philosophy, LMU Munich, until filled.

PHD POSITION: in the area of developmental robotics and robot learning, INRIA, Bordeaux, until filled.

PHD POSITION: in Bayesian Decision Theory, School of Computer Science and Statistics, Trinity College Dublin, until filled.

PHD POSITIONS: in Forensic Statistics, University of Twente / University of Amsterdam, until filled.

TWO PHD POSITIONS: in the project "Designing and Understanding Forensic Bayesian Networks with Arguments and Scenarios", Utrecht University / University of Groningen, to be filled asap.

PHD POSITIONS: in the Statistics & Probability group, Durham University, until filled.

PHD POSITIONS: in Statistical Methodology and its Application, University College London, until filled.

PHD POSITIONS: in Philosophy and History and Philosophy of Science, University of Leeds, deadline 23 January.

PHD POSITION: Knowledge Management research group at the Institute of Applied Informatics and Formal Description Methods, Karlsruher Institute of Technology, deadline 31 January.

PHD POSITION: at the Institute for Logic, Language and Computation (ILLC), deadline 1 February.

PHD STUDENTSHIP

To work on the relationship between Bayesian epistemology and inductive logic.
Philosophy, University of Kent, deadline 15 February

PHD POSITIONS: at the Department of Computer Science of the University of Liverpool, UK, deadline 15 February.