
THE REASONER

VOLUME 16, NUMBER 5
JULY-AUGUST 2022

thereasoner.org
ISSN 1757-0522

CONTENTS

Guest Editorial

Features

Dissemination Corner

What's Hot in ...

Events

Courses and Programmes

Jobs and Studentships

37
38
40
41
44
44
45

soning under Coherence in System P, *Annals of Mathematics and Artificial Intelligence*, 5–34). Moreover, he contributed substantially to the coherence interpretation of conditional random quantities, which has been exploited to developing a theory of compound and iterated conditionals. Angelo's work is highly relevant for the mathematics, philosophy, artificial intelligence, and psychology of reasoning under uncertainty. Giuseppe met Angelo when he was a student in Mathematics at the University of Catania because he chose Angelo's course on Probability. Then, Giuseppe even if he liked the course of Angelo very much, tried to do a thesis in a different field because of a scholarship offered by a multinational company. After this detour, studying FORTRAN coding, Giuseppe quit that scholarship and decided in 1999 to write his (graduate degree) thesis with Angelo, on the implementation of some algorithms related with probabilistic results of Thomas Lukasiewicz on conjunctive events. This was the starting point of an ongoing fruitful collaboration and friendship between Angelo and Giuseppe.



GUEST EDITORIAL

It is with great pleasure to present this July issue of *The Reasoner*. We have chosen our friend [Angelo Gilio](#), full professor at the Department of Mathematics and Computer Science, University of Catania (1994–2000), and later at the Department of Basic and Applied Sciences for Engineering, University of Rome “La Sapienza”, Italy (now retired), as our interview partner because of his well-known work on coherence-based probability theory. Based on Bruno de Finetti's pioneering work on subjective probability, Angelo generalized the notion of coherence to conditional probability in terms of the penalty criterion (and later in terms of proper scoring rules). He also introduced the notion of *generalized-coherence* (g-coherence) for managing imprecise (conditional) probabilities and algorithms for coherence checking and coherent probability propagation. This led to many applications in the domain of reasoning under incomplete knowledge and uncertainty, including a coherence-based probability semantics for the nonmonotonic reasoning System P (see, e.g., Gilio 2002: Probabilistic Rea-

Niki met Angelo for the first time at the *1st Salzburg Workshop on Paradigms of Cognition* 2002, when he presented his first research talk on experimental-psychological results on selected probability propagation rules Angelo proved for the basic nonmonotonic System P. It turned out that Angelo's coherence-based probability semantics of System P is not only formally adequate, but also psychologically highly plausible: the majority of human interval responses are within the best possible coherent probability bounds. Of course, Angelo was pleased with the experimental results. When Niki met Angelo at WUPES (Liblice, Czech Republic) in 2009, he introduced him to Giuseppe as one of the future experts on coherence. His FWF project allowed Niki to invite Giuseppe for a research stay in Salzburg. Then, joint projects and the friendship among the

three of us emerged.

NIKI PFEIFER

Department of Philosophy, University of Regensburg

GIUSEPPE SANFILIPPO

Department of Mathematics and Computer Science,
University of Palermo

FEATURES

Interview with Angelo Gilio

NIKI PFEIFER & GIUSEPPE SANFILIPPO: Angelo, thank you very much for agreeing to be interviewed for this issue of *The Reasoner*. Could you please tell us something about your intellectual history? How did you become interested in probability theory?

ANGELO GILIO: I was born in Acerenza, a small town in the province of Potenza, Basilicata, home to a beautiful Roman Catholic cathedral, one of the most notable Romanesque structures in this part of Italy. In 1968 I moved near Rome, as a guest of my brother Domenico, to enroll in Mathematics at “La Sapienza” University of Rome. I was in my second year of mathematics when one day a friend of my brother, captain of the Italian army, came to lunch. During lunch he talked to me about de Finetti’s subjective theory of probability... a subject then mysterious to me. In the third year of university, having chosen the degree program in Applied Mathematics, when attending lessons in the Probability course I discovered the extraordinary personality of professor Bruno de Finetti and the fascinating world of subjective probability. In the fourth year I wrote my thesis with de Finetti. With him I also had a scholarship for undergraduates, and after a scholarship for graduates, at the National Research Council (CNR). In that period I studied Markov processes and dynamic programming, with applications to economics, under the guidance of professor Giandomenico Maione (assistant to de Finetti, teaching Mathematical Statistics).

Later, I had a four-year contract at the Faculty of Engineering where, beside some research on Bayesian-Adaptive methods for the classification of segments of speech signals (with Paolo Mandarini, Giordano Bruno, Maria G and Maria D Di Benedetto), I continued to cultivate my passion for the coherence-based probability theory.

In those years, in collaboration with G Bruno, I wrote a paper (in Italian) that applied linear programming to de Finetti’s *fundamental theorem of probabilities*, for the computation of the interval of coherent extensions of an initial assignment of probabilities (Bruno, Gilio 1980: *Applicazione del metodo del simplesso al teorema fondamentale per le probabilità nella concezione soggettivistica*, *Statistica*, 337–344).

NP & GS : Were you still in contact with de Finetti in these days?

AG: I recall that we visited de Finetti at his home and received a positive evaluation on our work. In order to publish the paper, de Finetti invited us to contact professor Giorgio Dall’Aglia (Faculty of Statistics), who told us to contact professor Romano Scozzafava (Faculty of Engineering), an expert in the field of subjective probability. That was the occasion where I met Scozzafava, who helped us to publish the paper in *Statistica*, an Italian journal. In that period I decided to move to the Institute of Applied Mathematics to start the scientific

collaboration with Scozzafava. A curious episode: in 1984 in a seminar at the Faculty of Statistics, Frank Lad illustrated his work on the application of linear programming to de Finetti’s fundamental theorem. Scozzafava, who was present at the seminar, pointed out that G Bruno and I had published a similar work in 1980.

In that period I also wrote, in collaboration with G Bruno, a paper on logical operations among conditional events (a topic that has become central in the last 10 years to my research activity, in collaboration with Giuseppe Sanfilippo).

In the same period I deepened the study of the penalty criterion of de Finetti, by proposing a slight modification for the definition of coherence, which provides a unified approach for conditional and unconditional events (Gilio 1990: *Criterio di penalizzazione e condizioni di coerenza nella valutazione soggettiva della probabilità*, *Bollettino U.M.I.*, 645–660). In those years I started to represent numerically the third value, void, of a conditional event $A|H$ by its conditional probability $P(A|H)$. Twenty years later, in a joint work with Sanfilippo, the notion of coherence based on the penalty criterion had been generalized by exploiting proper scoring rules (Gilio, Sanfilippo 2011: *Coherent conditional probabilities and proper scoring rules*, *ISIPTA’11*, 189–198).

NP & GS : How long did you stay in working contact with de Finetti, and with whom have you been working thereafter in the field of subjective probability?

AG: Actually, I had contact with de Finetti mainly as a student, developing my research on subjective probability alone and, in the following years, with Romano Scozzafava, Giulianella Coletti, Pier Giorgio Marchetti, Fulvio Spezzaferrri, Salvatore Ingrassia, Veronica Biazzo, Rosalba Giugno, Alfredo Ferro, Thomas Lukasiewicz, Didier Dubois, Gabriele Kern-Isberner, Tommaso Flaminio, Lluís Godo, with Giuseppe Sanfilippo over the past 23 years, and with David Over and Niki Pfeifer for more than 10 years now.

NP & GS : Did you enjoy working at Catania University?

AG: Catania is a beautiful city in Sicily, close to Vulcano Etna. The experience with the scientific community at the Department of Mathematics and Computer Science was very significant. The only difficult aspects for me were the trips from Rome to Catania and vice versa, almost each week, for more than six years (in that period I also taught in Rome).

NP & GS : What are in your view the key features of coherence?

AG: The coherence principle of de Finetti appears as the only axiom by which we can develop a consistent probabilistic theory, with no need of using algebraic structures for the family of conditional events and/or random quantities; moreover, we can properly manage conditioning events of zero probability. By exploiting the coherence principle, I developed with full generality the theory of Adams on conditionals, by computing the lower and upper bounds for the consequents of some inference rules in probabilistic nonmonotonic reasoning (Gilio: 2002). I recall that Adams, when defining the notions of p-consistency and p-entailment, requires that probability distributions be *proper*, while in the setting of coherence we can avoid this constraint.

With Sanfilippo (as well as with Over, Pfeifer, Flaminio, and Godo), we obtained many developments on compound conditionals and iterated conditionals, in the setting of coherence, with applications to probabilistic nonmonotonic reasoning and to the psychology of reasoning under uncertainty (e.g.,

Gilio, Over 2012: The Psychology of Inferring Conditionals from Disjunctions: A Probabilistic Study, *Journal of Mathematical Psychology*, 118–31; Sanfilippo, Gilio, Over, Pfeifer 2020: Probabilities of Conditionals and Previsions of Iterated Conditionals, *Int. J. Approx. Reas.*, 121, 150-73; Flaminio, Gilio, Godo, Sanfilippo 2022: Compound conditionals as random quantities and Boolean algebras, *KR* 2022).

NP & GS : Could you let the readers know, in not too technical terms, the differences among the different notions of coherence (coherence, g-coherence, and total coherence)? Which of these notions is in your view the most important one for investigating reasoning under uncertainty?

AG: The basic (and hence the most important) notion of coherence, introduced by de Finetti, concerns *precise* probability assessments, like $P(A|H) = 0.8$, $P(B|K) = 0.5$ for given conditional events $A|H$ and $B|K$, which in the approach of many authors also formalize the probability of the conditionals “if the hypothesis (conditioning event, antecedent) H is true, then the conditioned event (consequent) A happens” and “if the hypothesis K is true, then the event B happens”. Due to possible logical relations among the events A, H, B, K the previous probability assessment may be incoherent, that is by a suitable combination of two bets on $A|H$ and $B|K$ you may obtain a sure loss, in all the cases where H and K are not both false (Dutch Book).

The notion of generalized coherence (g-coherence) concerns *imprecise* probability assessments, like $0.6 \leq P(A|H) \leq 0.8$ and $0.4 \leq P(B|K) \leq 0.7$. This assessment is g-coherent if there exists a precise coherent assessment $P(A|H) = x$, $P(B|K) = y$, with $0.6 \leq x \leq 0.8$ and $0.4 \leq y \leq 0.7$. In real applications, due to vagueness or partial information, it may result difficult to specify precise values for $P(A|H)$ and $P(B|K)$; then, a more practical approach could be to start by an imprecise probability assessment. Notice that g-coherence is more general than (but, for interval-valued imprecise assessments, equivalent to) the *avoiding uniform loss* property of lower previsions studied by Peter Walley. Based on g-coherence, a generalization of the fundamental theorem of de Finetti has been given in a joint paper with Biazzo (Biazzo, Gilio 2000: A generalization of the fundamental theorem of de Finetti for imprecise conditional probability assessments, *Int. J. Approx. Reas.*, 251–272). The relationship between g-coherence and probability logic has been studied, e.g., in Biazzo, Gilio, Lukasiewicz, Sanfilippo (2002: Probabilistic logic under coherence, model-theoretic probabilistic logic, and default reasoning in System P, *Journal of Applied Non-Classical Logics*, 189–213). Moreover, the notion of g-coherence has been applied with full generality in some papers by Pfeifer & Sanfilippo to give a probabilistic interpretation of squares and hexagons of opposition and of Aristotelian syllogisms (e.g., Pfeifer, Sanfilippo 2017: Probabilistic Squares and Hexagons of Opposition Under Coherence, *Int. J. Approx. Reas.*, 282–294).

The notion of total coherence is mainly of theoretical interest. I hence I will devote a small note to it below, along some further technical remarks.

NP & GS : What are key challenges for the future of coherence-based probability theory?

AG: The coherence-based probability theory is the most flexible and general approach to the probabilistic treatment of uncertainty. In my opinion, some key challenges for its future applications are:

- to develop user friendly tools for probability elicitation based on the incomplete, vague, uncertain, and partial information available to individuals;
- to implement efficient algorithms for managing logical relations, for checking coherence and for propagation, by solving problems related with the complexity of computations;
- to develop methods for managing the relationship between qualitative and quantitative probabilities;
- to deepen the study of new concepts and methods for a suitable interpretation and formalization of human reasoning in the area of psychology, philosophy, conditional logics, machine learning and artificial intelligence; an instance of these new concepts and methods is represented by conjoined and disjoined conditionals, iterated conditionals, and compound conditionals in general (e.g., Gilio, Sanfilippo 2014: Conditional Random Quantities and Compounds of Conditionals, *Studia Logica*, 709–29; and 2019: Generalized Logical Operations among Conditional Events, *Applied Intelligence*, 79–102), with the associated aspect of determining lower and upper probability bounds for the consequents of general inference rules in probabilistic reasoning (Gilio, Pfeifer, Sanfilippo 2020: Probabilistic entailment and iterated conditionals. *Logic and Uncertainty in the Human Mind: A Tribute to David E. Over*, 71–101).

As a final remark, a relevant aspect concerns meaningful real-world applications of compound conditionals, for instance, to economics and decision theory.

NP & GS : Which paper of yours would you recommend to a student, who would like to learn more about coherence?

AG: I would recommend the paper “Probabilistic reasoning under coherence in System P” published in 2002. The first time I presented the results of this paper was in 1994, at a workshop near Paris, with Ernest Adams, Didier Dubois, Giulianella Coletti, Romano Scozzafava, Philip Calabrese, Irwin R. Goodman, Hung T. Nguyen, and Elbert A. Walker.

NP & GS : Which work of other authors (besides de Finetti) would you also suggest to read to deepen understanding of coherence?

AG: There are a lot of works by various authors that could be read; in particular, from the Italian school, I would recall: Eugenio Regazzini, Silvano Holzer, Lucio Crisma, Luciano Daboni, Pietro Rigo, Giulianella Coletti, Romano Scozzafava. Other relevant authors are Dennis V. Lindley, Leonard J. Savage, Peter M. Williams, and Frank Lad. I will limit myself to mentioning the books “Coletti G, Scozzafava R, Probabilistic logic in a coherent setting, Kluwer Academic Publishers, 2002”, and “Lad F, Operational subjective statistical methods: a mathematical, philosophical, and historical introduction, Wiley Series in Probability and Statistics, 1996”.

NP & GS : Do you remember a funny moment with de Finetti?

AG: A curious episode that I remember with pleasure is the following: we were still students or just graduated and, in a group of 5–6 people with professor Bruno Rizzi (assistant of de Finetti), we went in the summer to de Finetti’s house in the Roman countryside. We cut the grass in the ground and then we all had lunch together.

NP & GS : What is your favorite quote from de Finetti?

AG: There are a lot of interesting, meaningful and paradoxical quotes from de Finetti; but what I prefer (a sort of oxymoron) is: “Probability does not exist”.

NP & GS : Grazie per l’intervista Angelo, we will never forget the time when the two of us (and also many times before Giuseppe) visited you in Genzano di Roma, not only because of the nice and fruitful discussion on our research projects, but also because of your and your wife’s kind hospitality, wonderful cooking, and playing *briscola* and *scopa*.

A note on total coherence and g-coherence

Consider that the precise assessment $P(A|H) = x, P(B|K) = y$ is coherent, for every (x, y) such that $0.6 \leq x \leq 0.8, 0.4 \leq y \leq 0.7$. Then, the imprecise assessment $0.6 \leq P(A|H) \leq 0.8, 0.4 \leq P(B|K) \leq 0.7$ is totally coherent (Gilio, Ingrassia 1998: Totally coherent set-valued probability assessments, *Kybernetika*, 3–15). However, it may happen that for some $(x, y) \in [0.6, 0.8] \times [0.4, 0.7]$, the assessment $P(A|H) = x, P(B|K) = y$ is not coherent, when there are logical relations among the events A, H, B, K ; in this case, the previous imprecise assessment is not totally coherent. For instance, given any real numbers x_1, x_2, y_1, y_2 , with

$$0 \leq x_1 < x_2 \leq 1, 0 \leq y_1 < y_2 \leq 1,$$

the imprecise assessment $x_1 \leq P(H) \leq x_2, y_1 \leq P(A|H) \leq y_2$ is totally coherent; but it cannot be extended to AH . Indeed, as $P(AH) = P(H)P(A|H)$, for every $0 \leq z_1 < z_2 \leq 1$, the imprecise assessment

$$x_1 \leq P(H) \leq x_2, y_1 \leq P(A|H) \leq y_2, z_1 \leq P(AH) \leq z_2 \quad (1)$$

is not totally coherent because, for each $(x, y) \in [x_1, x_2] \times [y_1, y_2]$, the assessment (x, y, z) , with $z \in [z_1, z_2]$, on $\{H, A|H, AH\}$ is not coherent in general. More precisely, it is coherent if $z = xy$, while it is not coherent if $z \neq xy$.

As a consequence, there exists no extension theorem for totally coherent probability assessments.

Notice that, if there exists $(x, y, z) \in [x_1, x_2] \times [y_1, y_2] \times [z_1, z_2]$ such that $z = xy$, then the imprecise assessment in (1) is g-coherent.

In addition, imprecise g-coherent probability assessments can always be extended to further conditional events. Indeed, given any coherent probability assessment $\mathbb{P} = (p_1, \dots, p_n)$ on $F = \{E_1|H_1, \dots, E_n|H_n\}$ and a further conditional event $E|H$, there exists a suitable interval $[z', z''] \subseteq [0, 1]$ such that the extension $z = P(E|H)$ of \mathbb{P} to $E|H$ is coherent if and only if $z \in [z', z'']$ (*de Finetti’s fundamental theorem of probability and its generalizations*). Then, given an imprecise g-coherent probability assessment

$$a_i \leq P(E_i|H_i) \leq b_i, i = 1, \dots, n, \quad (2)$$

there exists a coherent precise assessment

$$a_i \leq P(E_i|H_i) = x_i \leq b_i, i = 1, \dots, n,$$

which can be coherently extended to $z = P(E|H) \in [z', z'']$. Then, for each $[z_1, z_2]$ such that $[z_1, z_2] \cap [z', z''] \neq \emptyset$, the imprecise assessment

$$z_1 \leq P(E|H) \leq z_2, a_i \leq P(E_i|H_i) \leq b_i, i = 1, \dots, n,$$

is a g-coherent extension of the assessment in (2).

Other details on coherence, g-coherence, and total coherence are explained in (Gilio, Pfeifer, Sanfilippo 2016: Transitivity in coherence-based probability logic, *Journal of Applied Logic*, 46–64).

ANGELO GIILIO

DISSEMINATION CORNER

BRIO

The project BRIO, introduced in *The Reasoner* Volume 16, Issue 1, has the purpose of achieving Trustworthy Artificial Intelligent (TAI) systems by avoiding issues of bias, risk and opacity. The importance of achieving such a goal is becoming more and more evident as the pervasiveness of autonomous decision systems increases in our everyday life. Common examples are easy to find and range from automatic credit scoring systems, to automated driving systems to recommender systems for healthcare support. The output of these systems can have a great (and possibly negative) impact on important aspects of our lives and for this reason it is of paramount importance to minimize the risks connected with their use. Most automatic decision systems rely on Machine Learning (ML) techniques and, although the level of performance reached so far is impressive, still their trustworthiness is threatened by the risks incurred by humans who interact with them within an uncontrolled environment, by the biases that could affect the collection and classification of the data used to train the ML algorithms, and by the opacity of the functioning mechanisms of the algorithms. Such threats have already been studied at length (see for instance Mehrabi, N. et al. 2021: “A Survey on Bias and Fairness in Machine Learning”, *ACM Computing Surveys* 54(6).) and tools like FairML are being developed to address them, but efforts to operationalize this knowledge so as to identify and combat threats already at the design stage are still missing. The accomplishment of this step is what motivates both the interdisciplinary approach embraced by the BRIO consortium and the structure of the specific objectives and activities of the project.

Already at the beginning of the 90’s Marvin Minsky (1991: “Logical Versus Analogical or Symbolic Versus Connectionist or Neat Versus Scruffy”, *AI Magazine* 12(2)) has pointed out how the integration of connectionist (like ML) and symbolic approaches could represent a promising hallway to pursue towards the solution of such problems. More recently, Marinucci L. et al. (2022: “Exposing implicit biases and stereotypes in human and artificial intelligence: state of the art and challenges with a focus on gender”, *AI & SOCIETY*) have highlighted the role that knowledge representation frameworks and formal ontologies in particular can have in reducing bias and opacity in AI systems.

In BRIO we propose formal ontologies as a means to explicitly represent important aspects of the development of AI systems, by focussing in particular on the input data provided and classified by humans as well as on the clear specification of the semantics of the output of the system. Making the assumptions and the objectives of an AI system precise is crucial for two reasons: firstly, it enables the users to understand at least what the system is supposed to do; secondly, it enables the designer to identify the kinds of risks and biases that may affect such system and to indicate the related possible negative

consequences.

In general terms, formal ontologies are computational artifacts that are built to make explicit the hidden assumptions behind the use of terms and notions in the model or representation of a system. Hence, they have the advantage of being both machine-readable and understandable by humans. A fundamental prerequisite to build a formal ontology is a phase of conceptual analysis, in which various definitions taken from many disciplines (cognitive science, epistemology, law) are thoroughly studied, ideally complemented with interviews to experts who can guide the modeler in understanding the core issues to be addressed with the help of the ontology. A further preliminary step is the construction of a taxonomy of the main notions involved, resulting from the conceptual analysis. Finally, the properties of the main entities to be represented, together with their relations, are expressed through the use of formal axioms, i.e. constraints written in some suitable language, like OWL (Web Ontology Language).

The construction of a comprehensive formal ontology is the main activity of Objective 2, to be carried out by the [Laboratory for Applied Ontology](#) of the [Institute of Cognitive Sciences and Technologies of the CNR](#), in tandem with the [DAFIST \(Department of Antiquity, Philosophy and History\)](#) of the University of Genova.

In the scope of BRIO, the conceptual analysis will complement the outcomes of the activities performed under Objective 1 (which has been described and discussed in [Volume 16, Issue 3](#) of the Reasoner). The main task of Objective 2 is to offer a systematic characterization of the types of bias identified in the literature and, guided by the results of Objective 1 and by well-known methodologies like OntoClean (Guarino, N. and Welty, C.A. 2009: “[An Overview of OntoClean](#)”, in *International Handbooks on Information Systems*, Springer, pp. 201–220), to make them viable for formal and automatic identification. A second aspect of this task is the definition of risks involved in the construction and use of possibly biased complex AI systems. The identification of possible sources of risk in the design phase has the purpose of avoiding harm for humans interacting with AI. A starting point for the construction of the ontology are already existing characterizations of risk, trust and connected notions (Amaral, G. et al. 2019: “[Towards a Reference Ontology of Trust](#)”, in *On the Move to Meaningful Internet Systems*, Springer, pp. 3–21) using foundational ontologies like DOLCE (Masolo C. et al. 2003: “[WonderWeb Deliverable D18](#)”; Borgo, S. and Masolo, C. 2009: “[Foundational Choices in DOLCE](#)”, in *Handbook on Ontologies*, Springer, pp. 361–381; Borgo, S. et al. 2022: “[DOLCE: A descriptive ontology for linguistic and cognitive engineering](#)”, *Applied Ontology* 17(1): 45–69) and UFO (Guizzardi, G. et al. 2022: “[UFO: Unified Foundational Ontology](#)”, *Applied Ontology* 17(1): 167–210), which will be extended and systematized in a formal framework, in order to facilitate the application to decision systems.

The formal ontology resulting from the activities in Objective 2 will be combined with logical methods, formal specification and verification methods, that will be developed in Objective 3. Stay tuned!

ROBERTA FERRARIO

Institute of Cognitive Sciences and Technologies – CNR

DANIELE PORELLO

University of Genoa

Statistical Relational AI

One of the great triumphs of classical formal logic is the insight that full first-order logic is undecidable. However, over the past hundred years, several decidable fragments have been isolated. Originally, the focus of those endeavours lay on prefix classes, where the alternation of existential and universal quantifiers is restricted to ensure decidability. Already by the 1930s, the classical solvable fragments have been found: The prefix class $\exists^*\forall^*$ in which no existential quantifier is allowed in the scope of a universal quantifier, and the prefix class $\exists^*\forall\exists^*$, in which any literal may not be in the scope of more than a single universal quantifier.

It took until 1975, however, for another very natural fragment to be discovered: Two-variable logic FO₂, the fragment of first-order logic in which only two different variables are allowed to occur. In some ways, this is very limiting — for instance, it only allows for binary relation symbols to be used meaningfully. On the other hand, by reusing the two variables, one can express many queries that are relevant to graph-theoretic applications. This includes questions about the existence and non-existence of edges between two nodes, but also by reusing variables connectedness of two nodes in less than n steps for any n .

Recently, the two-variable fragment has been rediscovered for statistical relational AI. In the March edition of this column, we gave a brief introduction to lifted probabilistic inference and weighted model counting, to which the interested reader is referred for details on the background. In essence, as inference from statistical relational frameworks is usually P#-complete (which is just worse than NP-complete), one looks for fragments of those frameworks in which inference tasks can be computed in polynomial time. Although there are many different types of statistical relational model, from probabilistic logic programs to Markov Logic Networks, common inference tasks can be reduced to *weighted model counting*, computing the sum of the weights of models of a first-order formula. Therefore, fragments of first-order logic that support tractable weighted model counting translate to corresponding fragments of statistical relational frameworks that support tractable inference.

Against this backdrop weighted model counting was shown to be tractable for universal first-order formulas with 2 variables in 2011, followed by tractability of weighted model counting for unrestricted FO₂ in 2014. This was recently extended by accommodating counting quantifiers (quantifiers of the form “there are at most n x such that”).

At this year’s annual AAAI conference, one of the most prestigious AI conferences on the circuit, FO₂ was again a focus of research on lifted statistical relational methods. Sagar Malhotra and Luciano Serafini seem to have given the weighted model counting for FO₂ with counting quantifiers its final form, as they provide a closed-form solution for this setting. This is a closed formula with variables for the sizes of domain sorts that can be computed directly from a finite first-order theory and the weights. Since the domain size there appears only as summation index, binomial coefficient or exponent, the actual weighted model counting task on a given domain reduces immediately to such arithmetic calculations, for which very fast algorithms and implementations are known.

Besides exact inference, a key operation on a statistical re-

lational framework is sampling from its induced distribution. This has a variety of different applications. For instance, it allows the user to specify a probabilistic model and then sample possible worlds for benchmarking and testing of algorithms. Or it can be used in dedicated simulation environments to model disease dynamics or other progressive situations, in which the changes at every timepoint are obtained by sampling from a distribution. Sampling is also used to implement approximate inference, used where exact inference is infeasible in practice or when the distribution is combined with continuous random variables, for which only sampling-based inference has been defined. Despite its widespread use in applications, sampling is theoretically somewhat understudied, and I am not aware of any explicit complexity guarantees for sampling from statistical relational frameworks.

It is very welcome therefore that Yuanhong Wang, Timothy van Bremen, Yuyi Wang, Ondrej Kuzelka showed that for universal first-order logic with 2 variables (and possibly with counting quantifiers), both uniform and weighted sampling is tractable; much as with inference and weighted model counting, sampling from statistical relational formalisms can often be reduced to weighted sampling from a first-order theory.

Of course, there is a clear gap between the two results presented here: While weighted model counting has been shown tractable for *full* two-variable logic, sampling has only been shown to be tractable for *universal* first-order logic. Wang et al. comment on this in their paper, showing that a direct application of their algorithm to formulas with existential quantifiers is intractable (modulo $FP = \#P$). They conclude that completely different methods would be necessary to address the more general problem, and interestingly they are cautious in conjecturing tractability at least of full two-variable logic with counting; such a result would indeed subsume partial results on the uniform generation of regular random graphs that themselves used quite advanced techniques.

So, just as the door seems to close on the FO_2 counting problem, another is flung open.

FELIX WEITKÄMPER
Computer Science, LMU Munich

Uncertain Reasoning

The Rényi–Ulam game, a variation of the Twenty question game, was independently introduced by the Hungarian mathematician Alfréd Rényi (1961: On a problem in information theory, *Magyar Tud. Akad. Mat. Kutató Int. Közl.* (in Hungarian), **6**: 505–516) and the Polish-American scientist Stanislaw Ulam in his autobiography *Adventures of a mathematician* in 1976 (p. 281). The game consists of two players, Questioner (Q) and Responder (R). R thinks of a number, referred to as “the secret”, between one and one million (the *search space*), which is just less than 2^{20} . Q aims at guessing the secret only by using questions whose answer is either yes or no, e.g. “Is the secret an even number?”. If R is forced to answer truthfully, then Q can guess the secret in less $\log_2(1000000)$ questions (which can be approximated to 20) by applying the same principle of the bisection method. In the Rényi–Ulam game, R can lie a fixed number of times (m) known to both players.

In the Twenty questions game, once a pair of questions and answers excludes the number x of the *search space* from being the secret (we say that x has been *rejected*), x can no longer be

considered a valid candidate for being the solution of the game. Instead, in the Rényi–Ulam game with m lies, if a number x has been rejected only once, then it is still possible for x to be the secret. In order to finally exclude x we need at least $m+1$ pairs of questions and answers that reject it. In this case, we say that x has been *finally rejected*.

In the execution of the game, after every pair of questions and answers, we can depict the state of the game by computing for every number in the search space its distance in units of $1/(m+1)$ from becoming finally rejected. A candidate x rejected by $n < m+1$ pairs of questions and answers will have a distance (from becoming finally rejected) of $1-n/(m+1)$. If x has been rejected more or at least $m+1$ times, then it will have distance 0 and if x has never been rejected, then it will have distance 1. Thus, before start playing, the distance of each candidate to become finally rejected is 1. The game ends whenever all the candidates except one have distance 0 from becoming finally rejected, i.e. every candidate *is* finally rejected except for one. The only number whose distance is strictly greater than 0 is the secret. At the end of the game, the distance of the secret from being finally rejected is not necessarily 1. In fact, it could be that there are pairs of questions and answers that reject it, but the number of these pairs cannot exceed m . Thus, if both players follow the rules of the game, the minimum distance the secret can have from becoming finally rejected is $1/(m+1)$.

At every stage of the game, after having computed the distance of each candidate from becoming finally rejected, we can identify the candidates that are closer to being finally rejected as those with minimum distance. By looking at these sets of candidates that are entailed from the sequences of questions and answers, it is possible to highlight the non-monotonicity of this entailment. In fact, it could be that after a sequence of questions and answers, a number x belongs to the set of candidates closer to being finally rejected while considering a longer sequence, the same number x could no longer be entailed. The non-monotonicity of this entailment lasts until a candidate is finally rejected. In this case, the set of numbers closer to being finally rejected has value 0, the minimum value a candidate can be assigned to. This means that they have been rejected at least $m+1$ times. Thus, in the continuation of the game, the pairs of questions and answers that reject candidates whose value is already 0, can only increase and their value will continue to be 0.

Whenever R does not answer truthfully, she either rejects a set of candidates that should not have been rejected or fails to reject a set of candidates that should actually have been rejected. Let us consider the search space $[1, 1000]$. Suppose that the secret is the number 8 and that Q asks the question “Does the secret belong to the set $[1, 500]$?” R should answer YES; but if for any reason she lies, and answers NO, then she is erroneously rejecting the elements of $[1, 500]$. This situation can be related to a Type I error in statistical hypothesis testing. Then, if Q asks the question “Does the secret belong to $[501, 1000]$?”, R should answer NO, but if -lying- answers YES, then she is failing to reject the set of candidates $[501, 1000]$ and we occur in a case of false-negative. The situation just described relates to Type II errors.

The similarities between statistical analysis and the Rényi–Ulam game are not related only to different situations in which R can lie, but also to the development of the game itself. In the Rényi–Ulam game, a sequence of questions and answers is more informative than another sequence of questions and an-

swers whenever the former rejects more candidates than the latter one. In the hypothesis testing method, after the null hypothesis has been rejected, then the alternative hypothesis is accepted. Similarly, in the Rényi–Ulam game, after all the candidates except for one have been finally rejected, then the only candidate not finally rejected is identified as the secret.

As proved by Mundici in (1992, The logic of Ulam’s game with lies. Knowledge, belief and strategic interaction, pages 275–284) and (Ulam games, Łukasiewicz logic, and (AF) C*-algebras. *Fundamenta Informaticae*, 18(2-4):151–161), the Rényi–Ulam game with a fixed number of lies is a sound and complete semantics for Łukasiewicz logic. Using the connections briefly sketched above, it would be interesting to investigate the role of fuzzy logic in an alternative formalisation of the hypothesis testing method and, more in general, in statistical inference.

ESTHER ANNA CORSI
LUCI Group, University of Milan

Multidisciplinary Reasoning

We’ll continue a discussion considering the debate between pessimism and optimism in cultural studies information-theoretically. Utilizing information theory, we were able to analyze this debate as a physical system. (Peterson II, Victor: 2022 *The Reasoner* Volume 16, Number 3. What constitutes a pessimist channel is what’s determined as not contributing information but is indirectly proven to exist by the information it contextualizes. Whatever information is received through this channel is, with absolute certainty, treated cynically, regardless of source. With certainty information is destroyed. This information paradox has been a cornerstone in physics and complexity theory.



This information paradox has been explored in physical terms before. (Hawking, S.W.: 1975, “Particle Creation by Black Holes,” *Commun. Math. Phys.*, 43, 199-220) If we suppose that nothing can travel faster than light, then a signal travelling at that speed or less generates no new information. If the volume of the universe can be calculated as the integral of the circumference of a sphere surrounding entities at various distances from a central point, then a black hole is a sphere within that volume that has infinite radius. As such, black holes seemingly destroy information, not even light escapes. However, if a photon falls into a hole that’s by definition a vacuum, the energy it carries should wholly transfer into the state circumscribed by that sphere. Yet, as black holes are defined as having a negative infinite radius, the fall to the surface never fully occurs, leaving a trace. (S. Haco, S. W. Hawking, M. Perry, A. Strominger: 2018, “Black Hole Entropy and Soft Hair,” arXiv:1810.01847) This is illuminating for the cynic must maintain (=record) that which allows it to *know* that the state is the way they’ve determined so to maintain their own position. The cynic cannot prove their cynicism as they would have to produce an entity from that domain to show it’s empty. However, the function of expressing some entity, if identical to that entity, means that its function of expression is an object in itself, showing that that domain is not empty after all.

Another example is that of interest convergence. (Bell Jr., Derrick: 1980, “Brown v. Board of Education and the Interest-Convergence Dilemma,” *Harvard Law Review*, 93:3, 518-533) This theory states that convergent interests lead to the maintenance of the channel, not the information, thereby reinforcing what’s understood as the norm or dominant position. Cynicism is understood by absolutism of frame. A channel, by definition, is the sum alternatives that don’t contribute information, setting the boundaries through which information flows, i.e. what’s fixed relative to what varies thereby contextualizing what can be known. Convergence converts everything flowing through a channel to the same information. Independent alternatives are combined to fit either a category that maintains the flow of information in the same direction or to a category that is fed back into the system.

If X is the category to be returned and Y to be retained, with $I(Z)$ contributing to X , and $f(Z)=\prod_{i=1}^n p(z_i \dots z_n)^{Np(x_n)} = \sum_{i=1}^n p(x_i) \log p(z_i)$, with what’s X expressed as a function of N occurrences, then

1. If $X=f(Z)$, then $I(f(Z) \text{ given } Z)=0$, and $I(Z, f(Z))$ evaluates X to $I(Z)+I(f(Z) \text{ given } Z)=I(Z)+I(Z \text{ given } f(Z))$, meaning that $I(f(Z)) \leq I(Z)$, diminishing the content of X to null before it produces output.

With f being recursive, Z diminishes towards 0 as it’s returned once through the channel and multiplied over X^N many iterations. The output of the channel is maintained insofar as

$$2. C(X, Y) = I(Y) + I(X) = \sum_y p(y) \log p(y) + \sum_{i=1}^n p(x_i) \log p(z_i)$$

where to the right of the addition sign, the first term of the product sum represents the probability of categorization and the second the information produced/predicated. X is received although Z was sent. Alternatives arise but are ignored. If accepted, the information contributes (=additive); if not, it’s recursively coupled with the unacceptable (=multiplied). The sum of the combination of probable alternatives reduces to 0 after a certain number of cycles. The channel annihilates information that doesn’t contribute to the current structure of the state, deeming certain sources as nonexistent regardless of what’s transmitted. This nihilistic outcome is derived from the need to evacuate a source of content upon arrival so as to maintain the current state. However, if 0 is received, this doesn’t mean that nothing was sent, in fact, we find that that is due to the channel, not the source.

This interpretation grounds Afropessimism, a term originating in International Relations that described the futility of investments in Africa. (Hawkins, Tony: 1990, “Black Africa: from aid dependence to self-sustaining growth,” *The World Today*, 46(11), 205–208.; Hitchens, Christopher: 1994, “Africa without Pity,” *Vanity Fair*, 43–52.) It’s now conflated with the impossibility of Blacks contributing to global socio-political affairs. (Wilderson III, Frank: 2020, *Afropessimism*, Liveright) However, we do find evidence of culturally and politically significant transmission despite its negation upon arrival. (Mineo L. and Patterson O.: 2018, “The Kerner Report on race, 50 years on,” *The Harvard Gazette*, <https://news.harvard.edu/gazette/story/2018/03/harvard-professor-reflects-on-the-kerner-report-50-years-on/>) It’s the historical formation and current structure of the channel—slavery, colonization, appropriation—and not the source that

EVENTS

is the blame. As a 0% possibility is undefined information-theoretically, it is not necessarily the case that Black-ness has no import, only that it's necessarily barred from that channel and, possibly, is integral to it if it is the case that that probability does not vary given changing conditions. With 0% possibility being undefined information-theoretically, the limits of this archaic channel indirectly proves the necessary possibility of alternative channels.

Reduction to statistical combinations of behavior as a proof of their cause denies the historical formation of structures projected to organize, allowing or disallowing, expressions as a function of prior output. The pessimist channel decides what doesn't contribute information, only attributing existence to what passes. Cynicism emerges. Pessimism cannot *know* its object is void. Their disposition would be that of a cynic, then, not a pessimist. We find that nothing is certain, even the certainty that everything means nothing. Channel fundamentalism results from the cynicism of the pessimist deciding nothing is to be gained from a vested interest in the very subject they analyze. Information-theoretically, the cynicism of the pessimist, and its resulting nihilism, are not default positions; they emerge given the channel. Attend the channel, alternatives arise.

VICTOR PETERSON II
New York University

JULY

IARML: Interactions between Analogical Reasoning and Machine Learning, Vienna, Austria, 23 July.

IRSI3: 3rd International Rationality Summer Institute, Landau, Germany, 24 July–5 August.

PCCR: Parameterized Complexity of Computational Reasoning, Haifa, Israel, 31 July–1 August.

KR: Principles of Knowledge Representation and Reasoning, Haifa, Israel, 31 July–5 August.

AUGUST

TIS: Trust in Science, Stuttgart, Germany, 3–5 August.

BMA: Bayesian Modelling Applications, Eindhoven, Netherlands, 5 August.

CAUSREP: Causal Representation, The Netherlands, hybrid, 5 August.

NMR: Non-Monotonic Reasoning, Haifa, Israel, 7–9 August.

TSG&P: Truthmaking, Semantical Grounding, and Paradoxes, University of Bristol, 22–23 August.

SCIProg: Scientific Progress: Individual and Collective, Amsterdam, 24–26 August.

SEPTEMBER

'TRUE': The Meaning(s) of 'True,' University of Bristol, 5–6 September.

FoCoRE: Formal and Cognitive Reasoning, Trier, Germany, 20 September.

E&EiS: Engaging Ethics and Epistemology in Science, Leibnizhaus, Hannover, Germany, 29–30 September.

COURSES AND PROGRAMMES

Courses

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

LAIS: Logic for the AI Spring, 12–16 July.

Programmes

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

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

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

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

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

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

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

MASTER PROGRAMME: in Statistics, University College Dublin.



JOBS AND STUDENTSHIPS

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

MASTER PROGRAMME: in Artificial Intelligence, Radboud University Nijmegen, the Netherlands.

MASTER PROGRAMME: Philosophy and Economics, Institute of Philosophy, University of Bayreuth.

MA IN COGNITIVE SCIENCE: School of Politics, International Studies and Philosophy, Queen's University Belfast.

MA IN LOGIC AND THE PHILOSOPHY OF MATHEMATICS: Department of Philosophy, University of Bristol.

MA PROGRAMMES: in Philosophy of Science, University of Leeds.

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

MA IN LOGIC AND THEORY OF SCIENCE: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.

MA IN METAPHYSICS, LANGUAGE, AND MIND: Department of Philosophy, University of Liverpool.

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

MA IN PHILOSOPHY: by research, Tilburg University.

MA IN PHILOSOPHY, SCIENCE AND SOCIETY: TiLPS, Tilburg University.

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

MA IN RHETORIC: School of Journalism, Media and Communication, University of Central Lancashire.

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

MRES IN METHODS AND PRACTICES OF PHILOSOPHICAL RESEARCH: Northern Institute of Philosophy, University of Aberdeen.

MSc IN APPLIED STATISTICS: Department of Economics, Mathematics and Statistics, Birkbeck, University of London.

MSc IN APPLIED STATISTICS AND DATAMINING: School of Mathematics and Statistics, University of St Andrews.

MSc IN ARTIFICIAL INTELLIGENCE: Faculty of Engineering, University of Leeds.

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

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

MSc IN COGNITIVE SCIENCE: University of Osnabrück, Germany.

MSc IN COGNITIVE PSYCHOLOGY/NEUROPSYCHOLOGY: School of Psychology, University of Kent.

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

MSc IN MIND, LANGUAGE & EMBODIED COGNITION: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.

MSc IN PHILOSOPHY OF SCIENCE, TECHNOLOGY AND SOCIETY: University of Twente, The Netherlands.

MRES IN COGNITIVE SCIENCE AND HUMANITIES: LANGUAGE, COMMUNICATION AND ORGANIZATION: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastián).

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

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

Studentships

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

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

Jobs

ASSOCIATE PROFESSORSHIPS: in Philosophy of Language/Logic and Epistemology, University of Oslo, deadline 31 August.

