# The Reasoner

Volume 19, Number 1
January 2025

ISSN 1757-0522
Milano University Press
riviste.unimi.it/index.php/thereasoner


*Conveniently for everyone, it turns out that dark energy is produced by subterranean parasitoid wasps.*
https://xkcd.com/2986/

# The Reasoner 19(1), January 2025.

## Contents

# Introduction
## :: Hykel Hosni

**Abstract**
The Editorial Team welcomes readers to *The Reasoner*'s new home.

**Keywords**
Book reviews; News; Uncertain Reasoning; History of Logic and Reasoning; Medical Reasoning and Medical Methodology; Reasoning and AI; Causal Reasoning.

Dear Reasoners, it is my pleasure to welcome you to *The Reasoner*'s new home!

The Reasoner was established in 2007 as a printable and screen-readable outlet for research ideas that were more structured than blog posts, but leaner than the usual journal paper. For the benefit of younger readers, 2007 was *before* academics spent a varying amount of their day on social media. 2007 was also the time when publishers began defending themselves against open-access, which was clearly becoming a thing.

Nearly twenty years later, the editorial niche that The Reasoner originated and developed within seems more relevant than ever. As research communities are experimenting with new ways to


http://doi.org/10.54103/1757-0522/27602

fix what's broken with scientific publishing, The Reasoner aims to catalyse multidisciplinary scholarly discussion on reasoning-related topics, while serving as a tool for community-building.

To attain this twofold objective The Reasoner now solicits both *Research Articles* and *Features*.

Research articles typically do not exceed 3000 words and report either original results, or new perspectives on the existing literature. We welcome research articles on logic, methodology of science, and history of scientific reasoning – all broadly construed to include foundations and applications to society and technology. Research Articles are peer-reviewed through a double-blind process.

In addition, The Reasoner also publishes Features that are not peer-reviewed, but subject to the decision of the editorial board. We are particularly interested in Features that appeal to a broad and multidisciplinary audience. Features do not exceed 1000 words and are reasonably self-contained.

To encourage debate, we solicit *Post-publication open reviews and comments* to previously published articles and features.

Please see our Guidelines for Authors for full editorial details, and do not hesitate to get in touch with Esther Anna Corsi, our Editorial Manager, if you have any questions on the submission procedure.

Some of our Area Editors will now give you a feel for the topics and takes we hope to find in submissions. I would like to take this opportunity to thank all of them for being part of this. But please note that The Reasoner is always work-in-progress. Please get in touch if you are interested in editing an area which you see missing!

Finally, on behalf of the Editorial Team, I would like to thank all staff members at Milano University Press. Their commitment to the principles and best practices in Open Science made possible for The Reasoner to migrate to its new home, where we all hope it will thrive.

Hykel Hosni

Logic Uncertainty, Computation, and Information (LUCI) Lab,
University of Milan

## Book Reviews
## :: Paolo Baldi

I am very pleased to act as editor for the book reviews section in this new course of the Reasoner.

Scientific research keeps moving fast, with its products mostly presented in the format of journal articles and conference proceedings, not to speak of the much faster publication of preprints in online open access databases.

We are long past the era of the few all-encompassing "tractati" and "principia", and yet books are still indispensable and unsurpassed when it comes to the systematisation and education in the sciences.

I would thus like to invite all the readers and supporters of the Reasoner to contribute to this section, submitting reviews on books on reasoning-related topics. As we know, the analysis and modeling of reasoning is relevant to many different disciplines, and thus I would welcome texts belonging to as various fields as artificial intelligence, computer science, logic, linguistics, statistics, philosophy, cognitive science, economics, digital humanities, and even better a combination of two or more of those. More generally, I would appreciate reviews on books that appeal to both the technically-minded and the philosophical-minded readers of the Reasoner (as well as to their non-empty intersection!).

Apart from focusing on the ample range of topics touching on reasoning and the preference for interdisciplinarity, I would not put much further restrictions on the kind of books considered.

Reviews of different kinds of academic books are welcome, ranging from research monographs to textbooks, from handbooks to edited volumes on specific topics.

I would also encourage reviews of books from well-reputed authors that are intended for a general public, and yet present non-trivial scientific insights.

To name a few relatively recent publications, let me mention, for instance, *The Book of Why* by Judea Pearl and *The Road to Conscious Machines* by Michael Wooldridge. Both texts are indeed very timely in providing important conceptual tools for understanding the current AI spring, a topic naturally within the interests of the Reasoner. Furthermore, they are both presented in a accessible way, but at the same time they do not fear to dig deeper, and introduce formal and technical details, when needed.

You have a vast set of choices, and you may now have a good excuse to finally read carefully those books that have been waiting for so long on your desk. ...I am looking forward to receiving your contributions!

PAOLO BALDI

University of Salento

## News
## :: Lorenzo Casini

As News Editor of The Reasoner, I'd like to welcome all of our readers to this new and exciting season, which marks the transition of The Reasoner from an online gazette to an online journal!

Certain features of The Reasoner will inevitably change through this transition — most notably, The Reasoner will devote more space to submissions of original research articles, due to the journal now being indexed and each contribution getting immediately a doi.

Other features, however, which have contributed to the gazette's initial success, will be preserved. If it's not broken, why fix it? In particular, not only has the News Section of The Reasoner never lost the interest of our readership; in addition, we think that the continuity with which we have received spontaneous reports throughout the years witnesses the need for dissemination of the results of scientific events in the field, for which our gazette has always offered an important, if not unique, venue. In this light, we'll keep soliciting contributions of reports from organisers as well as accept spontaneous ones.

Moreover, our Editorial Team strongly believes that the dissemination task should also extend to advertising new doctoral dissertations, freshly funded scientific projects, novel scientific societies, etc. In other words, we welcome submissions that aim to give visibility to the results of young scholars, and thus pinpoint novel proposals, still unknown to the wider public, or to advertise

the future directions of academic research by giving a heads up to the community on what's going on where — which is key to foster the interactions between distant individuals or research groups that would otherwise operate in isolation.

In a world where information has never been so accessible but at the same time too vast to prune, there's a growing need for the right kind of filtering aid. We hope our News Section will be instrumental to fulfilling this need. Moreover, given that modern academia is becoming more and more specialized, there's also need for a tool that can trigger the cross-fertilisation between fields. In this respect, we hope that The Reasoner will keep helping to build bridges between academic disciplines which, albeit often institutionally separate, nonetheless share a common interest in reasoning-related topics.

Lorenzo Casini

University of Bologna

## Uncertain Reasoning
## :: Jürgen Landes

Our ability to make decisions based on good reasoning processes in an uncertain world have been a core part of what sets humans apart from all other known life forms. And yet, the now widely-accepted formalisation of probability theory is less than 100 years old. As Strevens put it in his book: "*our ancestors did not escape the sabertooths by thrashing them at roulette*" (Tychomancy, 2013, p. 217).

Not ever since, but especially since, Kolmogorov put forward his now-famous axioms of probability, probabilities – and models of uncertainty more widely – appear everywhere. Appearances of models of uncertain reasoning range from formal investigations in statistics, mathematics and philosophy over applications in the social sciences, physics and medicine to their didactics and history.

What then is the point of covering uncertainty and reasoning under uncertainty, which includes but is not restricted to probabilities, in *The Reasoner*? The point is precisely the vastness of the field, field**s** even!, that a) prevents the flow of ideas, b) directs and constricts one's interests to ever more-specialised niches and c) frustrates making general observations applying broadly across disciplines. *The Reasoner* aims to bring together researchers and ideas from various areas by publishing accessibly contributions that are of interest to a broad range of reasoners. Contributions in the *The Reasoner* also give readers a general sense of what is happening in a variety of areas and reduce the cost of entry when entering a new field.

When preparing a contribution to *The Reasoner* think of its readers. Wondering "how can I make my contribution fit neatly in one of the covered areas as identified by the editorial from the area editors?" or "how can I make my piece fit with the paradigm in my area?" is asking a *wrong question*. Do instead stress aspects that make your work relevant and stand out while remaining accessible to a wide range of readers. While we would all love to have our writings to be understood and treasured by everyone, but relevance for and appreciation by can only ever be achieved among a finite set of people.

Rather than telling you (at length) about my own research, which can of course be found at my website, I leave you with three of my favourite quotes:

> I only believe in statistics that I doctored myself. – (Winston S. Churchill)

> I'm not concerned about the rise of artificial intelligence; I'm worried about the decline of human intelligence. – (Unknown)

> To be uncertain is to be uncomfortable, but to be certain is to be ridiculous. – (Chinese proverb)

JÜRGEN LANDES

Munich Centre for Mathematical Philosophy,
LMU Munich

## Reasoning and AI
## :: Felix Weitkämper

Reasoning has been a core part of artificial intelligence from the very beginning. Since its first blossoming in the 1970s, the AI Winter of the 1980s and 1990s and rapid resurgence of the field since the 2000s, the scientific study of artificial intelligence has exploded into a vast area of research, with numerous publication venues stretching from completely comprehensive (such as the prestigious IJCAI conference series or the equally renowned *Artificial Intelligence* journal published by Elsevier) to highly specialised (with journals such as *Constraints* or the *Journal on Satisfiability, Boolean Modeling and Computation* serving their respective communities).

Why then would any researcher submit a communication on artificial intelligence to the *Reasoner*? First and foremost, The Reasoner has a rather unique audience. While the vast array of traditional artificial intelligence journals have a readership of computer scientists, The Reasoner aims to be interesting and accessible to everyone with an interest in Reasoning. Indeed, philosophers make up a sizeable proportion of the Reasoner's readership and editorship.

In my own research, I have oftentimes made use of both general ideas and concrete techniques and results developed in mathematical and philosophical logic to tackle issues in artificial intelligence that have resisted approaches more typical of computer scientists working in the field, and I feel that there is still very much to be gained from bringing over approaches that are salient to the

topic, but new to the respective field. I therefore hope that this cross-fertilisation between groups of researchers that often work on very similar topics from different angles, but with almost no institutionalised point of contact, will prove fruitful to everyone operating in the space of reasoning and artificial intelligence.

As a way of suggesting the many opportunities that arise for a dedicated space for reasoning in artificial intelligence, let us just briefly recall the wide usage that reasoning has found in contemporary artificial intelligence systems.

Firstly, there are the classical subfields of artificial intelligence that have crucially relied on reasoning techniques ever since they emerged. This encompasses assisted and automated theorem proving in all its variations. While arguably the mechanisation of mathematical reasoning was already the focal point Hilbert's ultimately doomed program of founding mathematics on a provably consistent and complete set of axioms, and Gödel's subsequent proof of its impossibility, positive progress on this semi-decidable task accelerated with Robinson's discovery of resolution theorem-proving in the 1960s. Since then, automated theorem proving led to the distinct discipline of logic programming, and its extension to various higher-order and non-classical logics of acute philosophical interest have captured the interest even of the general public, such as Benzmüller's work on Gödel's and Leibniz's argument for the existence of god, captured in modal logic.

Reasoning is also a crucial component of query answering systems. The close relationship between the predicate calculus and the relational algebra underlying the relational database model motivated plenty of work in finite model theory, and led to very fruitful investigations of the interplay between logic programming and databases. Even though they ultimately have very different

priorities, the joint roots of the fields of databases and artificial intelligence in reasoning can be very illuminating for a mutual transfer of results and techniques, as is happening now with probabilistic databases and statistical relational artificial intelligence.

Another traditional application area of logic in is automated planning and scheduling, which established itself very early on as a key part of artificial intelligence. Just to pick out one particular aspect, the vibrant field of epistemic planning, which takes explicit account of the knowledge of the agents involved in a situation, relies heavily on the advancement made in reasoning in epistemic logic.

Beyond those classical fields of symbolic artificial intelligence, which are almost unthinkable without formal reasoning, logic also interplays very well with the notions of learning and uncertainty that have dominated discourse on artificial intelligence since over the last 20 years.

The most direct combination of logical reasoning and learning is learning explicit logical connections from data. This has been studied in great depth by the field of inductive logic programming, which aims at learning logic programs that describe relationships in provided data. From its beginnings in the 1980s, there rapidly emerged both sophisticated theory and practical tools, some of which such as the learning systems FOIL and Aleph are still in use as research benchmarks today. After a brief lull in interest in the early 2000s, recent work in meta-interpretative learning has led to renewed interest in inductive logic programming approaches, particularly for classic program induction tasks from a small number of input-output pairs.

Combining logical rules with probabilities leads to statistical relational artificial intelligence, which developed since the 1990s from

roots in probabilistic logic programming and in template graphical models. There are now a wide range of formalisms and a unifying theory rooted in logical reasoning, supporting both probabilistic reasoning and structured machine learning (probabilistic inductive logic programming).

Of course, in addition to probabilities, research on reasoning has brought forth a range of other, qualitative approaches to incorporating uncertainty in logical frameworks such as belief functions or approaches based on fuzzy logic. A very promising, more recent application area of reasoning is neuro-symbolic artificial intelligence, which combines a symbolic reasoning layer with a subsymbolic component. This could directly enhance the reasoning capabilities of connectionist systems, or be used for a division of labour between a subsymbolic system responsible for low-level perception tasks and a reasoning layer responsible for higher tasks like planning or event recognition.

Hopefully, this very brief and woefully incomplete overview suffices to show the enduring relevance of reasoning in artificial intelligence. However, arguably the greatest potential for reasoning approaches in the current climate is the huge potential of reasoning for human-oriented, explainable artificial intelligence. The necessity to move beyond inscrutable black boxes to engage with human users on their own terms, in a language they can understand and with the certainty that this indeed captures the behaviour of the artificial agent, is becoming clearer with every new use case that artificial intelligence takes on. This opens up another application of reasoning: beyond applying reasoning *within* artificially intelligent systems, we can and should reason *about* artificial intelligent systems. Formal logic is uniquely placed to contribute to humanity's search for powerful, but safe and human-oriented artificial intelligence systems. However, for this challenge in par-

ticular, many disciplines have something to contribute, from computer science to cognitive science and philosophy. Hence, I am very excited to serve as the area editor for a journal that has the potential to bring all those communities together in a regular, but less formal setting, and look forward to your contributions, disseminating timely ideas from reasoning in artificial intelligence throughout the communities involved in reasoning research.

Felix Weitkämper

Chair of Programming Languages and Artificial Intelligence,
LMU Munich

## Medical Reasoning and Medical Methodology: :: Sarah Wieten

I am excited to join The Reasoner as the new subject editor in Medical Reasoning and Medical Methodology, I would like to re-invite papers related to reasoning, inference and scientific method as it pertains to medicine to submit here.

By way of introduction, my formal training is in philosophy of medicine, a relatively young discipline which is some ways takes it cues from philosophy of science. I have also worked as a clinical ethicist in a hospital, developing and applying work in bioethics to real time cases, and in a meta-research center, producing research on the institutions and methods that produce the contemporary science (particularly the health science) literature. These different positions throw different light on the area of medical reasoning, and I look forward to reading work related to all.

It is my hope that these works might come from a wide variety of disciplines (philosophy, medicine, public health, meta-research, economics, history, medical anthropology, sociology, psychology) and address a wide range of audiences (academic, practitioner, patient, public). Interdisciplinary work is very welcome.

However, of particular interest to me as an editor will be supporting work which pushes philosophy of medicine and work on medical reasoning in a somewhat new direction-towards integration with thinking in social and political philosophy. While work on the individual diagnostic encounter and shared decision making

about medical intervention remains important, investigations of the roles played by the social and political in medical reasoning, not just at the individual, but also at the population level, are extremely welcome.

Recent changes at The Reasoner mean that is this publication is indexed in DOAJ, WoS and Scopus and that each article has a doi, making the work more durable and accessible online. The Reasoner retains magazine-style interviews and features but is also an exciting option for placing academic work of interdisciplinary interest. Publishing in The Reasoner is **open access and free of charge, unlike much of the current academic publishing system.**

I would also like to close by pointing to a section after my own heart – The Reasoner Speculates. If you have an argument, correction, complaint, or question related to the goals of The Reasoner, but for whatever reason this is not a good fit for a traditional full-length article, consider it for The Reasoner Speculates. You can see more information about The Reasoner Speculates and other types of submission here.

I look forward to receiving your work in medical reasoning.

Sarah Wieten

University of Durham

## Causal Reasoning
## :: Jon Williamson

Causal reasoning has become a hugely important area of interdisciplinary research and we very much welcome submissions on any aspect of causal reasoning.

There are important descriptive questions about causal reasoning that would be of interest to readers of *The Reasoner*. For example, how do we infer causal relationships? How do we use causal relationships? How does causal cognition develop in childhood? The Oxford Handbook of Causal Reasoning provides a valuable introductory resource on topics related to causal cognition in humans. Descriptive work on causal enquiry in the sciences, medicine and law is also of key importance. Case studies of causal inferences in these fields can be very illuminating, for example, as Gillies (2019) demonstrates.

There are also important normative questions about causal reasoning. For example, how should we establish causal relationships? How should we evaluate causal claims? Can we automate causal discovery?

Normative approaches to causal reasoning can broadly be divided into three camps. Some theories make strong theoretical assumptions but offer methods that guarantee success in cases where those assumptions are satisfied. For example, Donald Rubin's potential outcomes framework and Judea Pearl's graphical causal modelling framework have both led to very fruitful streams of theoretical results.

A second group of theories make strong philosophical assump-

tions about the nature of causality and use these metaphysical presuppositions to motivate theories of causal enquiry. For example, probabilistic theories of causality can be used to motivate certain statistical approaches to causal inference, while agency or interventionist theories of causality steer one towards specific kinds of causal models, and dispositional or mechanistic theories are sometimes used to motivate methods such as process tracing. According to Andersen et al. (2019), such approaches are prone to 'philosophical bias', but may nevertheless be productive.

A third strategy is to try to avoid strong assumptions—both theoretical and metaphysical—in the hope of developing accounts of causal enquiry that are more general and less controversial, but nevertheless informative. The evidence-based evaluation movement can be thought of as situated in this camp, though it has by no means avoided controversy. Evidential Pluralism is an alternative approach in this camp that attempts to overcome some of the shortcomings of orthodox evidence-based methods.

What is clear is that while research on causal reasoning is enormously productive, the central questions are far from settled. Moreover, there is a great deal that each discipline can learn from others: some approaches dominate in certain fields for largely sociological reasons and there is value to be had in learning of good practice in other fields. *The Reasoner* is a perfect forum for this sort of cross-disciplinary engagement. Please submit features on research on causal reasoning that has caught your attention, or on your own research projects or research events in this area.

Jon Williamson

University of Manchester

# Non-Monotonic Bounded Reasoners

:: Paolo Baldi and Fabio D'Asaro

**Abstract**

We put forward a proposal for logics handling both non-monotonic reasoning and boundedly rational agents. We propose a hierarchy of depth-bounded non-monotonic logics. As the parameter $k$, which controls the level of reasoning depth, increases, non-monotonicity decreases while inferential power, in a classical sense, increases. Ultimately, these logics aim to provide non-monotonic approximations of classical logic. We present arguments and evidence supporting the adoption of this framework.

**Keywords**

Logic programming; depth-bounded boolean logics; non-monotonic reasoning; bounded rationality;

**Sumbitted** 11 December 2024
**Revised** 15 January 2025
**Accepted** 16 January 2025

**How to Cite**

Baldi, P., & D'Asaro, F. A. Non-monotonic bounded reasoners. The Reasoner, 19(1). https://doi.org/10.54103/1757-0522/27548

## Introduction

This paper focuses on two key issues that have attracted significant attention in non-classical logical investigations: the *non-monotonic* nature of cognitive agents and their *bounded* reasoning capabilities.

Cognitive agents typically infer, from their available information *more* than what is allowed by classical logic. In this "jumping to conclusion", the information that they infer should also be reasonably retracted upon the arrival of new information. This is the *non-monotonic* aspect of common sense reasoning and contrasts with classical (Tarskian) logic, where once a conclusion is reached, it remains valid regardless of the additional information. Logic programming in particular has proved to be a powerful computational tool to model this dynamic of non-monotonic reasoning.

On the other hand, cognitive agents also infer *less* than what is prescribed by classical logic. This is because classical logical inferences are computationally hard and real reasoning agents have *bounded* cognitive resources. Among the various logical formalisms introduced to model bounded reasoning agents, here we focus in particular on the approach based on *Depth-Bounded Boolean Logics* (DBBLs in the following), see D'Agostino, Gabbay, et al. (2024: *Depth-Bounded Reasoning: Volume I: Classical Propositional Logic*, College Publications). This approach provides a whole family of logical systems that are meant to approximate classical logic, with each system tractable, in computational terms.

In this paper we propose a formal framework for combining non-monotonicity with bounded reasoning, and more precisely the ideas of logic programming with those of DBBLs.

DBBLs classify logical inferences in a hierarchy, based on how much they employ information that goes beyond the actually available information.

As a typical example of such inferences involving additional information, consider reasoning by cases: to prove $\psi$, an agent may need to show that $\psi$ is the case both under supposition that $\phi$ is false, and under the supposition that $\phi$ is true. In this example, the supposition that $\phi$ is true or that $\phi$ is false does not count as information *actually possessed* by the agent, but is only *virtual*, i.e. instrumental for the derivation of $\psi$.

A logic $\vdash_0$ which does not manipulate any virtual information at all has been presented proof-theoretically in a system which is a mixture of classic natural deduction and a tableau method. A (unsigned) presentation of this system is provided in Figure 1, see D'Agostino, Gabbay, et al. (2024) for more details, and for a corresponding information-based semantics. This calculus consists of classical introduction and elimination rules for connectives (*intelim rules* for short). Intelim rules induce in a natural way the deducibility relation $\vdash_0$.

Completeness with respect to classical propositional logic is obtained by adding to $\vdash_0$ the rule of bivalence (RB), shown in Figure 2.

The rule of bivalence is taken to be the "difficult" part of an inference, both from a cognitive and computational point of view. Indeed, from the cognitive point of view, RB requires "guessing" the truth value of a formula $\phi$, beyond the information that is actually given. Computationally, it is only the unbounded use of RB which makes classical logic intractable.

$$\frac{a \quad b}{a \wedge b}\,\mathrm{I}\wedge \qquad \frac{\neg a}{\neg(a \wedge b)}\,\mathrm{I}\neg\wedge_1 \qquad \frac{\neg b}{\neg(a \wedge b)}\,\mathrm{I}\neg\wedge_2$$

$$\frac{\neg a \quad \neg b}{\neg(a \vee b)}\,\mathrm{I}\neg\vee \qquad \frac{a}{a \vee b}\,\mathrm{I}\vee_1 \qquad \frac{b}{a \vee b}\,\mathrm{I}\vee_2$$

$$\frac{a}{\neg\neg a}\,\mathrm{I}\neg\neg \qquad \frac{a \quad \neg a}{\bot}\,\mathrm{I}\bot$$

$$\frac{a \vee b \quad \neg a}{b}\,\mathrm{E}\vee_1 \qquad \frac{a \vee b \quad \neg b}{a}\,\mathrm{E}\vee_2 \qquad \frac{\neg(a \vee b)}{\neg a}\,\mathrm{E}\neg\vee_1 \qquad \frac{\neg(a \vee b)}{\neg b}\,\mathrm{E}\neg\vee_2$$

$$\frac{\neg(a \wedge b) \quad a}{\neg b}\,\mathrm{E}\neg\wedge_1 \qquad \frac{\neg(a \wedge b) \quad b}{\neg a}\,\mathrm{E}\neg\wedge_2 \qquad \frac{a \wedge b}{a}\,\mathrm{E}\wedge_1 \qquad \frac{a \wedge b}{b}\,\mathrm{E}\wedge_2$$

$$\frac{\neg\neg a}{a}\,\mathrm{E}\neg\neg \qquad \frac{\bot}{a}\,\mathrm{E}\bot$$

*Figure 1: Introduction and Elimination rules for Boolean operators $\neg$, $\vee$, $\wedge$. To make it easier to refer to the rules in the remainder of the paper, each rule has been given a unique label shown on the right side of each rule.*

$$\frac{}{\phi \quad \neg\phi}\,\mathrm{RB}$$

*Figure 2: The* rule of bivalence *(RB) for short), also known as the* cut rule*, where $\phi$ is some propositional formula.*

The hierarchy of DBBLs $\vdash_k$ is then defined by the integer *depth* parameter $k$, standing for the maximum number $k$ of nested applications of RB allowed. As the depth $k \to \infty$ we obtain in the limit classical propositional logic.

Let us provide a few more formal notions, for introducing $\vdash_k$. We denote by $\mathcal{L}$ our language, consisting of the set of propositional variables $\mathrm{Var}_{\mathcal{L}}$ and the connectives $\neg, \vee, \wedge$ and by $\mathrm{Fm}_{\mathcal{L}}$ the set of formulas obtained by closing off recursively as usual the variables in $\mathrm{Var}_{\mathcal{L}}$ under the connectives. For any subset $T \subseteq \mathrm{Fm}_{\mathcal{L}}$, we denote by $\mathrm{Var}_{\mathcal{L}}(T)$ the subset of $\mathrm{Var}_{\mathcal{L}}$ of propositional variables occurring in $T$.

We then present the definition of $\vdash_k$, slightly simplifying from D'Agostino, Gabbay, et al. (2024) as follows.

**Definition 1.** *Let $\mathcal{T} \cup \{\phi\} \subseteq \mathrm{Fm}_{\mathcal{L}}$. We say that $\mathcal{T} \vdash_{k+1} \phi$ for $k \geq 0$ if and only if there is $a \in \mathrm{Var}_{\mathcal{L}}(\mathcal{T} \cup \{\phi\})$ such that $\mathcal{T}, a \vdash_k \phi$ and $\mathcal{T}, \neg a \vdash_k \phi$.*

Note that RB is only implicit in the recursive definition above. Nevertheless, we find it useful for readability to include RB in proof trees as an explicit rule of $\vdash_k$, despite the abuse of notation.

An important result of DBBLs is that for any fixed $k \geq 0$, deducibility and refutability of formulas in $\vdash_k$ is polynomial, hence tractable.

As a running example of the application of DBBLs to cognitive agents, consider the following variation on a well-known logical puzzle:

**Example 1** (Children)**.** *Let Anne, Martha, and Jane be children. We know that Anne is blonde, while Jane is not. Additionally, Anne looks at Martha, and Martha looks at Jane (see Figure 3 for a diagrammatic representation).*

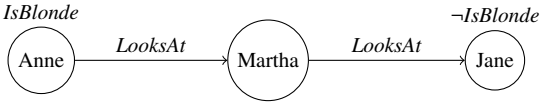


*Figure 3: Anne looks at Martha, who in turn looks at Jane. Is any blonde girl looking at a non-blonde girl?*

*The puzzle involves answering the question:*

*Is it the case that a blonde child is looking at a non-blonde child?* (1)

*This puzzle was informally and empirically examined in Stanovich (1999:* Who Is Rational? Studies of Individual Differences in Reasoning*, Lawrence Erlbaum). Participants could choose one of the following possible answers:*

*(A) Yes*

*(B) No*

*(C) Cannot be determined*

*It has been observed that approximately* 80% *of subjects selected option (C) "Cannot be determined." However, the correct answer is actually (A) "Yes." This can be illustrated by examining two scenarios based on whether Martha is blonde or not. In the first scenario, if Martha is blonde, then she is looking at Jane, who is not blonde. In the second scenario, if Martha is not blonde, then Anne, who is blonde, is looking at Martha. In both cases, the required relationship is maintained. In a first approximation, a first-order axiomatization of this scenario is*


*Figure 4: DBBL Intelim derivation for the children example.*

$\mathcal{T}$ = {*IsBlonde*(*Anne*), ¬*IsBlonde*(*Jane*), *LooksAt*(*Anne*, *Martha*), *LooksAt*(*Martha*, *Jane*)}.

*Note that 0-depth logic does not allow to infer such conclusion. The empirical results thus suggest that people tend to behave as* 0*-depth reasoning agents. A* 1*-depth reasoning agent, however, would be able to correctly apply RB on Martha's hair color and solve the puzzle. This is shown in the derivation in Figure 4, where we decorate the initial assumptions with* $*$ .

The example provides motivation for adopting DBBLs to model similar situations that require reasoning about hypothetical scenarios.

In the remainder of this paper, we demonstrate that DBBLs are

inadequate if individuals are not permitted to choose the incorrect option (C). We argue that to better understand their behavior in this alternative scenario, it is necessary to incorporate some elements of non-monotonicity into DBBLs. This approach is especially relevant in decision-making contexts where individuals face time constraints and encounter counterintuitive or cognitively challenging problems.

### Non-monotonic Reasoning and Logic Programming

Logic programming implements non-monotonicity largely based on a non-classical principle known as *Negation as Failure* (*NF* for short). This principle states that if a certain fact cannot be proven from a given theory, then its negation may be derived. To exemplify its relevance for commonsense reasoning, if a train from Lecce to Verona is not listed in a station's timetable, we typically – and reasonably – conclude that such a train does not exist. Later, if we receive new information that such a train does, in fact, exist – say, we access an updated timetable – we must retract our previous conclusion. Inferences based on NF are, therefore, tentative and subject to change when new information is made available, that is, non-monotonic.

Logic programs consists of *facts* of the form

$$\texttt{f.}$$

and *rules* of the form

$$\texttt{h :- b}_1,\ldots,\texttt{b}_n,\texttt{not b}_{n+1},\ldots,\texttt{not b}_{n+m}. \tag{2}$$

to be read as "$h$ if $b_1$ and $b_2$ and...$b_n$ and not $b_{n+1}$ and... not $b_{n+m}$". At first glance, it may thus seem that $\texttt{f}, \texttt{h}, \texttt{b}_i$, $\texttt{not b}_i$ may be read

as literals, with the `:-` sign serving as an "implication" ← connective and the comma representing a "conjunction" connective.

There are however some significant differences between logic programming and the corresponding classical logic notions.

First, rather than a classical negation, `not p` is to be interpreted as a Negation as Failure, i.e. lack of information concerning `p`. Second, *modus tollens* does not apply to the `:-` operator. Specifically, consider a rule `p :- q`. While the truth of `p` can be inferred when `q` is true, the falsity of `q` cannot be inferred from the falsity of `p`.

A consolidated approach to interpret these constructs is due to Vladimir Lifschitz and consists in computing the *stable models* or *answer sets* of a given logic program. Informally, an answer set is a minimal and consistent set of literals that represents the information implicitly provided by the rules and facts of the program. Here we illustrate the concept only with examples, referring the reader to e.g. Lifschitz (2019: *Answer Set Programming*, Springer Cham) for a formal description. Consider first the following program.

```
q :- not p.
```
(3)

This program has a unique answer set, i.e. {`q`}. This is because the program does not provide any information about `p` and states that, by default, if the agent does not hold positive information about `p` then it may infer that `q` is the case. This latter is the minimal information inferrable from the program and is thus its unique answer set.

Assume now the program (3) is extended with the fact `p`. The rule of the program will now cease to fire and the unique answer set of the resulting program will be this time {`p`}, retracting the previ-

ously endorsed information q. This shows how answer sets capture the non-monotonic nature of negation as failure in logic programming.

Note that, in the light of the mechanism of negation as failure, answer sets need not be unique. To briefly illustrate the concept, consider now the following program :

```
p :- not q.
q :- not p.
```

This is well-known for having two answer sets, i.e. :

{p} and {q}.

Again, {p} is an answer set since, if we were to fire the first rule of the calculus, we would obtain p, due to the lack of information about q. This would then inhibit the application of the second rule where the negation of p is present in the premise. A similar reasoning leads to obtaining the answer set {q}. Both answer sets provide self-consistent minimal models, satisfying the original rules without contradiction. Note that, if we were to add both literals p and q to the program, none of the rules would fire and the only answer set would turn out to be {p, q}.

This latter is an instance of a more general fact concerning answer sets: if a program contains as facts all the literals in its language, i.e. the program represents an agent holding *complete information*, then non-monotonicity ceases to play any role in inference. Formally, in these scenarios either the program has no answer set, or there is a unique answer set amounting to all the facts of the program.

Finally, it is important to highlight that while our discussion of logic programming primarily focuses on propositional programs,

the stable set semantics can also accommodate first-order theories through a process known as *grounding*. This process involves substituting (or *grounding*) variables with constants within the logic program. This allows us, in the following, to use examples in a first-order formalism, but apply on them exclusively propositional reasoning. This convention is particularly useful also for allowing us to focus only on DBBLs, which are propositional logics, rather than their first-order counterparts, which are considerably more intricate, see e.g. D'Agostino, Larese, and Modgil (2021: "Towards depth-bounded natural deduction for classical first-order logic", *IfCoLoG Journal of Logics and their Applications*, pp. 423–451).

## A hierarchy of non-monotonic bounded agents

In this section, we present the main proposal of this paper: the combination of non-monotonic reasoning and bounded reasoning.

This merging is both technically and conceptually challenging. Broadly speaking, the key conceptual issue at hand here is determining whether the non-monotonicity observed in logic programming is: (i) a suboptimal feature of inference that ideal agents should avoid whenever possible, or (ii) a neutral or even beneficial characteristic that cognitively capable agents may choose to utilize. Recalling that non-monotonic reasoning is based on supra-classical inferences, or "jumps to conclusions" and that we related RB with cognitive resources, we may thus either: (i) view RB as inhibiting supra-classical inferences, meaning that higher-depth agents are "less" non-monotonic than their lower-depth counterparts, or (ii) view RB as being neutral towards such inferences or even enhancing them, allowing thus for "more" non-monotonicity, on the basis of virtual information.

Let us note that in work D'Asaro, Baldi, and Primiero (2021: "Introducing k-lingo: a k-depth Bounded Version of ASP System Clingo", *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning*, DOI: 10.24963/kr.2021/65, pp. 661–665) we followed option (ii), albeit implicitly, taking non-monotonicity and depth-boundedness as two orthogonal phenomena. There we preliminarily developed a hierarchy of bounded non-monotonic systems that approximate standard logic programming in the limit. In other words, we developed *bounded non-monotonic logics* approximating a *unbounded non-monotonic logic*.

Instead, the proposal we put forward here joins non-monotonicity and depth-boundedness, under the perspective that more powerful agents will also be less non-monotonic. We introduce thus *bounded non-monotonic logics* that approximate *classical logic*.

The idea underlying this proposal is that at least some form of non-monotonicity is due to the limitations in cognitive power of agents on the one hand and the real-world pressure of providing answers or making decisions on the other.

A thorough analysis of the actual and virtual information available to an agent should then result in an inhibition of her jumps to conclusions, characterizing non-monotonic reasoning. Let us consider a variant of Example 1 as an illustration of this phenomenon.

**Example 2** (Children, continued)**.** *Recall from Example 1 that Anne, Martha, and Jane are children, where Anne is blonde, Jane is not, and Anne looks at Martha, who in turn looks at Jane.*

*In this example, we pose a new question: how would people respond if presented with only two options (A) "Yes" and (B) "No," when the commonly chosen but incorrect option (C) "Cannot be determined" is omitted? We hypothesize that, under strict time*

*constraints, most individuals would select (B) "No." This hypothesis is based on 1) the riddle does not explicitly present a blonde child looking at a non-blonde child and 2) the agents lack the means to infer such relation from the available data. Thus, they may infer that no such relationship exists. We tested this hypothesis in a small-scale experiment. Participants were asked to respond to Question* (1) *within a* 1*-minute time limit, choosing between options (A) "Yes" and (B) "No." A total of* 63 *participants took part in the study. Of these,* 23 *participants (*36.5%*) selected option (A), while* 40 *participants (*63.5%*) selected option (B). The observed preference for option (B) was statistically significant, differing from random choice (*$\chi^2 = 4.0635$*,* $p = 0.04382$*), with a statistical power greater than* 80% *for a medium effect size (*0.5*).*

Our preliminary experimental evidence suggests that bounded agents typically choose the erroneous answer (B). On the other hand, classical agents, i.e., agents correctly capable of reasoning by cases via RB, thus going for the classically correct answer (A). We formalize this behavior in Example 2, assuming that 0-depth agents perform non-monotonic inferences, while 1-depth agents inhibit such inferences through reasoning by cases.

Before proceeding with our proposal, let us recall that rules of logic programs are only deceptively similar to implication, since they fail to satisfy modus tollens.

On the other hand, reading implication $\phi \to \psi$ as $\neg\phi \vee \psi$, it can be easily shown that modus tollens hold even for the logic $\vdash_0$.

For these reasons, following, e.g., Miller (2022: "A Survey of the Proof-Theoretic Foundations of Logic Programming", *Theory and Practice of Logic Programming*, DOI: 10.1017/S1471068421000533, pp. 859–904), we opt for encoding the clauses of logic programs directly as *rules* to be added to the proof

system in 1. Consider a program containing a rule

```
q :- p.
```
(4)

Rather than the formula $\neg p \vee q$, a theory including this rule will add to a logical system, such as the one for $\vdash_0$ in Figure 1, the rule

$$\frac{p}{q}. \tag{5}$$

Similarly, a rule

```
q :- not p.
```
(6)

will be rendered as

$$\frac{\neg p}{q}$$

Note that we are here conflating the standard negation $\neg$ with the negation as failure `not` from logic programming, thus implicitly assuming a closed worlds assumption for the formulas in the premises. More sophisticated logic programming formalisms allow for a fine-grained control of negation, including two distinct operators, the non-monotonic `not` interpreted as negation as failure, but also a classical negation $\neg$.

Contrary to the usual logical rules, $p$ and $q$ are here taken to be concrete atoms, and may not be substituted with arbitrary formulas. Henceforth, by a *theory* $\mathcal{T}$ we will mean a set of formulas, together with additional rules of the form (5) above.

An analog of modus tollens does not hold for $\vdash_0$ extended with rules of the form above, but only for higher DBBLs. Indeed, assume one has a theory containing both the literal $\neg q$ and the rule (5) , i.e., the theory $\mathcal{T} = \left\{\frac{p}{q}, \neg q\right\}$. In $\vdash_0$, it is not possible to infer $\neg p$ using rule (5) and the fact $\neg q$ . On the other hand, this is again possible in 1-depth logic, as follows:

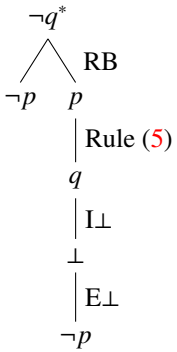


In other words, a 0-depth agent can only apply the rule (5) forward. Nevertheless, a 1-depth agent, conjecturing about the information about the antecedent, may simulate modus tollens. This is coherent with the failure of modus tollens in logic programming on the one hand, and the empirical results attributing more cognitive weight to modus tollens over rules such as modus ponens, on the other, see e.g. Stenning and Lambalgen (2008: *Human Reasoning and Cognitive Science*, MIT Press). We thus believe that our formalism can shed light on the conceptual debate concerning the controversial status of modus tollens in non-monotonic reasoning.

Henceforth, we denote by $ST(\mathcal{T})$ the set of answer sets of a theory $\mathcal{T}$, which may possibly include rules corresponding to rules of logic programs.

We are now ready to introduce our depth-bounded non-monotonic consequence relations, denoted by $\mid\!\sim_k$ for $k = 0, 1, \ldots$, in analogy to the Depth-Bounded Boolean consequence relations $\vdash_k$. Unlike their DBBLs counterpart, $\mid\!\sim_k$ exhibits non-monotonicity, and furthermore, in line with (i), such non-monotonicity will reduce as $k$ increases.

Let us start from defining the $\mid\!\sim_0$ consequence relation.

**Definition 2.** *We write $\mathcal{T} \mid\!\sim_0 \phi$ iff either $ST(\mathcal{T}) = \emptyset$ and $\mathcal{T} \vdash_0 \phi$, or $ST(\mathcal{T}) \neq \emptyset$ and $S \vdash_0 \phi$ for each $S \in ST(\mathcal{T})$.*

This means that $\mid\!\sim_0$ permits to entail from a theory only those formulas that are 0-depth consequences of *all* its answer sets, if there are any, and otherwise it reduces to $\vdash_0$.

Note that while $\vdash_0$ is weaker than classical logic, $\mid\!\sim_0$ is *incomparable* with classical logic. Indeed, as for $\vdash_0$ we have $\emptyset \not\mid\!\sim_0 p \vee \neg p$. On the other hand, given the non-monotonicity of the answer sets, one may infer from a theory also more than what is classically inferable. Consider indeed the theory $\mathcal{T}$ encoding Equation (3).

$$\mathcal{T} = \left\{ \frac{\neg p}{q} \right\} \tag{7}$$

We have that $ST(\mathcal{T}) = \{\{q\}\}$, hence there is a unique answer set for $\mathcal{T}$. One has thus $\mathcal{T} \mid\!\sim_0 q$, since $q \vdash_0 q$, while $\mathcal{T} \not\vdash q$ and in particular also $\mathcal{T} \not\vdash_0 q$.

Before proceeding to the introduction of the hierarchy, let us denote by $Dec(\mathcal{T})$ the subset of the variables in $\mathrm{Var}_{\mathcal{L}}$ that are *decided* by $\mathcal{T}$ at depth 0, i.e. the variables $a$ such that either $\mathcal{T} \vdash_0 a$ or $\mathcal{T} \vdash_0 \neg a$.

**Definition 3.** *We write $\mathcal{T} \mid\!\sim_{k+1} \phi$ for $k \geq 0$ iff there exists $a \in \mathrm{Var}_{\mathcal{L}}(\mathcal{T} \cup \{\phi\}) \setminus Dec(\mathcal{T})$, such that $\mathcal{T}, a \mid\!\sim_k \phi$ and $\mathcal{T}, \neg a \mid\!\sim_k \phi$.*

A few words on the restriction on the virtual information $a$ occurring in the application of RB in the definition above are in place. First, DBBLs typically allow for a restriction on the set of virtual information. Here we have opted for a rather standard restriction (see Definition 1), namely the variables occurring in the premises

$\mathcal{T}$ and the formula $\phi$ that we want to prove. On the other hand, we are also imposing the additional requirement that the virtual information is not already decided by $\mathcal{T}$. Note that, if this were not the case, RB would be useless. Unpacking the recursive definition above, this means that each application of RB at depth $k$ requires the use of a literal *distinct* from the ones used at depths $j$ for $j < k$.

While this additional constraint is harmless for DBBLs, it makes a difference in our non-monotonic hierarchy, as we will see in the next examples.

First, let us recall our former example (7). We have showed that $\mathcal{T} \mid\!\sim_0 q$. To show that $\mathcal{T} \mid\!\sim_1 q$, we need to apply RB either on $q$ or $p$ i.e. show that either both $\mathcal{T}, q \mid\!\sim_0 q$ and $\mathcal{T}, \neg q \mid\!\sim_0 q$ or both $\mathcal{T}, \neg p \mid\!\sim_0 q$ and $\mathcal{T}, p \mid\!\sim_0 q$. Let us start from applying RB on $q$. Clearly $\mathcal{T}, q \mid\!\sim_0 q$ since $\mathcal{T}, q \vdash_0 q$. On the other hand, $ST(\mathcal{T} \cup \{\neg q\}) = \emptyset$, hence $\mathcal{T}, \neg q \not\mid\!\sim_0 q$, since $\mathcal{T}, \neg q \not\sim_0 q$. Let us consider now an application of RB on $p$. On the one hand $\mathcal{T}, \neg p \mid\!\sim_0 q$, since $\mathcal{T}, \neg p \vdash_0 q$, but on the other hand we have $\mathcal{T}, p \not\mid\!\sim_0 q$. Indeed $ST(\mathcal{T} \cup \{p\}) = \{\{p\}\}$ and $p \not\sim_0 q$. In summary, we have obtained that $\mathcal{T} \mid\!\sim_0 q$, while $T \not\mid\!\sim_1 q$. Thus, in contrast with DBBLs, where $\vdash_0 \subseteq \vdash_1$, here we have $\mid\!\sim_0 \not\subseteq \mid\!\sim_1$.

As a final note, recall from the children Example 2 that people seem to show a preference for option (B) "No." We now briefly show that this behavior can be modeled in our framework.

**Example 3** (Children, continued (ii))**.** *Consider the theory* $\mathcal{T} = \{IsBlonde(Anne), \neg IsBlonde(Jane), LooksAt(Anne, Martha), LooksAt(Martha, Jane)\}$ *from our running Example 1, and extend it to* $\mathcal{T}'$*, by adding the rule*

$$\frac{\neg Answer(A)}{Answer(B)}$$

*encoding a* closed world *assumption, and the rule*

$$\frac{LooksAt(x, y) \wedge IsBlonde(x) \wedge \neg IsBlonde(y)}{Answer(A)}$$

*This rule includes in its premises variables that can be eliminated through grounding, considering that our language is finite. At depth* 0*, a non-monotonic DBBL agent that does not make assumptions about whether Martha is blonde or not would fail to obtain answer A, and thus non-monotonically infer that B is the correct answer. In other words, we have* $ST(\mathcal{T}') = \{\{\mathcal{T} \cup \{Answer(B)\}\}\}$ *and thus* $\mathcal{T}' \mid\!\sim_0 Answer(B)$*.*

*This is consistent with the preliminary experimental results in Example 1. Note that an agent able to apply RB on IsBlonde(Martha) would be able to correctly infer that, in fact, answer (A) "Yes" is the correct answer to the puzzle, as its reasoning would reduce to the classical derivation in Figure 4.*

The examples illustrate the feature distinguishing our hierarchy from other works integrating DBBLs and non-monotonic reasoning, such as D'Asaro, Baldi, and Primiero (2021: pp. 661–665) and Chapter 3 of D'Agostino, Gabbay, et al. (2024). In those works we have, for instance, that a 0-depth non-monotonic consequence relation is included in the 1-depth one, while in our case the consequence relations $\{\mid\!\sim_k\}_{k\in\mathbb{N}}$ are all incomparable. In particular, when focusing on those formulas that are not classically derivable, as $k$ increases the consequence relations allow us to infer *less*.

This is a crucial point for proving that the $\{\mid\!\sim_k\}_{k\in\mathbb{N}}$ hierarchy approximate classical logic. By "approximate" we mean, more precisely what follows.

**Theorem 1.** *Let* $\mathcal{T}$ *be a theory consisting of classical formulas*

*and rules of the form* (5)*. Then,*

- *If* $\mathcal{T} \vdash \phi$ *then there exists a* $k \in \mathbb{N}$ *such that* $T \mid\sim_k \phi$*.*
- *If* $\mathcal{T} \nvdash \phi$ *then there exists a* $k \in \mathbb{N}$ *such that* $\mathcal{T} \nmid\sim_k \phi$*.*

*Proof.* The first claim is immediate. Indeed, if $\mathcal{T} \vdash \phi$, by the approximation result in D'Agostino, Gabbay, et al. (2024), there exists a $k \in \mathbb{N}$ such that $T \vdash_k \phi$ . But this fact, by the definition of $\mid\sim_k$, implies $\mathcal{T} \mid\sim_k \phi$ .

Assume now that $\mathcal{T} \nvdash \phi$, and by contradiction, that for each $j \in \mathbb{N}$, we have $\mathcal{T} \mid\sim_j \phi$. Now, as discussed above, each application of RB in Definition 3 requires the choice of a different piece of virtual information. Assume that $|\mathrm{Var}_{\mathcal{L}}(\mathcal{T} \cup \{\phi\})| = k$. Unpacking the recursive Definition 3, we get that $\mathcal{T} \mid\sim_k \phi$ if and only if for each maximal set of literals $S$ built from $\mathrm{Var}_{\mathcal{L}}(\mathcal{T} \cup \{\phi\})$, we have $\mathcal{T}, S \mid\sim_0 \phi$. Now since $S$ contains all the literals of the language, for what we discussed in the previous section, we have either $S\,T(\mathcal{T} \cup S) = \{S\}$ or $S\,T(\mathcal{T} \cup S) = \emptyset$. Hence in both cases $\mathcal{T}, S \mid\sim_0 \phi$ implies $\mathcal{T}, S \vdash_0 \phi$. We thus obtained that $\mathcal{T}, S \vdash_0 \phi$ for each maximal set of literals $S$ from $\mathrm{Var}_{\mathcal{L}}(\mathcal{T} \cup \{\phi\})$. But this in turn means that $\mathcal{T} \vdash_k \phi$, hence $\mathcal{T} \vdash \phi$, contradicting our initial assumption.

□

To conclude, let us note a possible criticism against our approach. It may be taken as too strong in suggesting, at least implicitly, that ideal agents should avoid non-monotonicity in any situation, at any cost. We do not endorse this claim and believe that we can still make our hierarchy compatible with a weaker position regarding non-monotonicity. In particular, we believe that one can take non-monotonic agents using logics of lower depths of the

hierarchy as rationally minimizing their "cost of reasoning", when facing decisions. We thus plan in future work to investigate how to embed our hierarchy in a wider decision-theoretic framework and how this can shed light also on different ways of combining bounded reasoning and non-monotonicity.


## Acknowledgements

We thank the anonymous reviewers for suggestions that helped us improve on the final version of the paper.

Paolo Baldi acknowledges support from the project "Ethical Design for AI", part of the Spoke 6 of the PNRR project FAIR (PE00000013, CUP H97G22000210007), funded by the European Union-NextGenerationEU and the Italian Ministry of University and Research (MUR).



Paolo Baldi
Università degli Studi del Salento
Fabio D'Asaro
Università degli Studi di Verona

# An Interview with Sébastien Rivat

:: Jürgen Landes

**Abstract**
Jürgen Landes interviews Sébastien Rivat.
**Keywords**
Philosophy of Physics, Interview, ERC
**How to Cite** Landes, J. An Interview with Sébastien Rivat. The Reasoner, 19(1). https://doi.org/10.54103/1757-0522/27213

*Sébastien Rivat is an assistant professor at the MCMP in the Stephan Hartmann group. Like many philosophers of science he initially trained as a scientist. He recently won an ERC grant; congratulations! Before leaving you to the interview, I'd like to warmly thank him for his time. Beyond the exciting research he talks about, he also has some sound advice on writing a successful grant proposal. I thus hope that even readers not (yet!) interested in the philosophy of physics have reasons to read the interview.*

JL: Tell us about your academic background

**SR:** I originally come from mathematics and physics. After finishing my undergraduate studies in France, I moved to Cambridge in the UK to complete my master in theoretical physics. I stayed one more year to complete a MPhil in history and philosophy of science, and I was lucky to be accepted right after that into Columbia University's PhD program in philosophy. It was a wonderful op-

portunity to explore a variety of philosophical topics. I eventually came to develop a thesis project on the philosophy of relativistic quantum field theory, the theoretical framework that unifies special relativity and quantum mechanics. And this project eventually brought me toward the end of my PhD to develop projects on the history of contemporary physics. After my defense, I spent two years at the Max Planck Institute for the History of Science in Berlin to work on these projects, and I joined the Munich Center for Mathematical Philosophy at LMU Munich in 2022 to develop other philosophical projects. Going back and forth between the two has proved to be really productive in my case.

JL: How did you get into the philosophy of physics?

SR: My master in theoretical physics really helped me to get a first footstep into the nitty-gritty details of physics. But I also got increasingly confused about the structure of more advanced theories and the underlying world picture I thought I was supposed to get out of them. I felt that it would make more sense to try to understand what these theories are all about before going any further. At the time, Jeremy Butterfield was giving a seminar on the foundation of quantum theories at the Center for Mathematical Sciences at Cambridge. I think he played a major role in me switching to philosophy of physics. In hindsight, I also think that the switch was a rather natural one. Master programs in theoretical physics usually require some serious engagement with the details of quantum field theory. But despite its remarkable empirical success, this framework still lacks a proper conceptual and mathematical foundation. Establishing the mathematical existence of realistic quantum field theories is even one of the Millennium Prize Problems selected by the Clay Mathematics Institute in 2000. So, upon encountering quantum field theory, I imagine that many students may also have stumbled on the fundamental question ‘how

on earth is it possible for such a mathematically ill-defined and seemingly inconsistent framework to be so empirically successful at the same time?'. I have been particularly interested in trying to understand how the approximation methods involved in this framework work and what the world must be like for them to work so well. But I should mention that many others have been more interested in the mathematical foundations of quantum field theory per se.

JL: What is hotly debated in the philosophy of physics?

SR: There are obviously many hot debates right now, and my selection will be very partial. The longstanding issues of explaining the emergence of the arrow of time (i.e., why macroscopic phenomena don't seem to go backward in time) and accounting for the metaphysics of quantum mechanics (e.g., what "superposed" quantum states correspond to) are still widely debated. Philosophers more interested in space-time theories have been concerned with issues of theoretical equivalence, reduction, and underdetermination in recent years, to paint them in very broad terms. For instance, there has been a bit of back and forth lately about how to best understand the sense in which the curved spacetimes of classical general relativity reduce to the flat spacetimes of special relativity in sufficiently well-behaved local spacetime patches. On the particle physics side, philosophers have been worried about isolating the sorts of entities and structures that underwrite the success of our best current theories. Taking seriously the idea that these theories are merely approximative raises surprisingly difficult interpretative issues (see, e.g., https://philsci-archive.pitt.edu/18056/). And there are of course a number of emerging debates that cut across theories of space-time and matter, including the black hole information paradox, the emergence of space-time in quantum gravity, and the

relative epistemic standing of competing models of cosmology.

JL: You recently won an ERC project with title: "The Scaling Revolution in Physics: Historical and Philosophical Perspectives". Could you tell us what this project is about?

SR: The project aims to examine the historical development and the philosophical implications of physicists' new way of studying physical systems "scale by scale" and "across scales". To give a simple picture, physicists have been very much concerned at least since the rise of modern science with developing theories of space, time, matter, and motion that could apply in principle everywhere and under any circumstances. Philosophers have also very much relied on this ideal of formulating a final theory to interpret the content of theories that physicists came to formulate along the way. This is particularly clear in my sense today with classical general relativity and non-relativistic quantum mechanics. But it seems that the situation has considerably changed since the 1950s. Physicists now widely recognize that the best theories we currently have work only at certain scales. To put it more precisely, physicists widely recognize that these theories are more reliably formulated as effective theories (on pain of being inconsistent and making unwarranted assumptions about unexplored domains). Physicists also now spend most of their time studying systems scale by scale and analyzing how physical effects depend on each other across scales, even when they work on speculative models of quantum gravity. This more cautious attitude toward the frontiers of physics and this general methodological shift raises two fundamental questions in my sense: (i) How should we adjust our interpretative practice in light of this historical transformation? (ii) What kind of world picture should we extract from this new way of doing physics scale by scale? The project will address these questions by taking two of the most sig-

nificant scale-based theoretical tools developed by physicists from the 1950s onward: namely, effective field theories and renormalization group methods. And we will integrate both historical and philosophical analyses to try to get as clear and precise as we can about the impact of this historical transformation on fundamental issues pertaining to scientific representation, reduction, scientific realism, and the metaphysics of science.

JL: What is it that people should be doing in the philosophy of physics but are not doing?

SR: Let me say at the outset that philosophy of physics is living a very healthy life at the moment. There is a large variety of different approaches on a wide range of topics, from groundbreaking conceptual analyses published in mathematical physics journals to naturalized metaphysical studies published in top-tier generalist journals. But if I had to highlight one important blind spot, it would be the lack of engagement with the fact that all the empirically successful theories that we have so far are "merely" approximative. Concretely, this means that philosophers of physics would benefit from engaging more thoroughly with the variety of approximation methods developed by physicists and thinking further about how to reconcile our favorite all-or-nothing philosophical categories with the approximate and partial character of the systems we wish to study. Kerry McKenzie has recently done some very important work in that respect (see, e.g., https://philsci-archive.pitt.edu/23033/, for her latest paper on the topic).

JL: Are you interested in other areas you want to tell us about?

SR: My work in the philosophy of physics has indeed led me to

think about a number of more general issues. For instance, I have been thinking lately about the epistemological status of discretization methods in science (e.g., what are the epistemic benefits gained by using discrete mathematics?). I started to think about this question from the perspective of lattice theories in physics, where the continuous background of standard theories is replaced by a discrete grid made of sites and links. And I quickly realized that there are actually serious theoretical obstacles if we want to put our best current theories on a lattice. The physicist David Tong actually gives a very accessible talk on this topic, suggesting that these obstacles put some pressure on the idea that the actual world could be simulated by a computer. Whether the argument is decisive or not, I highly recommend watching the recording of his talk online (https://www.youtube.com/watch?v=QPMn7SuiHP8).

JL: I recently conducted in an interview with editors on writing good papers. As a successful grant writer I'm now turning to you for some tips/tricks/advise for our readers on how to write a good grant.

SR: My experience is very limited. But from what I have seen so far, a good grant proposal starts with a simple and exciting idea that one will eventually be able to summarize in three minutes. This is actually the time given to present one's project in many panels of the ERC grant program. Typically, this simple idea takes the following form in the case of a philosophical project: (i) most philosophers in the field F have assumed (or done) X; (ii) but Y suggests that X is wrong (or wrongheaded); (iii) the project will show how taking Y seriously has a fundamental impact on the field F. Then, to write a good grant proposal, there are three challenges to overcome: (i) finding what is missing in the field while remaining charitable and exhaustive; (ii) finding Y (interdisciplinarity helps but usually raises additional methodological

issues); (iii) having a detailed and concrete plan for the implementation of the project (which usually requires having already done some work on the impact of Y).

JÜRGEN LANDES
MCMP Munich

# Representativeness Heuristic is a Potential Danger to the Diagnostic Process

:: Paolo Luigi Balzaretti

**Abstract**

Similarly to the ordinary, the diagnostic process is characterized by the formulation of initial hypotheses to which probability judgments are associated, sometimes generated through heuristics and biases. The use of mental "shortcuts" can negatively affect the patient's management.

**Keywords**

Heuristics; biases; diagnosis; representativeness.

**How to Cite**

Balzaretti, P. L. Representativeness heuristic is a potential danger to the diagnostic process. The Reasoner, 19(1). https://doi.org/10.54103/1757-0522/27372

Dear Editor,

I read the article *Role of heuristics in diagnostic reasoning in practice* (Full Issue) with interest, and I would like to propose some criticisms of the conclusions reached by the Author.

The Author claims that diagnostic reasoning is articulated in the following steps: first, the physician examines the patient and "suspects" a possible diagnostic hypothesis; next, he chooses a test to evaluate this hypothesis; and finally, he combines the prior probability of disease as estimated from its population prevalence with

the relative likelihood ratio (LR) to calculate a posterior probability of disease. He affirms that such early hypotheses, being just suspicions and not probability judgements, can be discarded if proved wrong without introducing any risk of error in the whole process.

In my opinion, to "suspect" something is to generate a personal belief about the actual cause of a patient complaint. Different diseases can usually cause a symptom, so multiple diagnostic beliefs are generated, and specific degrees of confidence, which can be formalized as probabilities, are associated with each.

In the clinical context, probability judgments are formulated based on uncertain data, precisely the context in which heuristics are used, and biases can occur. The use of mental shortcuts is further confirmed by the evidence according to which the first diagnostic hypotheses (and the respective pre-test probabilities) arise within a few seconds from the beginning of the encounter with the patient (Pelaccia et al. 2014: How and when do expert emergency physicians generate and evaluate diagnostic hypotheses? A qualitative study using head-mounted video cued-recall interviews, Annals of Emergency Medicine, 64(6):575-85, doi: 10.1016/j.annemergmed.2014.05.003). The degree of belief in these early hypotheses, which seldom relies on the knowledge of population prevalence, drives the following decisions about further testing.

It is possible that, among others, the representativeness heuristic is employed, according to which the probability that object A belongs to class B is estimated from the degree of similarity of the former to the latter. If A is highly representative of B, the probability that the former belongs to the latter is considered high. At the same time, if A has no characteristics that make it simi-

lar to B (thus is "atypical" with respect to B), the probability of A belonging to B is considered low (Tversky and Kahnemann 1974: Judgment under Uncertainty: Heuristics and Biases, Science, 185(4157):1124-31, doi: 10.1126/science.185.4157.1124). Consequently, a clinical presentation that bears little resemblance to the description of the disease contained in the physician's mental model could be associated with such a low degree of confidence that it is not further analyzed.

Empirical evidence about the use of heuristics and bias in clinical medicine is scarce due to the intrinsic difficulties in defining both the population of possibly affected subjects and the misdiagnosis cases. Nonetheless, some accounts of the use of representativeness heuristics exist. According to Kulkarny and colleagues, victims of trauma defined with objective scales as moderate-severe but who do not present "typical" characteristics for the condition are transferred less frequently to dedicated medical centers (Kulkarny et al. 2019: Defining the representativeness heuristic in trauma triage: A retrospective observational cohort study, PLoS One, 14(2):e0212201, doi: 10.1371/journal.pone.0212201).

Contrary to what the Author of the essay claims, at least one case of death related to the representativeness heuristic has been reported in the literature. In a 70-years-old woman, the absence of characteristic features of the disease (such as hypotension, severe chest pain, and increased diameter of the mediastinal profile on chest X-rays) meant that a diagnosis of aortic dissection was not considered and led to an unfortunate outcome (Moffat et al 2022: Dual-processing theory helps to explain delay in diagnosis of Stanford type A aortic dissection, BMJ Case Reports, 22;15(4):e242036, doi: 10.1136/bcr-2021-242036).

In conclusion, there is evidence that the diagnostic pro-

cess is similar, in most cases, to ordinary reasoning. It includes an initial phase of hypothesis generation, in which the doctor may resort to biases and heuristics, including representativeness. Using the representativeness heuristic may correlate with worse outcomes for patients.

Paolo Luigic Balzaretti, M.D
Azienda Ospedaliera Ordine Mauriziano di Torino.

# Fourth BRIO Meeting & Idea League Ethics

:: Giacomo Zanotti

**Abstract**

Report on the Fourth BRIO Meeting  Idea League Ethics Working Group Workshop held at the Lecco Campus of Politecnico di Milano on 4-5 July 2024.

**Keywords**

BRIO; Trustworthy AI; Ethical aspects of AI.

**How to Cite**

Zanotti, G. Fourth BRIO Meeting & Idea League Ethics Working Group Workshop. The Reasoner, 19(1). https://doi.org/10.54103/1757-0522/27117

The *Fourth BRIO Meeting & Idea League Ethics Working Group Workshop* was conceived as a collaboration between the project BRIO – Bias, Risk and Opacity in AI (PRIN MUR), and the *Ethics Working Group of* IDEA League.

The workshop took place at the Lecco Campus of Politecnico di Milano on the 4th and 5th of July 2024, and its main aim was to create a fertile ground for interdisciplinary discussions around the ethical, technical, and societal aspects of Trustworthy AI.

The first speaker was Saskia Nagel (RWTH Aachen University), who gave a talk on "Trusting relationships – the case of human-technology interaction in medicine". Looking at the specific context of human-technology interaction in medicine, Nagel analysed

the conceptual tenability and the moral costs of extending the notion of trust to individual technological artifacts. After that, Giacomo Zanotti gave a talk on risk and uncertainty in AI, presenting the developments of the research conducted at Politecnico di Milano in collaboration with Daniele Chiffi and Viola Schiaffonati, and Emanuele Bottazzi Grifoni investigated the issue of understanding in LLMs by drawing on later Wittgenstein's philosophy of language. Karl de Fine Licht (Chalmers University of Technology) was the last speaker of the day, presenting a work on "Generative AI in higher education". Starting from an analysis of the actual capabilities of current Generative AI tools, with a focus on their prospects of deployment in higher education, de Fine Licht considered some potential positive and negative aspects of such deployment, discussing some strategies to collectively evaluate its legitimacy.

The first talk of the second day was given by Ben Wagner (TU Delft), who presented his work "Interlocking governance mechanisms: towards robust accountability of AI systems". Wagner argued that effective AI governance needs interlocking mechanisms – e.g., transparency databases, audit mechanisms, and financial and corporate reporting – to ensure accountability. After that, Ekaterina Kubyshkina presented her work (in collaboration with Mattia Petrolo and Giuseppe Primiero) on hyperintensionality in Trustworthy AI, distinguishing on a formal basis two sources of hyperintensionality in trust. The last talk of the morning presented the developments of the collaboration between Alkemy (Davide Posillipo) and BRIO's unit at the University of Milano (Greta Coraglia, Giuseppe Primiero, Francesco Genco), discussing the theory behind the BRIOxAlkemy tool for the evaluation of fairness-related risk.

The last two talks of the workshop were given by Salvatore

Giugliano and Daniele Porello. Giugliano presented a work (co-authored with Andrea Apicella, Francesco Isgrò and Roberto Prevete) on the improvement of the performance of ML through XAI techniques, while Porello focused on the issue of classification biases and presented an ontological framework for protected predicates.

The event was successful and fostered fruitful discussions of different aspects of Trustworthy AI at the intersection of computer science, logic, and philosophy.

Giacomo Zanotti
Università degli Studi di Pavia

# SMARTEST 1st Project Meeting

:: Alessandro Acciai and Rossella Suriano

**Abstract**

Report on the first official meeting of the SMARTEST project which took place at the CUMO in Noto, Sicily, on 9-10 September 2024

**Keywords**

SMARTEST; Digital Twins; Machine Learning; Ethics of AI.

**How to Cite**

Acciai, A., & Suriano, R. SMARTEST 1st Project Meeting. The Reasoner, 19(1). https://doi.org/10.54103/1757-0522/27141

On September 9 and 10, 2024, the first official meeting of the (SMARTEST) project took place at the CUMO in Noto, Sicily. The project, as presented in Volume 18, Issue 4, centers on studying and analyzing Digital Twin (DT) simulation through an ontological, philosophical, and epistemological approach, involving academic and researching institutions from Trento at CNR-ISTC Laboratory for Applied Ontology , University of Milan at Logic Uncertainty Computation and Information , and at Department of Cognitive Science at University of Messina. The event gathered key members of the SMARTEST research team and various guests from the academic world.

The workshop opened with a presentation by Balbir S. Barn from Middlesex University, who discussed a study conducted with the

London Digital Twin Research Centre, examining DTs from a sociotechnical perspective. The study highlighted the methodological aspects of this approach through various public policy case studies. Key issues raised included challenges and limitations of DT modeling, such as managing simplification in complex systems, the margin of error in sociotechnical models, the role of DTs in understanding internal system operations versus acting as ”black boxes,” and the complexity of modeling language semantics.

Next, Giuseppe Primiero from the LUCI Laboratory at the University of Milan presented on the formal logical relationships between DT simulation models, underscoring the importance of formal verification in assessing the reliability of digital replicas of machine learning models. This evaluation process goes beyond a binary classification of reliability, focusing on more nuanced assessments based on specific validity conditions compared to the original models. This framework aims to ensure that digital replicas maintain both reliability and ethical integrity.

Alessandro Buda from the University of Milan introduced GAMERA: Generative AI Mitigation for Ethical and Responsible Algorithms. GAMERA proposes to expand the landscape of trust by providing a Kripke-style counterpart to Carnap-style semantics.

From the School of Computing at the University of Leeds, Vania Dimitrova, in her presentation titled “Developing City Infrastructure Ontologies to Support Complex Decision Making,” prompted critical reflection on developing urban infrastructure in a smarter and more sustainable manner. She proposed holistic models that integrate knowledge and data, using the latter to validate and optimize cognitive models. Dimitrova emphasized the importance of adequate risk assessments, human oversight, and traceability, un-

derscoring the need for a "human-in-the-loop" approach involving stakeholders from the early stages to ensure effectiveness and participation.

Representing ISTC-CNR in Trento, Roberta Ferrario and Luca Bicchieri presented their work. Ferrario's presentation, "Foundational Issues for an Ontology of Digital Twin," focused on the ontological nature of DTs as artifacts and information objects. The final perspective of the work is that Digital Twins collectively form multiple models, encompassing both their physical counterparts (cyber-physical systems) and their social sphere influence (sociotechnical systems). Luca Bicchieri presented a study titled "Digital Twins for Trustworthy Human-Robot Interaction," addressing trust issues in Human-Robot Interaction (HRI) and proposing an ontological approach. He demonstrated results from a robotic case study in a fruit-sorting experiment, using the DT approach as an alternative setting for evaluating trust in HRI compared to classic game-theory contexts.

John Symons from the University of Kansas explored the theme "Domain-Specific Machine Learning in Scientific Inquiry," emphasizing how finite agents require support to progress in various tasks. Symons discussed the importance of developing systems that can reliably execute algorithms and overcome cognitive limitations. He addressed the No Free Lunch theorem, which posits that no universally superior learning algorithm exists for all problems. Furthermore, he considered the role of optimization in scientific inquiry and the principles of algorithmic learning, such as efficiency and hypothesis revision. In conclusion, he emphasized that a well-defined domain can become fertile ground for machine learning, promoting effective optimization methods.

Representing the University of Messina, Nicola Angius proposed

an epistemological and methodological analysis in a work titled "The Simulative Role of Neural Language Models in Brain Language Processing." He discussed the trend of applying the synthetic method of cognitive science through the use of neural language models to simulate brain language processing. He proposed a co-simulation approach, highlighting how the epistemic opacity of the transformer architecture is tackled by using the brain itself to simulate the neural language model, giving the opportunity to acquiring knowledge and testing hypotheses about both system.

Lucia Guerrisi, also from the University of Messina, introduced a modeling approach involving Deep Learning (DL) as a simulative method. The case presented concerned the Meteorological Neural Network 2, proposing a DL approach for model verification and validation. The study demonstrated how modeling in scientific contexts can be an approximate simulative method, especially in cases where structural similarities are identified through network interpretation.

The workshop concluded with Alessio Plebe's "Determinism in Neural Language Models" on the benefits of non-determinism in neural language models and its importance as a strategic element in their design. He highlighted the role of noise and variability in enhancing neural network performance. Exploring concepts like stochastic resonance, he discussed how noise influences biological processes and suggested similar mechanisms for computational models. Through nuclear sampling he illustrated how variability and non-determinism can enrich deep learning applications.

Alessandro Acciai
Università degli Studi di Messina
Rossella Suriano
Università degli Studi di Messina

# Probabilistic Reasoning in the Sciences :: Michal Sikorski, Alexander Gebharter and Barbara Osimani

**Abstract**
Report on the Conference on Probabilistic Reasoning in the Sciences which took place at the Marche Polytechnic University in Ancona, Italy. 29-31 August 2024.
**Keywords**
Probabilistic reasoning; Science; Methodology.
**How to Cite**
Sikorski, M., Gebharter, A., & Osimani, B. Probabilistic Reasoning in the Sciences. The Reasoner, 19(1). https://doi.org/10.54103/1757-0522/27544

The *Conference on Probabilistic Reasoning in the Sciences* took place at the *Marche Polytechnic University* in Ancona, Italy, from August 29th to 31st 2024. It was hosted by the *Center for Philosophy, Science, and Policy (CPSP)* and organized by Michał Sikorski, Alexander Gebharter, and Barbara Osimani. The main objective of the conference was to bring together philosophers of science and philosophers interested in using probabilistic tools in science. Probabilistic reasoning was discussed from both a foundational and a practical perspective.

The event had more than 30 participants, most of whom were based in Europe, with some attending from overseas. The conference was kicked off by welcoming words from Abele Donati,

the director of the *Department of Biomedical Sciences and Public Health*, and Stefano Staffolani, the dean of the *Economics Faculty*. There were six plenary lectures and 24 contributed papers. The plenary lectures were given by Saana Jukola (University of Twente), David Papineau (King's College London), Jan-Willem Romeijn (University of Groningen), Elliott Sober (University of Wisconsin), Jan Sprenger (University of Turin), and Katya Tentori (University of Trento).

Saana Jukola's talk "Evidence, Uncertainty, and Public Health Policies" investigated the question of how epistemic risk could and should be addressed in scientific reasoning and evidence-based policy. Based on cases from nutrition science and public health policy she argued that institutions rather than individuals should be responsible for such risks and that the collective responsibility literature might be useful to shed further light on these issues.

David Papineau's talk entitled "Causal Inference and the Metaphysics of Causation" explored the metaphysics of causal inference. He outlined a reductionist probabilistic account of causation as a metaphysical basis for causal inference. He then discussed how this theory naturally explains single-case causation, causal asymmetry, and the relation between causation and rational action. Finally, he compared his account with "interventionist" approaches to causation.

Jan-Willem Romeijn, in his talk "Statistical Evidence", argued that statistical methods are inevitably relative. Bayesian inference relies on the choice of a prior, while frequentist statistical methods—referred to here as "Bernoullian"—depend on the structure of the sample space, which exposes them to violations of the likelihood principle. In the second part of the talk, he discussed what

this means for the epistemic authority of statistical science. He offered a conception of evidence that is both relativized and objective, combining insights from logical empiricism and pragmatism. Finally, he argued that this conception of relativized objectivity can counter post-truth views on statistics.

Elliott Sober's talk "Darwin, Common Ancestry, and the Law of Likelihood" focused on Darwin's views concerning common ancestry and the Law of Likelihood. Darwin argued that if two species share a trait $T$, this is evidence for common ancestry. He also believed that adaptive similarities provide only weak evidence for common ancestry, whereas neutral and deleterious similarities offer stronger evidence. Sober evaluated both claims using the Law of Likelihood. In addition, he described how Darwin tested hypotheses about natural selection by inferring the characteristics of ancestors from the observed traits of present-day organisms.

Jan Sprenger's talk "Counterfactual and Probabilistic Reasoning" addressed a gap in the counterfactual reasoning literature concerning the relationship between causal and semantic approaches. He sketched a unified picture based on a "division of labor" between the truth and assertion conditions of counterfactuals, where assertion conditions correspond to truth in all relevant worlds within a given context. Finally, he argued that this framework brings the semantic analysis of counterfactuals closer to their use in scientific reasoning, particularly in relation to the *do*-operator in causal Bayes nets.

In her talk "Promoting a Culture of Uncertainty to Reduce Biases in Medical Reasoning" Katya Tentori (University of Trento) argued that the misunderstanding and misuse of scientific evidence are often rooted in a small number of probabilistic fallacies. She

highlighted that both scientists and the general population frequently struggle to draw accurate conclusions from statistical evidence, as exemplified by misunderstandings of vaccine efficacy results reported during the COVID-19 pandemic. Her main conclusion emphasized that enhancing probabilistic knowledge is essential for the more effective use of scientific evidence.

The conference was funded by Alexander Gebharter's and Lorenzo Rossi's PRIN PNRR project *Controlling and Utilizing Uncertainty in the Sciences* (CUP I53D23006890001), financed by the European Union (Next Generation EU).

The conference featured a Best Paper Award for students and PhD candidates. All submissions were evaluated by an expert panel consisting of Serena Doria, Hykel Hosni, and Jan Sprenger. The prize was awarded to Giacomo Molinari (University of Bristol, UK) for his paper "Disagreement over Imprecise Evidence" in which he argues that scarcity and ambiguity warrant different doxastic responses, develops deference principles for scarce and ambiguous evidence, and applies the resulting account to situations of peer disagreement.

We would like to thank the plenary speakers, the members of the best paper award expert panel, and all speakers and participants for making this event a success and are looking forward to possible future instalments. Finally, we would like to take this opportunity to invite submissions to a topical collection in Synthese devoted to the topic of the conference. Details and the call for papers can

be found at https://link.springer.com/collections/eaaaejgebj.

Michal Sikorski
Università Politecnica delle Marche
Alexander Gebharter
Università Università Politecnica delle Marche
Barbara Osimani
Università Università Politecnica delle Marche