
THE REASONER

VOLUME 18, NUMBER 1
JANUARY 2024

thereasoner.org
ISSN 1757-0522

CONTENTS

Guest Editorial	1
Features	1
Dissemination Corner	5
News	6
What's Hot in ...	7
Courses and Programmes	8

GUEST EDITORIAL

Dear reasoners,
we all inhabit a publish or perish environment. In recent times, I've been thinking a fair bit about the publication process and the decision making of journals. On the one side of the proverbial coin I joined [Logique et Analyse](#) as an associate editor, while on the other side I had some of my manuscripts desk-rejected. It occurred to me that there really should be a way to reduce disappointments of authors and editors. So, I set out to recruit a group of seasoned editors to talk about this. Fortunately, Wendy Parker, Kenny Easwaran and Thomas Reydon accepted my invitation for a discussion.

I hope that this discussion will provide some insights on the editorial process from "the other side". One thing to keep in

mind is that all three editors are philosophers. Readers from other disciplines should take what is said with a bit of salt. Although, I believe that many of the points they make also apply in other disciplines.

JÜRGEN LANDES
Logic, Uncertainty, Computation and Information Group,
University of Milan

FEATURES

Interview with Wendy Parker, Kenny Easwaran, and Thomas Reydon	8
---	---

JÜRGEN LANDES: **How and why did you become an editor?**

THOMAS REYDON: It was all a coincidence. While I was still a PhD student I was asked to become Managing Editor of the journal that was edited at the institute where I was doing my PhD. The journal was *Acta Biotheoretica*, a smaller journal that publishes in mathematical biology and philosophy of biology. (The journal represents a Dutch tradition of integrating theoretical/mathematical and conceptual work.) I liked doing that and I saw that it was an important way to contribute to the functioning of the academic community. If no one manages journals, after all, no one gets to publish. So, when I was asked later to become Associate Editor of *Acta Biotheoretica*, I did not hesitate long. I stayed Associate Editor for about ten years and when I was ready to leave (as editors shouldn't stay in place for too long, and ten years probably already was too long), I was asked to become one of the Editors in Chief of the *Journal for General Philosophy of Science*. I have been doing that now for almost seven years and I still like doing it. In the meantime, I also started editing a book series with two colleagues and have been doing that for thirteen years.

KENNY EASWARAN: A few weeks after I submitted my tenure file, I got an e-mail inviting me to be an editor. I don't recall which journal invited me first, but I believe that by the time my tenure was approved, I was an editor at *Episteme, Ergo*, and



Philosophy Compass. It seemed natural to me at the time to take up the role, because editing journals is an important sort of service work, and it is important to shield junior academics from it, so taking it up as I got tenure seemed appropriate. I have also throughout my career found that I have a greater interest in reading and commenting on work than many other people, so I figured it was worthwhile for me to give back to the community in this way more than through organizing conferences or running institutes or the like. In the years since then, I have given up my role at *Philosophy Compass*, taken on and later ended a role at the *Journal of Philosophical Logic*, and recently taken on a role at *Philosophy of Science*. One thing I have noted is that there is some difference in the roles I have taken on at different journals, and I have found myself more suited to some than others. At most of these journals, I have had a role like “Associate Editor”, in which I receive submissions from an Editor-in-Chief (who decides which Associate Editor might be best suited to the individual paper). But at *JPL* I was co-Editor-in-Chief, and at *Philosophy Compass* the role involved soliciting authors with planned topics rather than dealing with submitted papers, and I found that I was much less suited to either of those roles.

WENDY PARKER: In 2016, the *British Journal for the Philosophy of Science* (BJPS) was looking for a Co-Editor-in-Chief. At the time, I was serving as the Program Chair for PSA16, which involved, among other things, setting up peer review of submissions, selecting papers for presentation and publication, and shepherding papers through the publication process. I found that I genuinely enjoyed these activities. Still, I wouldn't have applied for the BJPS post without the encouragement of other people; it seemed like a job for someone more senior than myself. But I'm really glad I applied. I've been at BJPS since 2017, first with Steven French as the other Co-Editor-in-Chief and now with Rob Rupert, and I enjoy it.



LJ: What are typical avoidable mistakes authors make?

TR: Sending papers to journals too quickly. Writing philosophy involves revising, rethinking, revising, rethinking, and revising once more before you have a draft that is halfway OK. There is a lot of pressure on people to publish, in particular on early-career scholars, which I think sometimes leads people to just send their paper out rather than doing the necessary re-reading, re-thinking and polishing. This is understandable, of course, and I'm not sure what to do about it. Another avoidable mistake is submitting papers that are too heavy or too light on the science. I'm editing a journal in philosophy of science. Some papers we receive are very science-heavy and hardly connected to all the relevant debates in philosophy of science and philosophy more broadly, whereas other papers are very philosophy-heavy and hardly connected to the relevant science. It's always a bit painful to see work by professional philosophers who write about a scientific topic about which they don't actually know enough, or scientists who don't know their way in the relevant philosophical debates. I see this as something to work on for the future: good philosophy of scientists needs more collaborations between philosophers and scientists, who can bring their own expertise to what are essentially interdisciplinary topics. Fortunately, we see more and more such work.

WP: I agree with Thomas's observations. I would add to the list the mistake of trying to do too much or too little in the paper. Sometimes a paper tries to cover three papers' worth of material, with the result that nothing is worked out in enough detail. Other times a paper is really nicely done but makes a very small point. The aim should be for a paper that makes a significant contribution and does so in a clear and worked-out way. Another common mistake is failing to engage sufficiently with existing literature. This is not only a scholarship issue but will irritate referees to no end if it's their own work that hasn't been cited. Though it might sound obvious, conducting a good literature search is important.

KE: The most common mistake I find is that authors are too wedded to their theory *X*, and they summarize their argument in the abstract and introduction as “I provide a novel argument for why theory *Y* is wrong”. Much of the time, I find that authors could double their audience if they instead phrased the abstract and introduction as “I show how the assumption of *X* leads to the rejection of *Y*” - with this format, they don't just give the *X* people one more reason to reject *Y*, but they also give the *Y* people one more angle on understanding what parts of *X* they object to. It can be painful to allow that someone might not be moved by your self-evident premises, but it is both practically good for you to double the number of readers your paper gets, and intellectually good for the community to present a valuable new argument in a way that more people will be willing to engage with.

LJ: How can authors improve their manuscripts?

KE: There are many different things that can be done to improve a manuscript, but almost all of them benefit from getting more eyes to look at it. When multiple people have looked at your paper, you can get a better sense of whether some passage is actually confusing, or one reader was just confused because they're coming from a different perspective. You can see when everyone tells you to say more about something, or whether half of them say to say more and the other half say to cut it. I can't tell you how to improve the paper until I've read it, but when more people have read it, you have more of a sense of how people who don't already have the view in their head are and aren't successfully understanding what you've said. One thing to keep in mind - if the paper ever gets published, no one else will be obliged to finish reading it. Think about who are the people that you *want* to read this paper, and try to write it in a way that will encourage them to read the parts you want them to see. In some cases, that means writing it to get them gripped from beginning to end, but in at least as many cases that means giving them the signposts to see that, for instance, section 5 says something weird and new about their pet theory and they can jump right there, and then get so excited that they flip back and read the whole thing. It's fine if 90% of the people who pick up the paper put it right back down after the first paragraph, as long as the ones you care about can find something that they want to hear more about.

TR: In my view an important thing is to be aware that a philosophical paper is not about what you say but also about how you say it. Philosophy to some extent is always an expression of a personal way of seeing things, and authors should use their own voice to express their thoughts. So: find your own voice (which may change as you go along) and be aware that a paper basically is you speaking to someone else (and because of that, don't use AI tools to try to perfect the language of your paper).

WP: Get various other people to read the manuscript. If they

can't understand what you're saying, or find it confusing, then referees are likely to have trouble too.

LJ: Are referees always right?

WP: No. Referees are human and thus subject to all the usual human foibles. But referee comments do need to be taken seriously, if the paper is to make it through the review process and be accepted for publication. That doesn't mean one has to do everything that a referee says should be done. But every referee request/comment should be addressed in some way, even if just by explaining why one has chosen not to do what the referee suggested.

KE: They definitely are not always right. But they are always a sign of what one person who read this paper thought. It's very rarely appropriate to completely ignore a remark by a referee - but it's sometimes appropriate to make the change the referee suggested, and sometimes appropriate to make a different change to prevent the misunderstanding the referee made. Remember that the referee was obliged to read the whole thing once they agreed - that makes them a somewhat non-representative reader, since they may have decided halfway through that they don't like it, while in the wild, only people who still like it will read the second half. But you can still get a sense of what turns off one particular audience by seeing their reaction, and hopefully you can use that to get a sense of what your intended audience will think.

TR: No, definitely not. The best indication of this is the considerable number of papers for which reviewers' judgments diverge widely, with one reviewer recommending acceptance, perhaps with some minor revisions, while the other reviewer strongly recommends rejection. Diverging verdicts occur so often that you can't consider it an exception to the rule anymore. For authors I think the message should be to take reviewers' comments seriously, while at the same time having a critical eye for possible mistakes and misunderstandings on the reviewer's part. And there often is something positive to gain here: if a reviewer puts forward an (in your view) unfair criticism, it may be that your presentation wasn't sufficiently clear. So, when you disagree with a reviewer I think a good approach is to first look whether your writing caused a misunderstanding and you can change that. If not, look for a profound difference of perspective, or something else that can clarify the criticism. It is always possible to explain in your replies to the reviewers' comments why you disagree without making changes in the paper, or alternatively you can explain the disagreement in a footnote if you feel that's appropriate. The main message is: you don't have to change things in your paper for every single comment a reviewer makes, you can also simply explain why you disagree and didn't change anything in response to a comment.

LJ: What do you enjoy most about editing?

KE: The thing I enjoy most about editing is seeing all the weird and wonderful variety of ideas that people are coming up with WP: Having an inside look into the entire publishing process helped me understand the whole thing a lot better. As a philosopher of science, I am interested in how knowledge is produced and journal editing is one of the steps in the production process. Besides that, editing is my way of performing



service to the community and helping authors get their work out in print just is a satisfying thing to do. TR: I enjoy seeing the many different papers that are submitted. I also enjoy seeing the finished product - a print copy of an issue of the journal, with all of the great papers that it contains.

LJ: What do you like least about editing?

WP: Finding referees. People often decline invitations to review, and the process of identifying suitable referees is unsystematic. I wish we could improve the process by which referees are identified and selected, to make it a little more systematic and inclusive. This is a problem not just for philosophy of science but across academic fields.

KE: Trying to recruit referees, and trying to get their comments in on a timely basis. I know that there are many people out there who would be great referees, who just aren't getting asked to do it, but that doesn't help me find them. But as long as everyone who turns me down makes multiple suggestions of new people to ask (particularly if they're young people in the specific field of the paper, who I might not be personally familiar with), I can get past this.

TR: That it is hardly appreciated. Editing a journal is a lot of work, and I like doing it, but there is not much institutional appreciation for it. I am a professor in Germany and in the German system journal editing does not count for anything career-wise (nor does reviewing papers, for that matter). Editing and managing a journal, and reviewing papers, are the backbone of academic publishing, so it seems to me that it should be much more visible as service to the profession and be valued when it comes to hiring and promotions. The publisher shows appreciation and so does the community (once in a while), but my issue is with the academic tenure and promotion systems in many countries.

LJ: What's the favourite paper you edited?

TR: I couldn't say. Papers to me are noteworthy when you clearly see them improve due to the reviewers' comments, or when they happen to be on something that I'm currently working on, or when they have a high literary quality. Beyond that I'm not in favor of highlighting papers. Some journals have an "Editor's Choice" for every issue, some journals have a "most cited" tab on their website. And then there is the *Philosopher's Annual*, which self-describes as aiming "to select the ten best articles published in philosophy each year." I tend to think that such highlighting of articles is harmful to the field and the people working in it. It adds to competition in the field, a field that already is extremely competitive and that, I believe, doesn't benefit from competition. It's bigging up a few at the cost of the rest. (And then there is the implicit assumption behind the *Philosopher's Annual* that the best philosophy is done in English.) So, no favorites here.

WP: I also don't have a favourite.

KE: There are several papers I've edited that I think are doing great work on topics that are close to my own research. But one that I particularly enjoyed working on, despite not being on a topic I work on, was "On the Epistemological Similarities between Market Liberalism and Standpoint Theory", by Raimund Pils and Philipp Schoenegger, in *Episteme*. It takes two epistemological theories that are often associated with opposing political views, and argues that they have some important theoretical connections, and can each show some greater importance for the other.

LJ: If/how does the refereeing process improve manuscripts?

TR: If things go well, reviewers provide a balanced assessment of the strengths and weaknesses of a paper that gives authors a clearer view of where they are on to something good and where they have to do more work. Reviewers of course have the task of making recommendations to editors about whether a paper can be published and, if not, whether it is worthwhile to pursue further. But reviewers also have the task of providing comments for authors that could help them improve their paper. Many reviewers do this well by providing detailed criticisms of arguments that challenge authors to develop their arguments further, questions for more clarification, suggestions for relevant literature, and even lists of typos and other minor corrections. Reviewing is not just gatekeeping, but also (and perhaps more importantly) providing constructive feedback to authors. That said, occasionally authors actively use reviewers as their helpers: sometimes authors submit somewhat rough drafts that are still in an early stage and don't actually seem to aim for publication yet, but rather hope for lots of helpful comments. I think this is bad practice: authors should only submit papers that they feel are ready for publication and not use reviewers and editors to help them get a rough draft to publication stage.

KE: The biggest way the refereeing process can improve manuscripts is by helping authors understand what is actually most essential and valuable to the community about their ideas. Sometimes the ideas need to be presented in a different order, sometimes some of the things the author is thinking about should be saved for the next paper, and sometimes there's something that the author thought was obvious that actually turns out to be the most important section to explain. Unfortunately, there are also sometimes cases where the refereeing process makes manuscripts worse, by making them grow epicycles, and having footnotes that are arguments with the referee.

WP: There are lots of ways that the refereeing process improves manuscripts. A particularly important one, in my opinion, is that referees can point out objections that the author simply hasn't thought of. Often it's not that hard to address the objection, and then readers of the published version who think of the objection will hopefully be more satisfied (or less dissatisfied) with the paper/position. Sometimes the objection is not easy to address, because there's a real problem; in that case, the author may be glad that they didn't publish the paper without doing a lot more work.

LJ: **Words of encouragement to authors of rejected manuscripts:**

KE: First, double-check whether you actually got a rejection. I was once talking to a grad student who got a "revise and resubmit" and thought this was a negative judgment, not realizing that this is actually just about the most *positive* judgment that an initial submission can receive. Note also that some journals have "reject with encouragement to resubmit" as a possibility for more cautious "revise and resubmit" judgments. You may still decide that the comments are negative enough that you'd rather take the paper somewhere else - but it's still worth identifying whether any of the comments help you see what misunderstandings to try to forestall in future revisions. It's helpful to have a friend or mentor who can give you the caring pep talk while going through the comments, to help you see which are worth attending to for your next submission, and which are worth laughing off.

TR: Keep at it, keep working on the paper, it will get accepted at some point. I think everyone has had their fill of -

sometimes harsh and unfair - rejections in their career (I certainly have). Don't expect it to get better when your career advances, know that the rejections and (sometimes unfair) criticisms will keep coming in, and just try to get used to it. It's part of the normal process of getting your work out there.

WP: *Everyone* has papers rejected sometimes. Sometimes it's for good reason, and sometimes it's mainly bad luck. Even if you think it's the latter - you got a grumpy and confused referee - it's not necessarily a good idea to immediately submit the paper elsewhere, without making any changes. You might get some of the same referees again, and if they see that nothing has been done to address their concerns, they're likely to be annoyed and recommend rejection again. Better to consider the comments you did get on the rejected paper and see what merits addressing. But then do submit it again - don't take a rejection to mean that a paper should just be scrapped.

LJ: **Your thoughts on blinding the reviewing process (open, single, double or triple blind):**

TR: My journal uses double blind reviewing, which I think works well for us. I'm not in favor of open reviewing. I think it's important that reviewers can speak their mind without having to consider that the entire community will be able to read their reports.



Reviewers have to be able to express concerns about a paper that pertain to whether it is publishable, and such concerns aren't necessarily interesting to the entire community. Review reports are not published comments: the latter have the role of advancing the debate on a particular topic and a commentator can contrast their view to the author's; the former have the role of assessing whether a paper makes a point that is worthwhile to make available to the community for discussion and further development. Treating review reports as commentaries is a mistake, it seems to me.

KE: I've never done work for a journal with an open reviewing process, but I've refereed for mathematical journals that are single blind (so the referee and editor both see the author's name, but the author doesn't know the referee's name), and both refereed and edited for journals that are double blind (where the editor knows the author's name but the referee doesn't) or triple blind (where the editor also doesn't know the author's name). I find that seeing a name, even if it's one I don't know, makes it much harder for me to think about the paper, and easier to think about the effect of a rejection or acceptance on the author. I don't think this is helpful.

WP: BJPS uses triple masking. So even the editors don't know whose paper it is when they issue a decision. While there are some drawbacks to this, on balance I am in favour of it, both for principled reasons and because of my own experience refereeing scientific papers where the author names are shared with referees; I think it's hard not to have that information about author identity affect one's judgment in ways that it shouldn't. Selfishly, triple-masking also makes my job easier: I don't dwell on the fact that we rejected so-and-so's paper, because I don't know that we rejected so-and-so's paper. I choose to only learn the identities of authors for accepted papers.

LJ: **Your thoughts on open access publication/article processing charges/subscriptions:**

KE: Although I have been an editor at several journals, and a referee at many, I have never been involved with the business

side of a journal. No matter how frugally a journal is run, someone needs to keep the submission/editing software running, and someone needs to host the server that stores the papers for public download. If authors use Microsoft Word, then we usually also want someone to do the typesetting (and sometimes copy-editing) to produce a document that looks final. I believe there are some math journals that get these costs down to zero by having hobbyists who enjoy maintaining submission software, and piggybacking off the arXiv for paper hosting, and relying on authors to write in LaTeX. But if we don't have that, then someone has to pay. Traditional journals make subscribing institutions pay, while many modern open access journals make authors pay. In the sciences, the cost is often split, with authors paying per color figure and readers paying the rest. *Ergo* and the *Stanford Encyclopedia of Philosophy* raise funds from other sources so that neither authors nor readers pay, but it's not clear how much funding is available to generalize this. Given that academic publishing is usually seen as a benefit to the reader, and sometimes as a benefit to the author, and given the inequality of resources available to potential readers and potential authors, it's not clear to me which model is best, but we should be trying them all out.

LJ: Resources for paper writing that you can recommend:

WP: The BJPS blog, *Auxiliary Hypotheses*, includes a section on Publishing Advice. A few of the posts include [How to read rejection: advice for the puzzled or peeved](#) by Rob Rupert, [What to do with a revise and resubmit](#) by Steven French and [How to anonymise your paper](#) by Beth Hannon. You can see additional posts [here](#).

LJ: Thank you all very much! Here are some closing thoughts I'd like to share with the readers of The Reasoner.

Given the competitive nature of the publish or perish environment, this interview will surely not be the last piece written on the editorial process. Neither is this interview the first approach to this topic. Mathew Mercuri (2020: Publishing your work: An editor's perspective, *Journal of Evaluation in Clinical Practice*, 3-6) -- for example -- would surely agree with most of what was said in this interview. Furthermore, we don't only disseminate research and gain academic Brownie points by producing written words, appearances and talks at conferences (and at job interviews!) matter, too. To those interested reading on presentations I recommend Irvine Loudon (1997: On Talks, *Medical History*, 1-5). Finally, not only are we academics writing papers, we ask our students to do so, too. For your students I can recommend Jim Pryor's [guide](#) and the [Harvard Writing Guide](#).

With all that being said, I hope you've learned something today. In particular, that you happily accept the next invitation to review a manuscript and that you will not be too disappointed, when the next rejection inevitably arrives in your inbox. Happy writing everyone and don't forget to hug an editor!

DISSEMINATION CORNER

BRIO

In my previous contribution to the Dissemination Corner (see The Reasoner, Volume 16, Number 6, September 2022), I delved into the transformative impacts of recent developments in artificial intelligence (AI), particularly focusing on the breakthroughs achieved by Large Language Models (LLMs) like

GPT-3. These models have not only amazed the AI community but also challenged our perception of machine intelligence, with some experts suggesting that some Machine Learning-powered softwares may have passed the Turing Test and exhibit forms of self-awareness.

Since that discussion, LLMs have been released worldwide, reaching a wider audience. A prime example is the release of ChatGPT, based on GPT-3, on November 30, 2022. This platform rapidly became one of the most popular web applications, achieving a record-breaking user base in an incredibly short time. As I write this article, GPT-4 has been made available to premium ChatGPT users. This iteration represents a significant leap forward, with enhanced capabilities such as web browsing, sophisticated data analysis, and an advanced understanding and generation of images, thanks to its seamless integration with DALL-E 3. To demonstrate its capabilities, I tasked GPT-4 with generating an image of a medical doctor, the result of which is displayed below:

This image shows both the prowess and the underlying challenges of AI. It portrays a male doctor in a white coat with a stethoscope, aligning with widespread human stereotypes. While inferring bias from a single instance is problematic, the occurrence hints at the potential of AI systems to inadvertently perpetuate societal stereotypes.



This underlines the critical need for ongoing scrutiny and evaluation of these technologies, especially considering their increasingly opaque algorithms.

Addressing these concerns, the BRIO team engaged in organizing the first BEWARE workshop last year, co-located with the 2022 edition of the AIXIA Conference in Udine (further information and the workshop's proceedings can be accessed at <https://sites.google.com/view/beware2022>). This inaugural event marked a pivotal step in establishing a diverse, multidisciplinary community comprising philosophers, logicians, computer scientists, psychologists, and other experts, united in a quest to navigate the emerging challenges presented by AI's evolution.

The workshop's inaugural edition was a melting pot of ideas and perspectives, providing a holistic view of the ethical dimensions of AI. This ranged from exploring practical tools and applications to probing into the theoretical foundations and methodologies of AI ethics. The positive reception and success of this first edition set the stage for its sequel, BEWARE-23, which took place in Rome on November 6, 2023 (for more details and forthcoming proceedings, visit <https://sites.google.com/view/beware2023>). BEWARE-23 continued to explore critical themes like Bias, Risk, Explainability, and the influence of Logic and Logic Programming in AI. The event was supported by notable organizations including AI Aware (<https://www.ai-aware.eu>), SIPEIA (<https://sipeia.it>), BRIO (<https://sites.unimi.it/brio/>), and in a new collaboration with the FAIR (Future Artificial Intelligence Research, <https://future-ai-research.it>) Foundation.

Teresa Scantamburlo from the University of Venice Ca' Foscari inaugurated the workshop, inspiring attendees with her insights into the role of moral exercises in AI. The presenta-

tions that followed spanned a diverse array of topics. These included strategies for unveiling opaque AI predictors, bridging the gap between human conceptual frameworks and computer vision, dissecting gender stereotypes in text, refining classification systems with innovative masking techniques, and examining the ethical and conceptual underpinnings of a human-centered symbiotic relationship with AI. Another area of focus was the logical basis of randomized computation, which poses unique challenges and opportunities in the realm of AI.

The BRIO team presented a tool developed with Alkemy (<https://www.alkemy.com>) designed for post-hoc analysis of classifiers. This tool, dubbed BRIOxAlkemy, aims to expose potential biases in AI systems. A video overview of the tool is available at https://www.youtube.com/watch?v=9W8_DTmXcts, and the open-source software can be downloaded from https://github.com/DLBD-Department/BRIO_x_Alkemy. Inspired by the TPTND Calculus (see <https://arxiv.org/abs/2206.12934>), this tool is especially useful for, e.g., assessing factors like age in credit default predictions. Users can input datasets, marking certain attributes as sensitive, choose a bias definition among the several available, and perform detailed analyses. The corresponding paper, soon to be published in the BEWARE-23 proceedings, offers an in-depth look at the tool's functionalities as well as a future development roadmap. As an open-source project, we encourage the community to engage with the tool, contributing through bug reports, suggestions for enhancements, or new features, which can be communicated directly to the BRIO or Alkemy teams.

This initiative is a part of BRIO's broader commitment to tackling the current challenges in Explainable and Trustworthy AI, as outlined at the beginning of this article. We invite you to continually visit our website (<https://sites.unimi.it/brio/>) for a comprehensive overview of the diverse research topics our team is exploring, our latest publications, and opportunities for collaboration. We are dedicated to advancing the field of AI, ensuring its ethical application and making it fair and accessible to all.

FABIO AURELIO D'ASARO
University of Verona

NEWS

“Replicability Crisis in Science?” 18-22 September 2023, Padua

The first summer school on *Replicability Crisis in Science?* was organized by the Department of Statistical Sciences in collaboration with the Philosophy Department. The primary objective of the summer school was to bring together graduate and undergraduate students from diverse backgrounds to engage in discussions about the replicability crisis facing science. While not all scientific disciplines are equally affected, the replicability crisis is considered a contributing factor to the erosion of science's authority among non-experts, making it a topic of interest to many.

The replicability crisis is a complex phenomenon with various facets including the dearth of replication studies, a sociological issue rooted in journals' reluctance to publish replications of established results, and researchers facing challenges in securing funding and allocating time for replication stud-

ies. When studies are replicated, social and medical sciences, in particular, encounter difficulties reproducing previously established results – a more intricate issue requiring expertise in statistics and philosophy to address effectively.

To tackle the issue of replicability, individuals with diverse backgrounds must collaborate, although effective communication is not always straightforward. To facilitate this, Professors Branden Fitelson (Northeastern University) and Giovanni Parmigiani (Dana Farber Institute) jointly lectured on the philosophical and statistical foundations of replicability during the first two days of the summer school. This aimed to establish a common ground for statisticians, philosophers, and scientists to communicate effectively. In addition, the afternoons were dedicated to an R laboratory led by Professors Filippo Garbarota and Gianmarco Altoè, providing hands-on experience with basic statistical tools, essential even for less quantitatively inclined students.

During the third and fourth days, students shifted their focus to a more field-oriented perspective. Researchers from diverse fields, ranging from physics to pharmaceuticals, assessed the current issues with replicability in their respective domains while exploring potential solutions. A key takeaway was the acknowledgment that there is no one-size-fits-all solution; each scientific discipline differs in terms of replicability challenges. While some issues are widespread and can be collectively addressed, others are inherently linked to the specific nature of each scientific field. Consequently, addressing the multidimensional challenge of replicability requires a diversified approach.

The University of Padua provided an ideal setting for this initiative due to its strong connections between statistics, social and psychological sciences, and philosophical studies. As a first-year endeavor, the summer school undeniably succeeded, thanks to the efforts of organizers Professors M. Carrara, F. Grigenti, A. Salvan, and B. Scarpa. Hopefully, this success will pave the way for its replication next year.

GIOVANNI DUCA
University of Milan

“Reasoning with Imperfect Information in Social Settings”, 26-28 October 2023, Pisa

The workshop “Reasoning with Imperfect Information in Social Settings”, held at Scuola Normale Superiore, provided an opportunity to tackle cutting-edge issues in logic and social epistemology, showing how to fruitfully use the toolbox of logic in understanding the opacity and intricacies of information flow. The workshop brought together established scholars and young researchers working in the broadly conceived areas of (i) proof theory, (ii) computational methods for social epistemology, (iii) deontic logics and (iv) philosophical logic.

(i) The keynote speaker Marcello D'Agostino (University of Milan) introduced natural deduction systems for logics which deal with inconsistent and incomplete information, providing a computationally tractable version of First Degree Entailment. The main insight of these logics is to isolate the actual information available to agents, while providing an upper bound on the number of pieces of virtual information agents may use in deductions.

(ii) The keynote speaker Gabriella Pigozzi (Paris Dauphine) presented a talk on computational methods for social epistemology. The talk dealt with *myside bias*, a relevant theme for

argumentative dynamics in scientific communities. *Myside bias* is the tendency to privilege arguments which promote one's views instead of the those who may challenge them. The effects of this phenomenon is analysed with computational models. The results showed both the argumentative dynamics generated by *myside bias* and the effect of consensus within the scientific community which contribute to mitigate it.

(iii) The keynote speaker Christian Strasser (RUB Bochum) touched upon themes connected to (i) and (iv). The talk (based on joint work with Kees van Berkel) focuses on deontic argumentation calculi (DAC), a formalism which can be used to produce explanations as to why a certain norm should or should not apply. The calculi provide a proof-theoretic platform suitable for a uniform presentation of unconstrained input-output logics (I/O logics) and of default logic. Therefore two key areas of non-monotonic reasoning are unified through a unique calculus which gives explicit reasons why a certain norm or default holds.

At the interface between (iii) and (iv), the keynote speaker Réka Markovich (University of Luxembourg) provided a theoretical framework in which to formalise relational normative positions such as claims and duties under the banner of the concept of epistemic right. Logics for epistemic rights combine deontic and epistemic logic so as to formalise the deontic status of epistemic states towards such normative positions (e.g. the right to privacy).

Contributions: Paolo Baldi; Lorenzo Casini & Jürgen Landes; Francesca Doneda; Chris Fermüller; Wojtek Jamroga; Felix Kopecky & Gregor Betz; Matteo Michellini & Javier Osorio; Carlo Proietti; Andrea Sabatini; Caterina Sisti & Lorenzo Rossi; Jan Sprenger, Lorenzo Rossi & Paul Egré. Organisation: Mario Piazza; Caterina Sisti; Matteo Tesi; Pietro Vigiani.

MATTEO TESI AND PIETRO VIGIANI
Scuola Normale Superiore

WHAT'S HOT IN ...

Statistical Relational AI

When you read this column, the year 2023 has already reached its close, and with it a great milestone in the history of reasoning research: Fifty years ago, Prolog was born, the first practical programming language based on automated reasoning. Half a century later, there are around a dozen well-maintained implementations, and Prolog is used in a great variety of projects in industry and academia. The Association of Logic Programming marked this “Year of Prolog” with a special collection on “Prolog: The Next 50 Years” (<https://doi.org/10.1007/978-3-031-35254-6>), a survey paper in the association's journal “Theory and Practice of Logic Programming” co-authored by many of the foremost maintainers and implementers of present-day Prolog systems (<https://doi.org/10.1017/S1471068422000102>), and a special session at the 2023 International Conference for Logic Programming (ICLP) dedicated to Prolog implementation and applications.

Reason enough for me to share some musings on the role contemporary Prolog can have in the context of statistical relational learning.

At this point you would normally expect a brief introduction to the Prolog language. However, at the ICLP we came to the

conclusion that a key priority is to improve the logic programming content on the online encyclopaedia Wikipedia. In this spirit, I invite all of you to head over to en.wikipedia.org/wiki/Prolog. Please give the page a read, and should anything remain unclear, remark that on the talk page, accessible from the top right of the main article. That would be of great help to our colleagues working on improving that page!

So, after this brief detour, let me introduce you to a Prolog project I am currently working on.

The context is structure learning for relation logistic regression networks, which are relational Bayesian networks in which the probability that a certain predicate holds for a tuple depends on whether certain clausal formulas involving other “parent” predicates hold in a structure.

In a 2021 journal paper, an efficient boosting algorithm was suggested for this task by Nandini Ramanan and others, which was implemented as a small part of the widely used Boost-SRL Java project. Recently, I decided that it would be worth reimplementing the relational logistic regression learner myself, giving us a nimble prototype system to play around with.

Immediately, the first advantage of programming in Prolog came into play: its almost legendary suitability for rapid prototyping. The general reasons for that are Prolog's brevity (Prolog are often even more concise than equivalent programs in functional languages) and the fact that Prolog is essentially untyped, requiring essentially no boilerplate code whatsoever. Sure enough, within a weekend, a working prototype was ready that faithfully reproduced the algorithm presented in the paper.

What took many thousand lines of Java code scattered among many different class files was encoded in just 230 lines of Prolog. Beyond its conciseness, the Prolog code was declarative and streamlined; any changes to, say, the regularisation parameters or search heuristics could be performed at obvious places without causing repercussions in other parts of the very small codebase.

There was another reason though, directly related to our problem domain: Most of the Java codebase was concerned with formalising various aspects of deduction, substitution and resolution theorem proving. In Prolog, all of this happens within the programming language interpreter. In essence, Prolog is a resolution theorem prover doubling up as a programming language. And all it needs to evaluate a rule on a set of examples is to call the rule as a goal, and see the examples as a program. This even allows for background knowledge including Prolog builtins, giving maximal flexibility and the efficiency of fifty years of implementation technique.

Even though the actual coverage check performed well, unsurprisingly overall was poor. Profiling quickly revealed that this due to a parameter estimation step, which involved some matrix calculations. Clearly, linear algebra in Prolog is a bad idea. So, I reimplemented this particular component in C, connecting it to Prolog via its foreign-language interface, and this performance bottleneck disappeared. The foreign-language integration of Prolog is another big plus for modern Prolog systems. Beyond an efficient C interface, the “best application paper” award at the 2023 ICLP went to the two-way Python bridge Janus for the XSB Prolog system. Within the past months, Janus has now also been ported to SWI-Prolog, arguably the most-used Prolog implementation in 2023. Python integration is absolutely essential for implementing a machine learning system, since all the downstream tasks such as evaluation and visualisation are available as Python libraries.

Lastly, there is one more feature of modern Prolog absolutely essential for this task: tabling. In our setting of boosted learning, it is necessary repeatedly to check which of the examples satisfy which clausal formula. In ordinary Prolog, each of these queries would be executed as a new resolution proof, spending most of its time recomputing earlier results. The classical solution is to use Prolog's dynamic database for manual caching, but that too is inefficient, as well as cumbersome and error-prone, tearing up the declarative structure. However, since it was introduced as part of XSB in the early 1990's, tabling provides a very elegant solution to this issue. In the background, answer tables are created and administered efficiently, and rather than destroying the declarative semantics, tabled Prolog programs are actually aligned more closely to their declarative meaning by the avoidance of Prolog's infamous tendency to get into infinite loops. This tool, which allows Prolog to be used as a highly efficient and versatile deductive database engine as well as a dynamic programming system, is what made the brief, clean implementation actually more performant than its imperative counterpart.

As "professional reasoners", we are perfectly placed to make use of Prolog's advantages. So, next time you have a programming project, consider whether the power of logic may also make that task a little less tedious.

FELIX WEITKÄMPER
Computer Science, LMU Munich

COURSES AND PROGRAMMES

Programmes

MA IN HUMAN CENTERED ARTIFICIAL INTELLIGENCE: University of Milan, Italy.

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: 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.

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 MIND, BRAIN AND LEARNING: Westminster Institute of Education, Oxford Brookes University.

MA IN PHILOSOPHY OF BIOLOGICAL AND COGNITIVE SCIENCES: Department of Philosophy, University of Bristol.

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 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.

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).

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

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

MA IN PHILOSOPHY: Dept. of Philosophy, California State University Long Beach.



IN A FIELD THAT'S BEEN AROUND FOR A WHILE, IT CAN BE HARD TO FIGURE OUT HOW MANY LEVELS OF REBUTTAL DEEP YOU ARE.