
THE REASONER

VOLUME 15, NUMBER 5
NOVEMBER 2021

thereasoner.org
ISSN 1757-0522


CONTENTS

Editorial	35
Features	35
News	37
What's Hot in ...	40
Events	43
Courses and Programmes	43
Jobs and Studentships	44

EDITORIAL

Dear Reasoners,

It is my pleasure to introduce Natalie Gold, Senior Director in the Behavioural Practice at Kantar Public, and Visiting Professor in Practice at LSE. Natalie is best known for her contributions to decision making, framing and team agency. In the interview which opens this issue, we chatted about her multidisciplinary academic background and inter-sectoral experience. In addition to this, Natalie was kind enough to share her thoughts on the rather poor academic career-happiness balance which many know all too well.



Before leaving you to the interview, you may have noticed that this issue is appearing with some delay. This owes to the fact that we are in the process of revamping The Reasoner. Next month you will see the first output of this process, with a *focussed issue* on “Evidential Pluralism” edited by Jon Williamson.

If you'd like to let the community know about an exciting topic through a focussed issue, please get in touch at features@thereasoner.org.

HYKEL HOSNI
University of Milan

FEATURES

Interview with Natalie Gold

HYKEL HOSNI: Can you tell about your background?

NATALIE GOLD: As an undergraduate I studied philosophy, politics and economics, and I was lucky to be able to do it in a truly interdisciplinary way. In Week 1 of microeconomics I was hooked by the rationality assumptions underpinning Expected Utility Theory but also, contrastingly, by the clarity that simple rational choice models could bring to explanations of behaviour.

HH: Why the contrast?

NG: The axioms seemed neither to be tenets of rationality nor to describe actual human decision-making, yet the models based on them were very powerful and offered a lot of insight into human behaviour.

HH: You were telling us about your undergraduate interests ...

NG: I got into political theory more gradually, as we got away from the first-year text-based syllabus, which I didn't en-



joy and which made me think I was not interested in philosophy—but then I realised I was so long as it was focussed on contemporary analytical questions.

HH: You didn't say where you studied.

NG: I was educated at Oxford. It was a tutorial environment where I could have a syllabus that was tailored to my interests, so I studied a lot of rational choice political theory. In my final year, I was taught economic theory by Michael Bacharach, which was the perfect combination for me, of philosophical foundations of economics and actual evidence about how people made decisions.

HH: How did you choose for your postgraduate work?

NG: I didn't know which of the two subjects to pursue, economics or political theory. Economics seemed to keep more options open—throughout my career people kept telling me that if I left economics I wouldn't get back in—and the idea was to pick up the technical skills that an economics education would give me but then apply them elsewhere. It was clear to me from Day 1 of my M.Phil in Economics I that I was never going to be an academic economist.

HH: Why?

NG: I wasn't enough of a mathematician. People only cared about solving problem sets; there were optional essays every week, but no-one wrote them and they were never discussed in classes. I was grateful to David Miller for inviting me to the Nuffield political theory group and to Andrew Glynn for agreeing to give me tutorials on history of economic thought on the side, so that I could spend some part of my week discussing ideas.

And of course Michael, who became my supervisor, was a huge influence. He really drove me towards the psychology literature because I would make claims about decision-making and he would write marginal comments asking for evidence to back them up. So I started spending copious amounts of time in the psychology library. Michael ran the Bounded Rationality and Economics Behaviour group. Behavioural economics in the UK was quite a small crowd back then, and we often interacted with people at University of East Anglia and Leicester. Everyone was pretty philosophically minded, so it was a great mix for me.

HH: What was your thesis on?

NG: I wrote my D.Phil on framing. It was a neat fit for Michael's work because he was interested in framing as a mechanism that would lead to group identification, and therefore cooperation and coordination. For me, his project was an example of why framing effects might not be irrational, if framing oneself as a member of a group could lead agents to make different choices. But that *if* is an empirical question. And I would say that is the key theme of my academic work, which has ranged over a number of topics: that there is no neat separation between the normative and the empirical. Pretty much all my work has involved bringing together (or actively contributing to) an empirical literature in order to bring it to bear on a normative question.

HH: But the normative/empirical is not the only boundary you crossed. It is relatively unusual for researchers in your field to be based only partly in academia. Can you tell us what motivated you to consider (partially) reorienting your career?

NG: So many things... the most proximate cause is that I wanted to live in London and I never got an academic job there. If I had got one, then I might still be unhappily in academia. The ultimate cause is that I was unhappy because I found uni-

versities to be terrible employers, conditions for academics bad and getting worse, and lots of academics unhappy and even at each others' throats.

HH: Do you think it is our nature, or is it rather a question of incentives?

NG: That is another current research interest, the huge influence that institutions have on people's behaviour and in what sense 'bad' institutions bear responsibility for bad outcomes. I resented that universities exploit academics' desire to do research, filling their paid time with other duties and relying on them to do research on their own time. I worked out that I could get an interesting job which actually fitted within the 37.5 hours of the week I was contracted for and then left me plenty of time for research on the top.

HH: Can you tell us what this job involves

NG: My first non-academic job was at Public Health England (PHE), an executive agency of the Department of Health and Social Care, in PHE Behavioural Insights. It was a good fit for my more empirical side. We would work with colleagues in policy teams to explore how behavioural science could contribute to issues they were working on, and then they would either commission us to conduct the research or we would support them to commission it externally. I'm now at Kantar Public, in its Behavioural Practice, so I've moved to the other side of the table, being commissioned by government departments to apply behavioural science to solve their problems.

HH: Do you think your academic side contributes to it?

NG: Part of my academic side involves running experiments on decision-making, so there are clear synergies with the applied research I conduct in my non-academic job. Indeed, I would go so far as to say there is no clear distinction between academic and non-academic research. There is sometimes a deference towards academic research, but knowing how academics spend their whole time disagreeing with each other—not to mention the replication crisis—I am more critical than my colleagues and more likely to push on the weak points in any research.

HH: Conversely, does your research benefit from being employed outside academia?

NG: Depends on the topic. Probably not for anything super-abstract. However, if I am making a more applied argument in political theory, it is helpful to have seen what is going on from the inside and to have an idea of how actors within the system are thinking about the issue. It is very hard to say anything impactful if one's viewpoint is completely abstract—unless one has the good fortune for one's ideas to be taken up by someone who is more situated.

HH: Going back to your career decisions, did you understand that academia made you unhappy after partly leaving it?

NG: I had realised I was unhappy for a while before I actually left academia, and I applied for and was offered several non-academic jobs before I took one. Something always came up that made it seem like the wrong time to make the move. Eventually, I precipitated my exit by moving from a permanent position at Edinburgh to a fixed-term contract at Kings College London and, as the FTC was drawing to an end, it was time to choose London over academia. But I didn't realise quite how unhappy I had been in academia until after I left, when my partner went round telling everyone how much happier I was now in my new job!

HH: What advice would you give to someone who after their PhD is unsatisfied by the academic career prospect or is

attracted to the non-academic path?

NG: Think about your transferable skills, and learn how to display them in a non-academic resume and how to use them in an interview context. Investigate how recruitment works, which varies by sectors. For instance, in the UK public sector, there are standard recruitment processes and lists you can sign up to to receive job ads; then there are standard application and interview processes, including types of questions you can prepare for in advance. In the private sector, the team that is recruiting often has much more freedom and there can be value in meeting people or being recommended by someone, who can send your CV straight to someone on the inside.

Most importantly, remember that, outside of academia, people move jobs every few years, so your first non-academic job is not a forever job: it is just a foot in the door somewhere, which gets you experience that you can use to make your next move. So choose something that you think you might enjoy and where you will be able to develop and grow. Then take the plunge.

NEWS

Prolog 2021: Combining Probability and Logic. 31 August - 3 September (virtual)

The 10th event in the Prolog series of conferences aimed at combining Probability and Logic took place online on 1-3 September 2021. The conference has been preceded by a one day Summer School. The Conference has been organized by Jürgen Landes (LMU Munich), funding was provided by the German Research Foundation (DFG) and virtually hosted by the Munich Center for Mathematical Philosophy.

In the Summer School, Peter Grünwald (CWI and Leiden University) gave a lecture on Safe Probability clarifying much of the discussion between strict Bayesians, imprecise probabilists and frequentists and pointing towards a unified view. Grünwald introduced a new pragmatic notion of probabilistic truth based on a hierarchy of safety notions and explained how safe probability hints towards interesting generalization of measure theory. Claudia d'Amato (University of Bari) presented how Machine Learning and Knowledge Graphs (KGs) can be related. In particular, in her lecture at the summer school, d'Amato showed how Machine Learning methods are suitable also in case of incoherent or noisy Knowledge Graphs and how they can be seen as an additional layer on top of deductive reasoning for introducing new forms of approximated reasoning capabilities. Vincenzo Crupi (University of Turin) explained how by taking a parametric family of scoring rules (including so-called logarithmic and Brier scores as special cases) as a basic building block, is possible to formalize notions such as epistemic inaccuracy, uncertainty, information and evidential support. In addition, Crupi showed how using a formal representation, many issues gain a much clearer presentation. In the final lecture of the summer school, Simon Huttegger (University of California Irvine) presented the problem of merging of opinions and convergence to the truth by Bayesian agents within the framework of algorithmic randomness.

During the conference, Ute Schmid (University of Bamberg) has introduced inductive logic programming (ILP) as a powerful interpretable machine learning approach that allows combining logic and learning. Schmid showed how ILP can be

combined with different methods for explanation, generation and proposed a framework for human-in-the-loop learning. Jon Williamson (University of Kent), presenting the results of joint work with Jürgen Landes and Soroush Rafiee Rad (University of Amsterdam) showed some progress on the problem of how to determine maximal entropy functions for objective bayesian inductive logic. Alena Vencovská (University of Manchester) presented results concerning pure inductive logic in the case of languages with binary and possibly unary predicates. Olav Vassend (Nanyang Technological University) addressed some issues concerning the principle of indifference using a novel accuracy-based argument.

Alexandru Baltag (University of Amsterdam) and Soroush Rafiee Rad presented a work showing that Leitgeb's method for extracting qualitative beliefs from probabilistic degrees of belief can be generalised to de Finetti's "qualitative probability", that takes the relation $A < B$ ("event B is at least as probable as event A") as the primitive notion. Vincenzo Crupi presented also during the conference and showed some results of joint work with Andrea Iacona (University of Turin). During the talk, Crupi suggested that incremental probabilistic support is key to the understanding of indicative conditionals and their role in human reasoning. Conor Mayo-Wilson (University of Washington) presented a joint work with Konstantin Genin (University of Tübingen) on Statistical decidability in confounded, linear, non-Gaussian causal models (LiNGAMs). In the presentation, Mayo-Wilson showed that the orientation of every edge in a LiNGAM is statistically decidable, even in the presence of confounding variables and hinted at profound implications of their result for methodology in the biomedical sciences.

On the following day of the conference, Claudia d'Amato gave a talk on the role that semantics and reasoning may play when developing numeric based solutions targeting KG completion at instance level. Furthermore, d'Amato discussed the role of semantics and reasoning on symbol-based machine learning solutions targeting KG enrichment at schema level. Giuliano Rosella (University of Turin) showed how to expand causal modelling semantics of counterfactuals in order to account for the probability of counterfactuals with complex antecedents. In particular, he showed how to calculate the probability of counterfactuals with disjunctive antecedents with respect to a causal model. Esther Anna Corsi (University of Milan) reported on a joint work with Tommaso Flaminio (IIIA - CSIC) and Hykel Hosni (University of Milan). Using geometrical investigation of de Finetti's Dutch Book method, she showed the existence of rich enough sets of events for which the coherence criteria for belief functions and lower probability collapse. Joe Roussos (Institute for Futures Studies, Stockholm) addressed the question on how an agent should update their beliefs when they encounter a completely new possibility. Roussos presented a new model of growing awareness which specifies how an agent's probabilistic beliefs ought to be extended to a new space of possibilities and which provides constraints for the formation of new beliefs about those new possibilities. Francesca Zaffora Blando (Carnegie Mellon University) talked of the phenomenon of merging of opinions and the phenomenon of polarisation of opinions for computationally limited Bayesian agents from the perspective of algorithmic randomness. Jonathan Vandenberg (Northwestern University) argued that the model of hypothesis testing can explain how people learn complex, theory-laden propositions like con-

ditionals and probability constraints.

On the last day of the conference, Peter Grünwald presented some new results on an alternative way for significant testing: the e-value. After having pointed out some non-negligible limits of the p-value Grünwald showed how the e-value can solve some of them. Niki Pfeifer (University of Regensburg) presented two probabilistic approaches to check the validity of key principles of connexive logic within the setting of coherence and showed that coherence-based probability logic offers a rich language to investigate the validity of various connexive principles. Paul Thorn (Heinrich Heine University, Düsseldorf) presented two case studies based on computer simulations of the connection between environment entropy and the reliability of non-deductive inference. Through these studies, Thorn showed that attending to variations in the reliability of inference methods across environments with different entropy is an excellent means of evaluating the reasonableness of the methods. William Peden (Erasmus University Rotterdam) showed how to compare the three main approaches to statistics (standard Bayesian, frequentist, and entropy-maximising Bayesian based on Jon Williamson's inductive logic) by simulating a decision problem involving bets on a series of binomial events. They coded three players, each based on each of the approaches and showed that all players have comparable performances. Tom Sterkenburg (LMU Munich) presented a work on the no-free-lunch theorem showing that every data-only learning procedure must possess some inductive bias. However, Sterkenburg pointed out that many standard learning algorithms are better conceived of as model-dependent, and can be given a general model-relative justification. Jeremy Goodman (University of Southern California) and Bernhard Salow (Oxford University) presented a new theory of knowledge and belief and showed that modelling belief with two closely related normality orders is really fruitful. Finally, Simon Huttegger showed how it is possible to extend the superconditioning theorem in two ways: a shift from a prior to a set of posteriors providing insights about the reflection principle, and a shift from a prior to two or more distinct sets of posteriors.

For more information, videos and slides of the talks visit the [conference website](#).

ESTHER ANNA CORSI
University of Milan

31st International Conference on Logic Programming 2021, 20-27 September (virtual)

Due to the pandemic, this year's International Conference on Logic Programming (ICLP 2021) was once again held virtually rather than in the wonderful city of Porto, where it was originally to be located. Nevertheless, there was an exciting and taxing program, filling 8 days from the 20th to the 27th of September.

Although traditionally the main focus of the ICLP lies in the areas of Prolog and Answer Set Programming, it has also been the premier venue for research on probabilistic logic programming ever since Sato introduced the distribution semantics at the 1995 edition of the conference.

This year, one session of the main conference was devoted to probabilistic logic programming, and additionally the 8th Workshop on Probabilistic Logic Programming (PLP 2021) was colocated with ICLP once again.

Furthermore, the considerable interest of the ICLP community in wider statistical relational AI was evidenced by the choice of invited talks, which included a presentation on the open-universe statistical relational language BLOG by Stuart Russell and a survey of neural knowledge representation featuring several statistical relational frameworks by William Cohen.

Four of the five original papers on probabilistic logic programming presented at the main track were dedicated to extensions of the distribution semantics.

Marco Gavanelli reported on joint work with Elena Bellodi, Riccardo Zese, Evelina Lamma and Fabrizio Riguzzi on adapting the distribution semantics to abductive reasoning with probabilistic integrity constraints. Abductive Logic Programming has been used with great success in applications such as fault diagnosis, where possible causes are hypothesised from an observed consequence. Since in practice the integrity constraints, the relationships known to hold in a situation, are often not absolutely certain, incorporating probabilistic integrity constraints is a natural extension. Furthermore, by incorporating classical techniques from abductive logic programming with the distribution semantics, their work allows abducting non-ground literals with free variables too.

Mario A. Leiva reported briefly on joint work with Alejandro J. Garcia, Paulo Shakarian and Gerardo I. Simari on Probabilistic Defeasible Logic Programming, in which a probabilistic model is combined with a defeasible model. The latter employs default reasoning based on argumentation to make inferences from presumptions and potentially contradictory information.

Damiano Azzolini presented joint work with Fabrizio Riguzzi on a syntax for hybrid probabilistic logic programs, which incorporate continuous distributions into probabilistic logic programming. This is of course vital to many putative application domains, in which continuous distributions over some types of numerical data occur together with the binary or categorical distributions typically associated with probabilistic logic programming. While semantical aspects have been studied before, and the semantics is known to be well-defined, this contribution suggested a concrete syntax for hybrid probabilistic logic programs and discussed necessary syntactical restrictions. Damiano Azzolini also presented a recently published joint paper with Fabrizio Riguzzi and Evelina Lamma in which the authors extend the semantics of hybrid probabilistic logic programming even further to allow for function symbols to occur in the program.

Finally, Damiano Azzolini presented joint work with Fabrizio Riguzzi on Probabilistic Optimizable Logic Programs, in which in addition to probabilistic facts and rules there are also optimisable facts, constraints and an objective function. The system then finds among those probability assignments on the optimisable facts that satisfy the constraints the assignment that minimises the objective function. From a statistical relational AI standpoint this contributes to a line of work on relational optimisation while methodically it is closely linked to parameter learning in probabilistic logic programs. Aptly, this work builds on the PITA engine which underlies the cplint probabilistic logic programming system, since the 10-year-test-of-time award was given to the very paper from ICLP 2011 which introduced the PITA system in the first place.

My own contribution was more analytically in nature and evaluated the asymptotic probabilities arising from probabilistic logic programs on large domain sizes. By leveraging techniques from classical finite model theory, I was able to show

that asymptotically, probabilistic logic programs degenerate to so-called determinate probabilistic logic programs. These are essentially propositional in nature and are well-known to be particularly simple to work with. This gives rise to interesting non-expressibility results, since it implies that non-trivial projective families of distributions, which do not degenerate on large domain sizes cannot be expressed by probabilistic logic programs.

Before the conference Rafael Penaloza and I chaired the Probabilistic Logic Programming workshop, in which new technical developments, applications and work in progress around the topic could be presented. There was an interesting mix of contributions, ranging from an application of the ProbLog 2 engine in epidemic modelling, which I presented as joint work with Beatrice Sarbu and Kailin Sun to theoretical results on the decidability of independence queries in open-universe probabilistic logic programs, presented by Kilian Rueckschloss (joint work with myself). In addition, Markus Hecher presented recently published joint work with Thomas Eiter and Rafael Kiesel on Algebraic Answer Set Counting, important to probabilistic logic programming through its applications to probabilistic inference.

Of particular interest were the invited talks by Stuart Russell and William Cohen, which highlighted frameworks that have been less-studied in the recent past. Stuart Russell presented BLOG, a statistical relational language based on Bayesian networks whose particular interest comes from its support for uncertainty not just about the rules and facts but about the domain itself. This is highly relevant in applications where the domain is not fixed or where the identity of putatively distinct domain elements is unclear (think for instance of different patient records that might in fact refer to the same person).

In a broad sweep of formalisms for neural and symbolic integration, William Cohen dwelled particularly on Stochastic Logic Programs and the closely related formalism TensorLog. They are of interest because they are among the few probabilistic logic programming approaches not to adopt the distribution semantics but instead to define probability distributions over the resolution proof trees, which is reminiscent of the program trace semantics used in other probabilistic programming languages.

For me one of the conclusions of this conference was the interest in languages and frameworks for probabilistic logic programming that lie outside the distribution semantics, but also the flexibility of the distribution semantics in admitting extensions to various other paradigms.

The list of papers accepted to the conference, with the titles of all the contributions mentioned here, can be found at <https://iclp2021.dcc.fc.up.pt/index-acceptedPapers.html>.

As is by now traditional for ICLP, selected papers are published in two forthcoming special issues of Theory and Practice of Logic Programming. They can be accessed as soon as they are processed at

<https://www.cambridge.org/core/journals/theory-and-practice-of-logic-programming/firstview>.

In addition, there is an EPTCS volume of technical communications available at <http://eptcs.web.cse.unsw.edu.au/content.cgi?ICLP2021>

which contains all those contributions that have not been selected for publication in TPLP.

The proceedings of the workshops, including PLP 2021, can be found in the dedicated <http://ceur-ws.org/Vol-2970/>.

FELIX WEITKÄMPER
LMU, Munich

New Work on Induction and Abduction, 29-30 September (virtual)

Inductive and abductive reasoning are indispensable not only in science, but also in philosophy, as the current work investigating these forms of reasoning clearly shows. The workshop *New Work on Induction and Abduction*, held online from September 29–30, 2021 brought together scholars from the fields of logic, epistemology, metaphysics, and philosophy of science in order to discuss the new insights and controversies regarding inductive and abductive reasoning.

The workshop focused on discussing four recent monographs: Igor Douven’s “The Art of Abduction” (2021), Ilkka Niiniluoto’s “Truth-Seeking by Abduction” (2018), John Norton’s “The Material Theory of Induction” (2021), and Gerhard Schurz’ “Hume’s Problem Solved” (2019). Each of these monographs was commented upon by a renowned specialist in the corresponding field of research, and after that the replies and reflections by the authors of the monographs followed. The workshop also hosted presentations from leading scholars in this field of research, exploring and clarifying different aspects of inductive and abductive reasoning.

Oliver R. Scholz and Ansgar Seide (both WWU Münster) gave the first talk on *Induction, Abduction and Inductive Metaphysics. Historical Background and Systematic Perspectives*, in which they explored similarities and differences between inductive metaphysics as a methodological or meta-metaphysical research program, on the one hand, and inductive metaphysics as a historical movement in the 19th and early 20th century philosophy, on the other hand.

Elke Brendel (University of Bonn) highlighted in her *Commentary talk on Gerhard Schurz’ “Hume’s Problem Solved”: Justifying Induction vs Justifying Deduction* the importance and originality of Schurz’ attempt to resolve Hume’s famous problem of induction by the help of optimality justification. However, she also argued that this form of justification cannot be expanded to deduction, because the putative optimality of classical logic in the sense of universal translatability of non-classical logics into classical logics does not work for well-established non-classical logics.

Gerhard Schurz (University of Düsseldorf) in his *Replies & Reflections* talk addressed this criticism by presenting new results about the translatability of four kinds of non-classical logics into classical logic: many-valued, intuitionistic, paraconsistent, and quantum logics. He argued for a generalization of optimality justifications towards a new program for foundation-theoretic epistemology, which can be applied for deductive, inductive, and abductive reasoning.

Adam Carter (University of Glasgow) presented his work on *Abduction, Scepticism, and Indirect Realism*, where he argued that abductive inference plays an important role in attaining perceptual knowledge. According to Carter, by making the transition from animal to reflective knowledge, a knower gains an epistemic perspective on her belief, from which she endorses the source of that belief as reliably truth-conducive, and thereby

improves the quality of antecedently attained perceptual knowledge.

Stathis Psillos (University of Athens) and Chrysovalantis Stergiou (American College of Greece) gave a commentary talk on John Norton's "The Material Theory of Induction". They noted that, according to Norton, there are no universal principles of induction: all inductive inferences are material and warranted by local background facts. Psillos and Stergiou argued that Norton's material theory makes presuppositions in order to account for a regress problem, which seem to them no more appealing than presuppositions made by John Stuart Mill in his account of enumerative induction. They also argued that Norton's solution to Hume's original problem is based on premise circularity.

John Norton (University of Pittsburgh) concluded the first day of the workshop with his *Replies & Reflections*. Norton stressed that his material theory of induction does not rely on a principle of uniformity of nature and also that he does not object to such a principle because it is universal in scope. Rather, his main objection to the principle is that it is either so vague as to be inapplicable or just factually false. Norton also presented his approach of abduction as a two step structure, where the first step consists in comparing a favored theory with its foils and demonstrates that it is better in terms of accuracy and "evidential depth"; the second step is the more challenging one (and in practice oftentimes neglected) and aims to demonstrate that the favored theory is not only better but actually the best.

Christian J. Feldbacher-Escamilla (University of Cologne) and Gerhard Schurz opened the second day of the workshop with their talk on *Epistemic Engineering: The interplay of meta-induction and abduction in the justification of laws of nature*. They argued that meta-induction is a prediction method that solves the problem of induction by, first, employing optimality justifications instead of reliability justifications; and, second, using the past track record of induction to justify that induction is an optimal choice for making predictions. They also considered recent objections to this method and argued that these can be answered by using a principle of cognitive coherence and a weak inductive uniformity assumption that plays also an important role in the justification of scientific laws.

Paul Thorn (University of D'usseldorf) in his commentary talk on Igor Douven's "The Theory and Practice of Abduction": *Abduction, Induction, and Direct Inference* considered some ways by which one form of non-deductive inference might depend on another and argued that the reasons given in Douven's book for thinking that abduction is not dependent on induction are not entirely conclusive.

Igor Douven (INSHS, Paris) in his *Replies & Reflections* responded to this criticism by providing a detailed example (an evolutionary simulation) of two competing strategies of reasoning, inductive vs abductive, running against each other in an expert prediction game, and getting different results (with abduction being more successful in the long run).

Alexandros Apostolidis (University of Athens) and Stathis Psillos gave a talk on *Why Formal Abduction is not IBE*. They argued that AKM (Aliseda – Kowalski, Kuipers, Kakas – Magnani, Meheus) models recently proposed to formalize Inference to the Best Explanation (IBE) by means of explanatory abduction and minimal abduction are unsuccessful. They are neither internally equivalent with IBE, because their criteria for determining the best explanation are different, nor externally equivalent with IBE, because there exists at least one class of

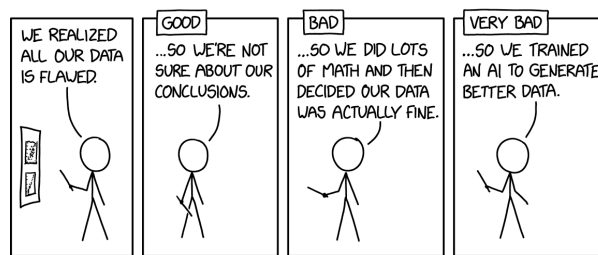
abductive problems where they end up with different solutions. Stephen Biggs (Iowa State University) and Jessica Wilson's (University of Toronto at Scarborough) talk was entitled *Does Anti-Exceptionalism about Logic Entail that Logic is Justified A Posteriori?* They argued that abduction is an *a priori* mode of inference, and considered the consequences that this thesis has for the proper understanding of anti-exceptionalism about logic. In particular, they argued that the justificatory status of logic turns not on the role played by abduction as such, but on the justificatory status of *a priori* or *a posteriori* data on which abduction operates.

Atocha Aliseda Llera (National Autonomous University of Mexico, UNAM) gave the last commentary talk of the workshop, namely, on Ilkka Niiniluoto's "Truth-Seeking by Abduction": *Truth-Seeking by Abduction: A Rule for Progress in Science?* She considered the place of this book in the context of research on scientific change in a post-Kuhnian era in the philosophy of science, and discussed in detail Niiniluoto's account of abduction and truthlikeness; as well as his notion of abductive belief revision.

Ilkka Niiniluoto's (University of Helsinki) *Replies & Reflections* was the final talk of the workshop. Niiniluoto provided a historical sketch of the debate about verisimilitude and truthlikeness and outlined how this debate can be and also was in fact linked to the discussion of abduction, stressing also that he considers the most recent development in this area of research as highly promising.

The workshop was organized by the research unit *Inductive Metaphysics*, supported by the German Research Foundation (FOR 2495). The goal of the research unit is to establish how empirical sources and inductive forms of inference play a role in metaphysical research. The particular workshop organisers were Christian J. Feldbacher-Escamilla, Oliver R. Scholz, Gerhard Schurz, Ansgar Seide and Maria Sekatskaya.

MARIA SEKATSKAYA AND CHRISTIAN J. FELDBACHER-ESCAMILLA



Calls for Papers

CAUSALITY: CONTEMPORARY APPROACHES: special issue of *Portuguese Journal Philosophy*, deadline November 30.

REASONING WITH INCONSISTENT, INCOMPLETE, AND UNCERTAIN KNOWLEDGE: special issue of *IEEE Intelligent Systems*, deadline 30 December.

THE EPISTEMOLOGY OF COMPUTER SIMULATIONS: special issue of *Balkan Journal of Philosophy*, deadline 30 December.

WHAT'S HOT IN . . .

Statistical Relational Artificial Intelligence

For most of its history, artificial intelligence was intimately connected with the art of automated reasoning. In a classi-

cal system of symbolic artificial intelligence, knowledge is encoded by experts as facts and rules in a suitable formal language. Using those, new knowledge can be derived by reasoning. Facts and rules could be highly sophisticated, involving intricate relations between different entities. Paradigmatic for this approach are expert systems and other artificial intelligence applications based on *logic programming*, and in particular on the language Prolog and its dialects. While this leads to a highly expressive way of representing knowledge, a major limitation soon emerged: The world is indeed complex, but it is also uncertain, and general rules hardly hold without exception.

Although probability theory and statistics are well-developed, time-honoured parts of science, incorporating probabilistic reasoning into artificial intelligence was long considered out of the question, since it would be impossible to specify all conceivable correlations and dependencies between factors. However, with the introduction of Bayesian Networks in the late 1980s, transparent independence assumption suddenly made this viable. Since then, statistical machine learning has developed into a large area of artificial intelligence. In the process, the emphasis shifted from reasoning about complex relationships to learning the dependencies between simple properties.

However, the world has not become any less interconnected since then. Therefore, since the mid-1990s, the new field of STATistical Relational Artificial Intelligence (StarAI) aims to bring those paradigms together. This can be approached from either direction: Either, one can expand logic programming to incorporate probabilities, or one can extend graphical models to incorporate relational information. In this short review, which will be followed by regular columns highlighting new research in the field, we will discuss the former approach to StarAI, *probabilistic logic programming*.

The key idea behind probabilistic logic programming is the *distribution semantics*, introduced by T. Sato in 1995. The distribution semantics neatly divides a probabilistic logic program into a simple list of probabilistic facts and an ordinary logic program which takes those probabilistic facts as input to compute more complex predicates. More precisely, a probabilistic logic program consists of a list of assertions, each annotated with a probability between 0 and 1, and a logic program whose extensional vocabulary coincides with the signature of those facts. We illustrate this idea with a common toy example from the literature:

The probabilistic logic program *Smokers and Friends* consists of the probabilistic facts

```
0.2 :: befriends(X,Y).
0.5 :: influences(X,Y).
0.3 :: stress(X).
```

and the rules

```
friends(X) :- befriends(X,Y).
friends(X) :- befriends(Y,X).
smokes(X) :- stress(X).
smokes(X) :- friends(X,Y), smokes(Y), influences(Y,X).
```

The semantics of this program reads as follows: For every domain entity (referred to as a person in the following), there is a 30% chance that this person is stressed (that is, `stress(person)` is true). Similarly, for every pair of persons, there is a 20% chance that one person befriends the other, and a 50% chance that one person influences the other. All

these random choices are made independently of one another. After these choices have been made, the rules of the program are brought to bear. In this case, the binary relation `friends` is evaluated as the symmetric closure of `befriends`, and the `smokes` relation is computed as follows: Firstly, every stressed individual smokes. Then, `smokes` predicate is spread recursively along `friends` influencing each other. Once this process has been completed and no more smokers can be added this way, the program terminates.

On the first glance, this formalism seems very restrictive, since it confines probability entirely to independent choices on factual assertions. However, the ingenuity of the distribution semantics is that probabilistic rules can be expressed by simply adding additional auxiliary predicates:

The probabilistic rule

```
0.5 :: smokes(X) :- friends(X,Y), smokes(Y).
```

is equivalent to the clauses

```
smokes(X) :- friends(X,Y), smokes(Y), influences(Y,X).
0.5 :: influences(X,Y).
```

of the program in the first example.

In fact, it has been shown that if only ground (propositional) clauses are considered, acyclic probabilistic logic programs have the same expressiveness as Bayesian networks, while general probabilistic logic programs go beyond that by adding recursion.

However, the real strength of probabilistic logic programs (and StarAI in general) lies in the ability to specify rules and facts with variables, and therefore to specify a particular probability distribution for any domain under consideration. From this viewpoint the probabilistic logic program associates with every domain of people a probability distribution over possible L -structures on this domain, where L contains all the predicates mentioned in the program.

This framework can be extended to consider as input not just plain domains but structures in an extensional vocabulary $E \subseteq L$; for instance, the network of friendships in the examples here could be taken as input rather than generated probabilistically. In that case, every input structure M gives rise under a probabilistic logic program to a probability distribution over all possible L -structures extending M . For a more detailed contemporary treatment of probabilistic logic programming under the distribution semantics, see the comprehensive textbook by F. Riguzzi (2020:Foundations of Probabilistic Logic Programming, River Publishers).

The tasks for any probabilistic logic program system can be divided into two main headers: inference and learning.

Among inference tasks, the most prominent are marginal inference, where the probabilities of a certain property are computed, and maximum-a-posteriori (MAP) inference, where the most likely configuration to arise from a probabilistic logic program is computed. In the example above, a marginal query could compute the probability of a certain individual smoking. MAP inference could answer the question: What is the most likely smoking behaviour of this group of individuals given some observations?

Unfortunately, the naive approach to grounding, that is, substituting the domain individuals into the probabilistic logic program, and then performing inference on the ground program does not scale well to large datasets. Therefore, current research is much concerned with lifted inference, which is performed to varying extent on the level of the original proba-

bilistic logic program (with variables) rather than on the large grounded program.

The problem of learning probabilistic logic programs from data is posed in different settings. In parameter learning, the rules of the program are fixed and an optimal set of probabilities for those rules are sought. In structure learning, the rules themselves are not given either and have to be found by the learning algorithm. Structure learning of probabilistic logic programs is generally considered a difficult problem, since it subsumes and significantly extends the problem of learning deterministic logic programs, known as Inductive Logic Programming. Indeed, early structure learning algorithms approached structure-learning in two parts, applying first an Inductive Logic Programming system to learn the rules and then a parameter learning algorithm for the probabilities. However, better results have recently been obtained by interleaving the two parts.

I would like to close by highlighting two particularly active and well-maintained systems, both of which have an easy-to-use online interface and tutorial.

ProbLog 2, an implementation of the ProbLog language maintained by the KU Leuven group available at <https://dtai.cs.kuleuven.be/problog/>, is a powerful Python-based system with many features. Its underlying language, ProbLog, has been designed particularly with usability and straightforward syntax in mind, and it has been used extensively in bioinformatics. A variety of academic and real-world applications are listed at <https://dtai.cs.kuleuven.be/problog/applications.html>.

cplint, a SWI-Prolog-based system supporting various languages and maintained by the University of Ferrara, is available at <https://cplint.ml.unife.it/>. It is particularly rich in configurations and supported algorithms, including state-of-the-art methods for structure learning and lifted inference (neither of which are supported by ProbLog 2 at present). On the other hand, it is not quite as straightforward to use as ProbLog 2.

FELIX WEITKÄMPER
LMU, Munich

Mathematical Philosophy

The proofs that appear in mathematics journals typically don't look much like the idealized proofs of philosophers' fancy. Ideal proofs are perfectly rigorous, complete, error-free, and known to be correct by any readers with the expertise to understand them. The proofs mathematicians actually produce are informal, sketchy and full of gaps, and usually not vetted in exhaustive detail even by their assigned referees.

So it may be philosophically scandalous, but shouldn't be too surprising, that published proofs contain lots of mistakes. Many of these are almost certainly never noticed. (And even when they are, they're rarely corrected in print—retractions and corrigenda aren't much of a thing in math.) In fact, some very famous and important proofs contain a great many known mistakes, and perhaps many more unknown ones.



The classification of finite simple groups is a case in point. Considered one of the crowning results of modern math, the proof is known for its extreme length, complexity and collaborative nature, comprising thousands of pages published over decades by dozens of mathematicians. What's less well known is that the proof is full of errors. These mistakes come in varying degrees of severity, the worst known so far being a gap in the classification of quasithin groups which required a new 1000-page monograph published in 2004. (The proof was originally announced to be complete in 1981.) I learned this and lots more about the classification theorem from Habgood-Coote and Tanswell's new paper "[Group Knowledge and Mathematical Collaboration: A Philosophical Examination of the Classification of Finite Simple Groups](#)" (forthcoming, *Episteme*).

In spite of these issues, the classification theorem is still widely believed to be true, and even to be *proven* true. But this can't be on the basis of any known completely correct proof. Mathematicians are still busy simplifying and fixing the original arguments—a process expected to yield a 3000–4000 page reworked proof once finished, which will itself probably contain its fair share of mistakes.

Suppose the group theory community is right to regard the classification theorem as proved. What would justify this verdict, if not the existence of an actually correct argument? The mathematicians involved with the classification theorem proof often appeal to its *fixability*. As Habgood-Coote and Tanswell understand the notion, "A proof is fixable when all of its errors could easily be corrected by experts within the relevant mathematical community, without needing to do any substantial new maths" (15). Group theorists' attitudes toward the classification theorem suggest they're willing to accept fixable proofs as successful proofs.

De Toffoli's "[Groundwork for a Fallibilist Account of Mathematics](#)" (2020, *Philosophical Quarterly*) makes similar points about mathematicians' tolerance for correctable errors. "Consider Gauss's original argument for the fundamental theorem of algebra; it contained minor mistakes. Strictly speaking it was not a proof [in the idealized, error-free sense]. However, the corrected version is a proof [in the error-free sense] and it is generally considered as being 'essentially another version' of the original one" (12-13). Mathematicians are willing to call Gauss's original argument a genuine proof because its problems were fixable.

I think this is an interesting and important insight. I wonder, though, whether Habgood-Coote and Tanswell's definition of fixability is exactly right. Here's a case. Shortly before its unfortunate and complete extinction, an alien civilization whose mathematics is a thousand years beyond ours broadcasts its newly discovered proof of the Riemann Hypothesis to the universe. The proof, which uses cutting-edge tools of 31st-century number theory, is not quite right—it contains a handful of mistakes that the aliens could have fixed, had they lived long enough to try. Alas, they didn't, and now there's no one left in the universe who understands 31st-century math. Is the aliens' argument fixable? (Assume we managed to receive and translate the transmission, so the work still exists.)

My intuition is that the argument is fixable, and indeed that it deserves to be called a genuine proof. But it's not clear that Habgood-Coote and Tanswell's definition predicts this result. After the aliens die off, there are no experts around capable of fixing the mistakes. This seems to incorrectly entail that the

proof isn't fixable. But there's perhaps another interpretation of HCT. When they write "could easily be corrected by experts within the relevant mathematical community", this might mean "could easily be corrected by the experts *if there were any*" rather than "could easily be corrected *by some existing group of experts*".

On this reading, the proof of the Riemann Hypothesis is fixable, because the aliens could have corrected it if they'd survived. But this notion of fixability may be too permissive. Consider the proof Fermat thought he had of his Last Theorem. Presumably Fermat's argument was far off the mark. But number theorists of the 21st century could correct it without having to invent any new math. This seems to fit the second, broader definition of fixability, but I don't think Fermat's proof should be called fixable just because we could now identify and circumvent its mistakes.

In any event, cases like the proof of the classification theorem are strange and surprising. I think we should be puzzled by them. How can a proof full of mistakes count as a success, even if the mistakes are never noticed or fixed? Does it matter that we think we *will* fix them sooner or later, as with the proof of the classification theorem? But then how is it reasonable to believe that mistakes can be easily fixed when we don't yet know what they are?

The moral of the story is a familiar one: scratch a supposedly simple mathematical practice and you'll find lots of fascinating epistemology underneath.

WILLIAM D'ALESSANDRO
MCMP, Munich

EVENTS

NOVEMBER

MoSCG: Debate: The Merits of Semi-Classical Gravity, online, 2 November.

CT&D: on Belief in Conspiracy Theories and Delusions, online, 3 November.

EPIPLUR: Epistemic Pluralism, online, 4 November.

BPOs: Bergen Philosophy of Science Workshop, online, 4–5 November.

AIXIA: Conference of the Italian Association for Artificial Intelligence, Milan, 29–30 November.

DECEMBER

TCiPoS: Thick Concepts in the Philosophy of Science, Hannover, Germany, 3–4 December.

IPS: Interdisciplinarity and Philosophy of Science, online, 6–8 December.

COURSES AND PROGRAMMES

Courses

CIe: Computability in Europe 2021: Connecting with Computability Tutorials, 5–9 July.

Programmes

MA IN REASONING, ANALYSIS AND MODELLING: University of Milan, Italy.

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

MASTER PROGRAMME: MA in Pure and Applied Logic, University of Barcelona.

DOCTORAL PROGRAMME IN PHILOSOPHY: Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.

DOCTORAL PROGRAMME IN PHILOSOPHY: Department of Philosophy, University of Milan, Italy.

LOGICS: Joint doctoral program on Logical Methods in Computer Science, TU Wien, TU Graz, and JKU Linz, Austria.

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

MASTER PROGRAMME: in Statistics, University College Dublin.

LoPhiSC: Master in Logic, Philosophy of Science and 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.

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 PROGRAMMES: in Philosophy of Science, University of Leeds.

MA IN LOGIC AND PHILOSOPHY OF SCIENCE: Faculty of Philosophy, Philosophy of Science and Study of Religion, LMU Munich.

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, SCIENCE AND SOCIETY: TiLPS, 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 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.

MSc IN COGNITIVE & DECISION SCIENCES: Psychology, University College London.

MSc IN COGNITIVE SYSTEMS: Language, Learning, and Reasoning, University of Potsdam.

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 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 Sebastián).

OPEN MIND: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.

RESEARCH MASTER IN PHILOSOPHY AND ECONOMICS: Erasmus University Rotterdam, The Netherlands.

JOBS AND STUDENTSHIPS

Studentships

DOCTORAL PROGRAMME IN PHILOSOPHY: Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.

LOGICS: Joint doctoral program on Logical Methods in Computer Science, TU Wien, TU Graz, and JKU Linz, Austria.

Jobs

POSTDOC: in Causal Discovery, RIKEN, Japan, deadline open until filled.

ASSISTANT PROFESSOR: in Philosophy of Science, University of Toronto, deadline 1 November.

ASSOCIATE OR FULL PROFESSOR: in Philosophy of Science, University of California, Los Angeles, deadline 12 November.

SENIOR LECTURER: in Theoretical Philosophy, University of Gothenburg, deadline 15 November.

PROFESSORSHIP: in Philosophy of Physics, Stanford University, deadline 19 November.

