The Reasoner

Volume 1, Number 6
October 2007

www.thereasoner.org

Contents

§1 Editorial

This month sees the launch of an interdisciplinary MA in Reasoning at the University of Kent. The course will appeal to those students seeking to indulge their interests in a subject that employers will find appealing, as well as to those students seeking the broad interdisciplinary background required for postgraduate research in reasoning. If you know of undergraduates who are showing an interest in the study of reasoning, please make them aware of this possibility.

Probabilistic reasoning is very much in the news this month, with coverage of the meetings on Bayesian Nonparametric Regression (Isaac Newton Institute, Cambridge) and Combining Probability and Logic (Centre for Reasoning, University of Kent). Keeping with this theme, Colin Howson talks about probability, logic and 'progicism' in the interview below. Probability and its connections with category theory are currently up for blog discussion in The *n*-Category Café—see here and here.

The editorial team are very grateful to Federica Russo for her excellent work as assistant editor, and we wish her well in her new research post in Belgium. We will endeavour to keep The Reasoner running smoothly; please help us by continuing to submit your thoughts on reasoning and inference. Submission instructions are available at http://www.thereasoner.org/.

Jon Williamson
Philosophy, University of Kent

§2 Features

Interview with Colin Howson

Colin Howson is Professor of Logic in the Department of Philosophy, Logic and Scientific Method at the London School of Economics.

Jon Williamson: Could you say a bit about your intellectual history—what you've worked on and why, and what directions your studies and career have taken, how you became interested in logic and probability?

Colin Howson: I started as an undergraduate at LSE studying for an economics degree, thinking that I ought to do something likely to land me a good job when I left university. I had no previous knowledge of economics and found it boring. Popper's small LSE Philosophy department had just started a Philosophy degree course, and though I didn't know who Popper was I thought it looked a lot more interesting than economics, so I

switched to Philosophy. What really attracted me was not Popper or philosophy of science, but mathematical logic. But when I started thinking about a doctoral dissertation mathematical logic was looking just too hard (this was the post-Cohen epoch of work on large cardinals etc.), and I was left for a time in the air. After a time Peter Urbach, who had just joined the department, got me interested in Bayesianism, and Peter and I used to enjoy ourselves showing how Bayesianism solved all the methodological problems Popper claimed to have solved but actually had not. After that I became a sort of card-carrying Bayesian, and Peter and I wrote a joint book on Bayesian philosophy of science, now into a third edition. Hume's famous discussion of induction, and its modern follow-up, Goodman's 'grue' construction, *demonstrate* that without a differentiating prior probability any body E of data can be extrapolated in an infinity of mutually incompatible ways each bearing exactly the same logical relation to E. But while Bayesianism, or something very like it,[1] seems the obviously true methodology, I have always found the orthodox, decision-theoretic foundations offered for it unsatisfactory. By contrast, the striking similarity of the Kolmogorov formalism, with its underlying possibility-space understood as a set of logical possibilities, to that of logic suggested to me at any rate that as a foundation logic promises to be a better option. In fact it solves a lot of the problems associated with the decision-theoretic approach. To take just one example, on a logic-based foundational approach the so-called problem of logical omniscience, viz. that a human agent seems obliged to obey rules that exceed the abilities of a Universal Turing Machine, simply falls away. In the last few years I have been trying to show how the logic-based approach can be worked out in detail.

JW: What's the connection between logic and reasoning, in your opinion?

CH: My own view is that logic consists of a set of formal models of valid reasoning, none of which is perfect but which by and large offer a fairly good average fit. The current paradigm for logic is still first order logic and various extensions of it (including modal logic), but of course that is only a model of *deductive* reasoning. The idea of seriously trying to model uncertain reasoning within a framework both logical *and* probabilistic is really only just being taken seriously, and even then by no means everywhere.

JW: You're an advocate of progicism—could you explain what that is and why you advocate it?

CH: The previous sentence presents the progicist programme, at least as I understand it, in a nutshell. To go beyond the nutshell horizon is quite challenging and brings up questions whose answers, even among devout progicists, is a matter of some controversy. My own personal stamp of progicism is as a formal generalisation of deductive consistency: deductive consistency is the solvability of truth-value assignments subject to the constraints of a classical truth-definition, and probabilistic consistency is the solvability of probability assignments subject to the constraints of the probability axioms (this seems to have been de Finetti's view also, at any rate at certain periods of his life). Just as in deductive logic, such a view of probabilistic consistency immediately provides a corresponding definition of consequence: an assignment to a proposition π is a consequence of a joint assignment to the members of a set Π of propositions just in case every solution of Π determines a solution of π. Though pleasing, this idea immediately raises several big questions: (i) why those probability axioms?—(ii) and what exactly are those probability axioms?—(iii) in particular, should they include countable additivity?—(iv) should conditional or unconditional probability be taken as primitive (or both)?—(v) what are 'propositions': are they sentences in some formalised language, or algebraic/set-theoretical entities? Given the expressive limitations of formalised languages possessing a reasonable proof theory (which rules out second order languages under standard semantics) the latter option is more promising if one wants to capture 'ordinary' probabilistic reasoning. People however are still hung up on the idea that logical validity must be recursively enumerable, so work in the area of logical and probability still concentrates overmuch on first order languages. But, possibly stimulated by the development of quantum logic, people are getting used to the idea of representing propositions purely algebraically, as members in the classical case of a complete Boolean algebra and in the quantum one the projection lattice of a Hilbert space.

JW: In your view what are the most exciting and important research directions in the area of reasoning? What topics would you recommend to graduate students starting out today?

CH: What excites one person can be without any interest for another, ditto for importance. Computer scientists would probably say that producing efficient reasoning algorithms is the most exciting and important research area. I myself would nominate a satisfactory unification of deductive and probabilistic logic.

JW: You're based at LSE—could you say a bit about what's happening in the study of reasoning there?

CH: I will be leaving LSE to take up a position in the University of Toronto from 2008. Both institutions treat reasoning as pretty much identified with deductive logic, including of course modal logic, and Toronto in particular has a fairly comprehensive range of courses from elementary to advanced. Given the increasing overlap between logic, computer science, AI, cognitive science and psychology I think the proposed University

[1] 'And almost every one when age,/ Disease, or sorrows strike him;/ Inclines to think there is a God,/ Or something very like him'. (Arthur Hugh Clough).

of Kent Masters programme in Reasoning broadly construed is an innovation which should be widely copied.

Why we shouldn't fault Lucas and Penrose for continuing to believe in the Gödelian argument against computationalism - I

Most reasoned critiques (e.g., 1996, *PSYCHE*, 2(23); 2000, *Journal of Experimental and Theoretical Artificial Intelligence* 12: 307-329; 1993, *Behavioral and Brain Sciences*, 16, 611-612) of Lucas' and Penrose's well-known Gödelian arguments against computationalism are unassailable. They fail, however, to satisfactorily explain why Lucas and Penrose—reasonable men, both—remain convinced of the essential soundness of their arguments.

One reason could be that Lucas and Penrose have unquestioning faith in, and uncritically follow, standard expositions of classical theory in overlooking what Gödel (1931: *On formally undecidable propositions of Principia Mathematica and related systems I*. In M. Davis. 1965. *The Undecidable*. Raven Press. New York. Theorem VI. p24) has implicitly proven; another, that they have similar faith in, and, as uncritically, accept Gödel's (1931, p27) own, informal, interpretation of the implications of this Theorem as definitive.

If so, they should not be taken to task on either count for their faith; it is standard expositions of Gödel's reasoning that are ambiguously silent on both issues.

Specifically, as we show below, the Tarskian *satisfaction* and Tarskian *truth* (A. Tarski 1936: *The concept of truth in the languages of the deductive sciences*. In *Logic, Semantics, Metamathematics, papers from 1923 to 1938*. Hackett Publishing Company) of the formulas of an Arithmetic, and the *soundness* of the Arithmetic itself, *can*, indeed, be *formally* defined in an effectively verifiable manner within the Arithmetic.

So, why do such definitions elude serious enquiry by Lucas and Penrose?

One factor could be that Tarski's Theorem (ibid. 1936) *appears* to implicitly suggest that the intuitive truth, of the formulas of Peano Arithmetic under its standard interpretation, *cannot* be formalised within the Arithmetic.

> **Tarski's Theorem:** The set of Gödel-numbers of the formulas of any first-order Peano Arithmetic, which are intuitively true in the standard model of the Arithmetic, is not arithmetical.

Lucas and Penrose rely on the implication unquestioningly in their Gödelian arguments. However, the implication does not withstand scrutiny.

Tarski's definitions of *satisfiability* and *truth*

To see this, we note the following, standard, definitions of the *satisfiability*, and *truth*, of formulas of a formal language, say L, under a well-defined interpretation, say M, due to Tarski (ibid. 1936).

For instance, a formula $[R(x)]$ of L is defined as satisfied under M if, and only if, its corresponding interpretation, say $R(x)$, holds in M for any assignment of a value s that lies within the range of the variable x in M.

> Tarski's definitions are mathematically significant only if we assume that, given any s in M, we can decide whether, or not, $R(s)$ *holds*, or must *hold*, in M. Where L is PA, the Church-Turing Thesis postulates that such decidability must be *algorithmic*. Note that, in principle, this *can* be weakened to effective, *instantiational*, decidability only—the minimum requirement of Tarski's definitions.

The formula $[(Ax)R(x)]$ of L is, then, defined as true under the interpretation M if, and only if, $[R(x)]$ is *satisfied* under M.

Moreover, the formula $[\neg(Ax)R(x)]$ of L is, further, defined as *true* under the interpretation M if, and only if, $[(Ax)R(x)]$ is *not true* under M.

Clearly, mathematical satisfaction and truth are defined relative only to decidability in an interpretation.

Both, Lucas and Penrose, quite reasonably, therefore, attempt to draw philosophical conclusions from the meta-logical status of the *intuitive decidability* given in mathematical reasoning to the formulas of Peano Arithmetic under its standard, *intuitive*, interpretation.

Defining *formal satisfaction* and *formal truth* verifiably

However, if we take M to be an interpretation of L in L itself, then we have the formalisation of the concepts of verifiable, and unarguable, *formal satisfaction*, and *formal truth*, of the formulas $[R(x)]$ and $[(Ax)R(x)]$ of L, respectively, in L, as:

> The formula $[R(x)]$ of L is defined as *formally satisfied* under L if, and only if $[R(s)]$ is *provable* in L for any term $[s]$ that can be substituted for the variable $[x]$ in $[R(x)]$.
>
> The formula $[(Ax)R(x)]$ of L is *formally true* in L if, and only if, $[R(x)]$ is *formally satisfied* in L.

Defining *formal soundness* verifiably

If we, further, define *formal soundness* as the property that the axioms of a theory are *satisfied* in the theory itself, and that the rules of inference preserve *formal*

truth, then, it follows that the theorems of any *formally sound* theory are *formally true* in the theory.

It is straightforward to verify that first-order Peano Arithmetic is, indeed, *formally sound*.

Gödelian propositions

Now, even if the formula $[(Ax)R(x)]$ is *not* provable in L, it would be *formally true* in L if, and only if, the formula $[R(s)]$ were provable in L for *any* well-defined term $[s]$ of L that could be substituted for $[x]$ in $[R(x)]$.

The existence of such a, Gödelian, proposition is, precisely, what Gödel (1931, p24) proves in his Theorem VI for a consistent Peano Arithmetic. He constructs a formula, $[(Ax)R(x)]$, of PA that is, itself, unprovable in PA, even though, for any given numeral $[n]$, $[R(n)]$ is provable in PA.

So, Gödel has constructed a formally *unprovable* Arithmetical formula that is not only *intuitively true* in the standard, *intuitive*, interpretation of the Arithmetic, but which is also *formally true* in the Arithmetic in a verifiable, and intuitionistically unobjectionable, manner that leaves no room for dispute as to its 'truth' status vis-à-vis the PA axioms!

Moreover, since the Arithmetic can be shown to be *formally sound*—again in a verifiable, and intuitionistically unobjectionable, manner—we no longer need appeal to the arguable assumption that the Arithmetic is *intuitively sound* under the standard interpretation.

> PA is said to be *sound* for a class S of sentences if, whenever PA proves $[\varphi]$ with $[\varphi]$ in S, then φ is *true* in the structure N of natural numbers.

Replacing the, philosophically debatable, concepts of *intuitive decidability* with, constructively verifiable, definitions of *formal decidability* should, thus, place Lucas' and Penrose's Gödelian arguments in better perspective.

Bhupinder Singh Anand
Mumbai

There is no question about it!

Martin Mose Mentzen is suffering from a dose of use-mention confusion when puzzling about "Is the answer to this question 'No'?" in The Reasoner Vol 1 No 5 (I supply the needed quotation marks he significantly omits). Of course, he is not alone in this but that is a much longer story.

Bentzen's 'Question 1', namely 'Is the answer to this question no?' contains a demonstrative 'this' which could be used to refer to many things (likewise 'question 1' itself). So if one attempts to specify the intended, self-referential question one must take the quotation marks away from the whole, and use the demonstrative (or some substitute for it, like 'question 1') with the supposed reference. Hence what Bentzen has in mind is the following expression in indirect speech:

Question 1 is whether the answer to question 1 is 'No'.

But clearly, as a definition of a question this is circular. Alternatively, it leads to an infinite regress, as William Kneale pointed out in 1972 with respect to the supposed proposition associated with Liar sentences such as 'this is not true' ('Propositions and Truth in Natural Languages', Mind 83, 225-243).

The point also holds with Bentzen's 'Imperative 1' i.e. 'Violate this imperative', which, upon unpacking the intended sense, leads to the regress in 'Imperative 1 is to violate imperative 1'. It even holds with the seemingly much more puzzling 'Is this a question?' For surely that is definitely a question, to which the answer consequently cannot be either the rhetorically expected 'no', or even 'yes', since seemingly there is no question about its being a question. But while there is the grammatical form of a question, nothing is asked, and so no answer is required. For again the quoted 'this', in the grammatical form, or some replacement for that demonstrative, must be used rather than just mentioned, to generate the supposed intention, and equally there is a regress in

Question 5 is whether question 5 is a question.

So there isn't even a question, in the self-referential case.

Barry Hartley Slater
Philosophy, University of Western Australia

'Can,' the Principle of Relevant Alternatives, and Moral Responsibility

When we say that "A ought to do x" we are saying something about the future actions of A, namely that it is permissible for her do x and that she can do it. According to standard usage when we claim that A can do x we are maintaining that 1) she is able to do x, 2) she knows how to do x, and 3) she has the right to do x. When we say that person A ought to do x, we are maintaining that A has a reason that explains and justifies her doing x. What we are not claiming is that she 'could have' performed an alternative action; ought does not imply could have. The philosopher who would go beyond standard usage in order to include the notion of relevant alternative options must provide reasons for doing so that does not beg the question. She might reason that the standard usage does not allow us to analyze why we hold a person morally responsible and that by expanding the definition to include the availability of alternative actions we are thereby able to perform this analysis. However, the standard usage does allow for this analysis if we focus on the third criterion as the

morally essential criterion that must be met if we are to maintain that it is permissible for a person to perform some action. Presuming that if a person performs an action that she is able to perform it and knows how to perform it, then the only morally relevant question is 'did she have the right to perform it.' This can be analyzed without referencing relevant alternative actions. If we claim that we can only be morally responsible if we could have performed some other action then the one we performed, we are misunderstanding an important feature of what it is that we want the concept of moral responsibility to accomplish, namely change peoples' future actions, as well as hold them accountable for what they did.

A person is morally responsible for an action if she performed the action and had a reason that she thought explained and justified her performing the action. If we claim that what she did was morally acceptable then we agree that her reason explains and justifies what she did. If we think that what she did was morally unacceptable, we are claiming that the reason does not justify her action even if it does explain why she performed it. We are not making any claim as to the availability of any alternative course of action that she could have performed. When we analyze the morally responsibility of a person it is the third criterion that is most important. But having the right to do x does not entail, or even imply, that there are any options available to A at the time she performs x. It simply means that A has a reason that she thinks justifies and explains what she did. If no errors are made in the reasoning process, this process creates the right for performing x that other rational agents will recognize and accept as sufficient for explaining and justifying A performing x. If A is acting as a rational agent then it is the reasoning process, not the choice that eliminates other options. At some time prior to performing x the option not to do x was available to A, but the reasoning process that A employed between the time x and -x were available and the time she performed x determined that she perform x when she did. The reasoning process determines what a person is committed to doing if that person is to be considered rational.

This form of determinism is not to be confused with causal determinism. Rational determinism is simply the willingness to follow the argument such that if the premises of the argument are believed to be true and the conclusion follows from the premises then the conclusion must be accepted and acted upon if one is to be considered rational. If we do not accept the reasoning of A, it does not follow that A could have done anything other then what she did. After all, A believes that she is justified in performing x so she performs it. We simply do not accept her reasoning and we think she was somehow mistaken when she reasoned that x was the correct action to perform. Somewhere along the 'chain of reasoning' that A performed she made a mistake and had she recognized this mistake she would not have performed x. In an important sense, 'morally responsibility' is not about the past. There is no doubt that A did x. We cannot change that fact. When we hold that a person is negatively morally responsible we are maintaining 1) that she did not have the right to do x and should be accountable for doing so (retributivism), but more importantly 2) that in the future we do not want her, or others, to perform x (deterrence). We want people to do a better job of reasoning regarding the actions they are considering performing. When we say that A ought not to do x we are talking about future actions, we are not debating that A did, or did not, do x in the past; that is given. "Can" refers to an action not yet performed, it does not pertain to actions already performed. Holding a person morally responsible for her action does not hold out any possibility of changing that action already performed, but it does hold out the possibility of changing future actions. Holding someone morally responsible for an action is a sanction, either positive or negative, relative to the acceptability of an action already performed (retributivism) and also an attempt to influence (determine) future actions (deterrence). 'Could have done otherwise' does not pertain to moral responsibility.

John K. Alexander
Grand Valley State University

Referential Usage and Gödelian Completions

Neale (2004, "This, That, and the Other", In M. Reimer and A. Bezuidenhout (eds.), *Descriptions and Beyond*, Oxford: Oxford UP, 68-182; 2005, "A Century Later", *Mind*, 114, 809-71) argues that Russellians ought to drop the traditional Gricean approach to referentially used definite descriptions in favor of a 'Gödelian' approach. On the former approach, a referential use of 'The F is G' implicates, but does not express, a singular proposition containing *a*, where *a* is *S*'s intended referent.[2] On the latter approach, referential usage is treated as a special kind of 'incomplete', elliptical, usage. In particular, *S*'s utterance of 'The *F* is *G*' is said to be elliptical for an utterance of 'The *F identical to a* is *G*'. Accordingly, *S*'s elliptical utterance expresses a general proposition, albeit one containing *a* as a constituent. Unlike the Gricean, then, the Gödelian assimilates referential usage to what is said rather than to what is meant.

To illustrate the difference between the two approaches, consider Donnellan's (1966, "Reference and Definite Descriptions", *The Philosophical Review*, 75, 281–304) original case of referential usage. Sally is

[2] See, for example, Neale (1990, *Descriptions*, Cambridge: MIT Press).

attending Jones's trial. Jones is accused of murdering Smith. Jones's frequent psychotic outbursts throughout the trial prompt Sally to claim that:

1. Smith's murderer is insane.

On the Gricean approach, Sally's utterance expresses (2) and generates (3) as a conversational implicature:

2. [The x: x murdered Smith] (x is insane).

3. Jones is insane.

In essence, Sally *says* (2) but *means* (3). On the Gödelian approach, however, Sally's utterance of (1) goes proxy for an utterance of 'The murderer of Smith *identical to Jones* is insane'. Accordingly, Sally's utterance of (1) expresses:

4. [The x: x murdered Smith & x = Jones] (x is insane). Here, Sally *simply says* (4).

Now, Neale contends that the Gödelian approach is preferable to the Gricean approach because, unlike the latter approach, it can withstand Devitt's (2004, "The Case for Referential Descriptions", In M. Reimer and A. Bezuidenhout (eds.), *Descriptions and Beyond*, Oxford: Oxford UP, 280—305) argument from convention. The argument from convention runs as follows. Referential uses of definite descriptions, Devitt observes, are regularized, systematic, and cross-linguistic. Referential usage, then, is *conventional usage*. As such, it is a semantic phenomenon. Thus, any non-semantic approach to referential usage is dubious. The Gricean approach, then, is dubious; it is a pragmatic approach. The Gödelian approach, however, survives Devitt's argument. The Gödelian approach is a semantic approach; it handles referential usage at the level of what is said. Accordingly, Neale urges that the Gödelian approach replace its more traditional counterpart.

The Gödelian approach, however, is not without its own problems. In particular, the Gödelian approach struggles with two issues its traditional counterpart readily handles. The first issue concerns misdescription; the second concerns what Ludlow and Segal (2004, "On a Unitary Semantical Analysis for Definite and Indefinite Descriptions", In M. Reimer and A. Bezuidenhout (eds.), *Descriptions and Beyond*, Oxford: Oxford UP, 420—36) call 'residual misdescription'. I consider each in turn.

Consider this twist on Donnellan's original case. Smith was indeed murdered. But it wasn't Jones; Bill murdered Smith. Now, Bill is sane. Jones, however, still remains insane. As Neale rightly claimed, in this scenario Sally's utterance of (1) presents us with an uneasy tension:

> We want to say that [Sally] did something right but also that [Sally] did something *wrong*. After all, the description [she] used *failed* to fit the person [Sally] wanted to 'talk about,' and to that extent the speech act was defective. (1990, 91; emphasis in the original).

Now, the Gricean has a straightforward explanation of the uneasy tension these 'misdescription' cases generate. Sally's utterance expresses a false proposition; the murderer, Bill, is not insane. Nonetheless, Sally conversationally implicates something true, namely (3). That is, Sally conveys something true while uttering a *falsehood*. Consequently, she does something right but also something wrong.

In contrast, the Gödelian seems unable to explain the uneasy tension. For the Gödelian, Sally's utterance expresses (4); she does not convey another proposition. So, on the Gödelian approach, Sally surely does something wrong; she expresses a false proposition. She does not seem, however, to do anything right. On the Gödelian approach, then, misdescription cases shouldn't produce an uneasy tension. The Gödelian must maintain that in misdescription cases S only does something wrong; she utters a falsehood. Surely, this position is counterintuitive.

Now, consider a second twist on Donnellan's case. Smith did not have just one murderer; two people had a hand in the gruesome act. Bill and Jones did the deed. Jones, however, is alone on trial; his insanity made him easy to catch. When Sally utters (1), we are again struck by an uneasy tension. In these 'residual misdescription' cases, one wishes to maintain that Sally has done something right but also something wrong.

Here, too, the Gricean has a perfectly intelligible explanation of the uneasy tension gripping us. Sally's utterance expresses a false proposition; there is more than one murderer. Thus, she does something wrong. Nonetheless, she also does something right; she, quite correctly, conveys (3). By exploiting the saying/meaning distinction, the Gricean can gesture towards a workable account of the uneasy tension residual misdescription cases generate.

The Gödelian, however, seems unable to explain the uneasy tension. According to the Gödelian, Sally's utterance simply expresses a true proposition! Sally correctly says that the individual who murdered Smith and who is identical to Jones is insane. On the Gödelian approach, Sally does nothing wrong at all. Thus, there should be no uneasy tension to explain. This position also seems counterintuitive.

So, although the Gödelian approach allows the Russellian to circumvent the argument from convention, it does so at an extremely heavy price. The Russellian who follows Neale can no longer provide accounts of the uneasy tension misdescription and residual misdescription cases generate. I doubt many Russellians are willing to pay that high a price.

Francesco Pupa
The Graduate Center, CUNY

§3
News

Report on the Programme: Bayesian Nonparametric Regression, Isaac Newton Institute, Cambridge, 30 July – 24 August 2007

Basic Theme and Background

Bayesian nonparametric inference is a relatively young research area, with the earliest papers appearing in the mid 1970's, with greatly increasing activity over the last 10 years. The group of researchers working in this field is still moderately small and there was a need for a focus opportunity, which the BNR programme provided, to survey the field, exchange ideas, identify gaps in the literature, and coordinate research efforts. Specifically in the area of regression, where there is a lot of current activity with many possible avenues of investigation.

Thus the purpose of the BNR programme was to review the current state of nonparametric Bayesian research, to foster collaborations between researchers with different focus areas, and to identify important open problems. One prominent example is the need to combine research concerned with the construction of random probability measures with another direction of research concerned with random mean functions in a regression problem.

There is a need to build and strengthen connections between emerging areas of nonparametric Bayesian inference. Besides the already mentioned connection of random distributions and regression functions, other examples are the relevance of asymptotic results to model choice and inference, generalizations of predictive probability functions and random clustering models. We were particularly pleased for a number of persons from the Machine Learning community to attend the Programme, who highlighted practical problems to be solved.

The four organizers met early in 2005 for an initial planning meeting when we agreed on the broad focus, on invitation lists and priorities. We also decided at that time to have the 2nd week workshop, and to leave maximum flexibility and opportunity for informal interaction for the remaining Programme.

Structure

In weeks 1, 3 and 4, the programme organization was very informal. We organized two informal talks each day. The nature of these talks varied, with some presentations being descriptions of current research ideas, some describing work in progress, and some that reported on recently completed research. The format of these talks was chosen to encourage interaction and "lively" discussions.

For week 2 (6 to 10 August, 2007) the Programme organizers were responsible for a Workshop on Bayesian Nonparametric Regression. The first day was dedicated to 4 tutorials, with speakers Subhashis Ghosal, Antonio Lijoi, Wee Teh and David Dunson. The remaining four days were regular talks. There were about 90 participants at the Workshop and over 40 speakers.

The first day of the workshop consisted of four tutorials, which provided an excellent survey of methods in nonparametric Bayesian inference. The tutorials were extremely successful. For similar conferences, usually only part of the audience attends introductory tutorials. But for this workshop we noticed that 80% and more of the participants stayed for all four tutorials.

The remaining days of the workshop were research talks, including some talks on innovative applications of nonparametric Bayesian inference (Mukherjee, Ruggiero, Laud, Herring, Popova, House, Williams, Mena) novel methods for survival analysis (Johnson, Guglielmi, De Iorio, Yin), issues of prior choice (Cox, MacEachern), new non-parametric models (Quintana, Griffin, Arjas, Pruenster), constructions of families of dependent random probability measures (Spano, Petrone, Steel, Dunson, Basu), and asymptotic results (Choi, De Blasi, Lee).

Outcome and Achievements

The Programme successfully surveyed the field and pointed towards specific directions for future applications of research. It generated interest in alternative models that generalize the most traditional ones which currently dominate practical nonparametric Bayesian methods. At the same time, several talks highlighted and discussed distinguishing features of these traditional models.

Most of the talks and presentations in the informal part of the Programme addressed the question that initially motivated the organization of the Workshop, namely the combination of random probability measures and random regression mean functions to construct real nonparametric regression models. Several talks focused on the construction of families of random probability measures indexed with a covariate. This formally provides a constructive definition of the desired combinations.

Aside from the organized collection of participants at the talks; feedback from participants made it quite clear that many collaborations are under way and many new projects initiated as a consequence of the intensity of the Programme.

A book is planned with the CUP as publishers. The contents would be the material from the 4 tutorials, given by Antonio Lijoi, Subhashis Ghosal, Wee Teh and David Dunson, with each of the organizers providing a chapter.

The organizers would like to express their deep gratitude to the INI for their outstanding hospitality and for accepting us as guests for a month.

Stephen Walker
Statistics, University of Kent

Progic 2007: the Third Workshop on Combining Probability and Logic, University of Kent, Canterbury, 5–7 September 2007

The southeast of England hosted two major events last month. One of them was the World Beard and Moustache Championships, the other was the third edition of what can now legitimately be called a series of Progic conferences. There have been two previous editions of Progic conferences, one in 2002 at Kings College London on combining probability and logic, and one in 2005 at the London School of Economics, of which the special focus was objective Bayesianism. The special focus of this year's Progic conference was probabilistic logic and probabilistic networks. Incidentally, this is the topic of the research collective Progicnet, whose members (Rolf Haenni-Sola, Jan-Willem Romeijn, Greg Wheeler, and Jon Williamson) took part in organising the conference, together with Federica Russo.

Naturally the reader may raise questions on the objectivity of this report, but here I am reporting that the conference was a success. Thanks to generous support by the Leverhulme Trust, the programme consisted of no less than ten invited speakers, next to some ten contributed papers, covering a wide variety of topics. The quality of these talks was invariably high, and despite the variety in topics all of the talks were interesting to most of the attendants. It will unfortunately not be possible to go through all of the invited talks, let alone provide an overview of all talks; the interested reader will find abstracts and slides of all the talks on the conference website. Instead, I will discuss a number of issues that came up repeatedly in the conference, or during the round table discussion at the end.

One such issue was beautifully brought forward by speakers from the computer science community, like Cozman, Domingos, Getoor, and Jaeger. Each of them brought out in a different way how network structures can be used to speed up and simplify probabilistic calculations. It was very instructive and inspiring, at least from the side of philosophers and logicians, to see probabilistic logic at work in the somewhat less abstract context of computer science.

Another issue concerns the choice that each probabilistic logic must make between sharp probability values or sets of such values as representation of uncertainty. Representatives of both sides were present at the conference: Hartmann, Howson, Makinson, and Paris all set up their probabilistic logics using sharp values, while de Cooman, Cozman, and Progicnet all choose sets of probabilities. As was indicated by the probabilistic logicians from the computer science community, sets of probabilities bring additional complexity to an already difficult area. In applications, we will have to have strong reasons to motivate the use of imprecise probabilities. But on the side of philosophical logic, the principal aim is not in the applications, but rather in getting the representation of uncertain opinion or belief right, and in coming up with a sensible notion of valid inference under uncertainty. It seems that philosophers remain divided on the issue of the correct representation of uncertain opinion.

As a third general issue then, what would determine the correctness of the representation and of the associated probabilistic logic? For the many probabilistic logicians who take De Finetti as starting point, the criteria for the representation of opinion derive from the behavioural consequences of opinion, such as buying a bet or running for the train you might not catch. One of the motivations for imprecise probability is that the lowest selling price for a bet is typically higher than the highest buying price, suggesting that the associated probability can be determined up to an interval. However, this invites the question whether a probabilistic logic should perhaps be supplemented with a decision theory, or even be designed in conjunction with a decision theory from scratch. It may be that as a normative theory, a standalone probabilistic logic misses a firm foundation.

At the end of the day, these abstract worries over belief representation did not prevent the participants of the other world conference in southeast England to come up with strong opinions, and even decide on a number of winners. I refer again to the Progic 2007 website, which links through to the beard and moustache champions.

Jan-Willem Romeijn
University of Groningen

Mechanisms and Causality

A 3-year interdisciplinary research project, Mechanisms and Causality, has just begun at the University of Kent. The principal investigator is Jon Williamson and I am a postdoctoral researcher. In the first year the project will focus on comparing the use of mechanisms in biology, economics, psychology and physics. In subsequent years the project will go on to apply the results to think about the metaphysics of causation, and the use of mechanisms in causal inference. It hopes to bring real insights into the working of the sciences themselves and into theorising about causality.

We have organised a reading group on mechanisms to run this semester, and would like to invite anyone interested to attend—including any interested scientists working in the relevant fields. See the web page for

readings and meeting times. We would also like to hear from anyone who works with mechanisms in any capacity who has any interest in the project, and welcome any suggestions. Please contact either Jon Williamson or Phyllis McKay.

Phyllis McKay
Philosophy, University of Kent

Out Now

I would like to announce my new book

The Minimum Description Length Principle, Peter Grunwald, MIT Press, 2007. 570 pages. USD 45.

For more information and sample chapters, please visit http://homepages.cwi.nl/~pdg/book/book.html

This book provides the first comprehensive introduction and reference guide to the minimum description length (MDL) Principle. The central concepts of this theory - both mathematical results and philosophical foundations - are explained in great detail.

The book consists of four parts. Part I provides a basic introduction to MDL and an overview of the concepts in statistics and information theory needed to understand MDL. Part II treats universal coding, the information-theoretic notion on which MDL is built, and Part III gives a formal treatment of MDL theory as a theory of inductive inference based on universal coding. Part IV provides a comprehensive overview of the statistical theory of exponential families with an emphasis on their information-theoretic properties. The book contains some new results that have not been published elsewhere.

Peter Grünwald
Information-Theoretic Learning Group,
CWI Amsterdam

Journal of formalized reasoning

Charter and Scope

The journal encourages submission of papers describing significant formalization efforts in any area, including classical math, constructive math, formal algorithms, and program verification. The emphasis of the journal is on proof techniques and methodologies and their impact on the formalization process. In particular, the journal provides a forum for comparing alternative approaches, enhancing reusability of solutions and offering a clear view of the current state of the field.

Editorial Policy

The Journal of Formalized Reasoning is a purely electronic publication. The Journal provides immediate open access to its content on the principle that making research freely available to the public supports a greater global exchange of knowledge.

The Journal encourages the submission of interesting, insightful, and informative papers, preferably accompanied by formal scripts. Although the journal is not meant to become a repository for proof scripts, an effort will be made to ensure that the "experimental data" backing formalisation papers will remain accessible. Each new contribution must be original, and not submitted before different journals for consideration. Detailed instructions can be found at the journal home page http://jfr.cib.unibo.it/.

Calls for Papers

Adaptive Monte Carlo Methods: Special issue of Statistics and Computing, deadline 1 October 2007.

Social Learning in Embodied Agents: Special issue of Connection Science Journal, alberto.acerbi@istc.cnr.it, deadline 30 October 2007.

Introducing . . .

The Reasoner would like to publish very short introductions to key terms, people and texts in logic and reasoning. Selected pieces will also be published in a book "Key Terms in Logic" by Continuum. If you would like to contribute, please contact TheReasoner@kent.ac.uk

Mathematics and Argumentation: , Special Issue of Foundations of Science, deadline 1 November 2007.

Machine Learning Algorithms: Machine Learning Algorithms for Event Detection, Special Issue of Machine Learning Journal, deadline 28 November 2007.

Spatial Statistics: Special Issue of Computational Statistics and Data Analysis, deadline 30 November 2007.

Evolutionary Intelligence: Special Issue on Artificial Immune Systems, deadline 1 December 2007.

Conditionals and Ranking Functions: Special issue of Erkenntnis, franz.huber@uni-konstanz.de, deadline 31 May 2008.

§4
Events

TBILISI: The Seventh International TBILISI Symposium on Language, Logic and Computation, 1–5 October 2007.

SUM 2007: First International Conference on Scalable Uncertainty Management, Washington DC Area, 10–12 October 2007.

Models and Simulations 2: Three-day conference at the Tilburg Center for Logic and Philosophy of Science, 11-13 October 2007.

Reason, Intuition, Objects: The Epistemology and Ontology of Logic, Buffalo, 13 October 2007.

LPAR 2007: Logic for Programming, Artificial Intelligence and Reasoning, Yerevan, Armenia, 15–19 October 2007.

Women in Machine Learning: Orlando, Florida, 17 October 2007.

Case Studies of Bayesian Statistics: The Ninth Workshop on Case Studies of Bayesian Statistics, Carnegie Mellon University, Pittsburgh, 19–20 October 2007.

ISDA'07: 7th International Conference on Intelligent Systems Design and Applications, October 22-24, 2007, Rio de Janeiro, Brasil.

MWPMW 8: Eighth annual Midwest PhilMath Workshop, to be held at Notre Dame, 27–28 October 2007.

ECSQARU'07: Ninth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, October 31, 1–2 November 2007, Hammamet, Tunisia.

Novembertagung: 18th Novembertagung on the history, philosophy and didactics of mathematics, Bonn, Germany, 1–4 November, 2007.

Infinity: International Conference on Infinity in Logic and Computation, University of Cape Town, South Africa, 3–5 November 2007.

Curves: Modern challenges of curve modelling: inverse problems and qualitative constraints, Bristol, 7–9 November 2007.

Aesthetics and Mathematics: University of Utrecht, 9–10 November 2007.

Context-Dependence: Context-Dependence, Perspective and Relativity in Language and Thought, Ecole Normale Supérieure, Paris, 9–11 November 2007.

URSW: 3rd Workshop on Uncertainty Reasoning for the Semantic Web, Busan, Korea, 12 November 2007.

EPSA07: 1st Conference of the European Philosophy of Science Association, Madrid, 15–17 November 2007.

LogKCA-07: ILCLI International Workshop on Logic and Philosophy of Knowledge, Communication and Action, Donostia, 28–30 November 2007.

M4M5: Methods for Modalities 5, Cachan, France, 29–30 November 2007.

NIPS: Neural Information Processing Systems—Natural and Synthetic, Hyatt Regency Vancouver, 3–6 December 2007.

Workshop: International Workshop on Applied Bayesian Statistics, EpiCentre, Massey University, Palmerston North, New Zealand, 10–14 December, 2007.

ISAIM 2008: Tenth International Symposium on Artificial Intelligence and Mathematics, Fort Lauderdale, Florida, 2-4 January 2008.

3rd IMS and ISBA meeting: The third joint international meeting of the IMS (Institute of Mathematical Statistics) and ISBA (International Society for Bayesian Analysis) will be held in Bormio, Italy from Wednesday, January 9 to Friday, January 11, 2008.

Graduate Conference: 1st Cambridge Graduate Conference on the Philosophy of Logic and Mathematics, St. John's College, Cambridge, 19–20 January 2008.

FoIKS 2008: Foundations of Information and Knowledge Systems, Pisa, Italy, 11–15 February 2008.

Artificial General Intelligence: The First Conference on Artificial General Intelligence, Memphis, Tennessee, 1–3 March 2008.

Constraint-Sac2008: Track on Constraint Solving and Programming, at the 23rd Annual ACM Symposium on Applied Computing, Fortaleza, Brazil 16–20 March 2008.

Causation: 1500-2000: King's Manor, University of York, 25–27 March 2008.

RelMiCS10-AKA5: 10th International Conference on Relational Methods in Computer Science & 5th International Conference on Applications of Kleene Algebra, Frauenwörth, Germany, 7–11 April 2008.

Reduction and the Special Sciences: Tilburg Center for Logic and Philosophy of Science, 10–12 April 2008.

Workshop: XVIII Inter-University Workshop on Philosophy and Cognitive Science, Madrid, luis.fernandez@filos.ucm.es, 22–24 April 2008.

SIG16: 3rd Biennial Meeting of the EARLI-Special Interest Group 16—Metacognition, Ioannina, Greece, 8–10 May 2008.

UR 2008: Special Track on Uncertain Reasoning, 21st International Florida Artificial Intelligence Research Society Conference (FLAIRS-21), Coconut Grove, Florida, USA, 15–17 May 2008.

COMMA'08: Second International Conference on Computational Models of Argument Toulouse, France, 28–30 May 2008.

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

HOPOS 2008: Seventh Congress of the International Society for the History of Philosophy of Science, Vancouver, Canada, 18–21 June 2008.

EPISTEME: Law and Evidence, Dartmouth College, 20–21 June 2008.

ISBA 2008: 9th World Meeting, International Society for Bayesian Analysis, Hamilton Island, Australia, 21–25 July 2008.

First Formal Epistemology Festival: Conditionals and Ranking Functions, Konstanz, 28–30 July 2008.

Conference: Language, Communication and Cognition University of Brighton, 4–7 August 2008, Brighton, UK.

Soft Methods for Probability and Statistics: 4th International Conference, Toulouse, France, 8–10 September 2008.

Valencia Meetings: Valencia / ISBA Ninth World Meeting on Bayesian Statistics, Spain, June 2010.

§5 Jobs

2-year Postdoc: Konstanz University, Germany. The Emmy Noether junior research group Formal Epistemology, two year postdoctoral research position in Philosophy, on the project 'Belief and Its Revision', deadline 1 November 2007.

1 PhD and 1 Postdoc Position: Institute for Logic, Language and Computation, University of Amsterdam, project on Computational Social Choice, deadline 9 November 2007.

Assistant Professor TiLPS: Tilburg Centre for Logic and Philosophy of Science, Tilburg University, deadline 12 November 2007.

Three Assistant Professorships: The Department of Cognitive Science at the University of California, invites applications for three faculty positions at the Assistant Professor level (tenure-track), deadline November 15 2007.

Assistant Professor: Texas A&M University, College Station, TX. Assistant Professor, tenure-track, beginning Fall 2008, deadline 15 November 2007.

Desicion Support Systems Project: Dynamic system analysis & uncertainty research position; one year position in Human Computer Interaction; two year position in Psychology of Decision Making, The Cork Constraint Computation Centre.

§6 Courses and Studentships

Courses

Andreas Weiermann: In the winter term 2007/8, lectures on phase transitions in logic and combinatorics at the Department of Mathematics in Ghent starting in the last week of September (26/09/2007).

Discrete Choice Modelling: Centre for Transport Studies at Imperial College London, 28–30 November 2007.

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.

Second Indian Winter School on Logic: IIT Kanpur, 14–26 January 2008.

Studentships

PhD position: The Department of Philosophy and the Tilburg Center for Logic and Philosophy of Science (TiLPS) invite applications for a three-year full-time PhD position, commencing January 1, 2008. The deadline for applications is October 25, 2007.

Two PhD Positions: The Faculty of Philosophy of the University of Groningen, the Netherlands, starting January 1, 2008. Applications should be sent by 1 November 2007 to T.A.F.Kuipers@rug.nl

Three PhD positions: Within the Swiss National Science Foundation Research Module Norms, epistemic, rational and social of the recently accepted pro*doc graduate programme in philosophy, the Philosophy Departments of the Universities of Geneva, Lausanne and Fribourg offer three PhD positions (respectively one per university), before November 31, 2007.

Logic and Philosophy of Science: 4 year PhD position or a 80%-funded 4 year post-doctoral research position, The Center for Logic and Philosophy of Science at the Vrije Universiteit Brussel, sonsmets@vub.ac.be, deadline 1 December 2007.

Acknowledgements

The Reasoner is supported by a Strategic Research Development Fund grant from the University of Kent. It is a development of the progicnet academic network supported by the Leverhulme Trust.