Progic 2025:
The 12th Workshop on Combining Probability and Logic

Center for Formal Epistemology,
Carnegie Mellon University

April 3-5, 2025
Baker Hall A36 (Adamson Wing)

Workshop Program and Abstracts

Thursday, April 3, 2025

8:30am-9:00am	Registration and Breakfast
9am-9:15am	Opening Remarks
9:15am-10:45am	Plenary Talk **Dan Roy** (University of Toronto) **Admissibility is Bayes Optimality with Infinitesimals**
10:45am-11am	Coffee Break
11am-11:45am	**Nate Ackerman** (Harvard), **Cameron Freer** (MIT), **Rehana Patel** (Wesleyan) **Learning Invariant Measures**
11:45am-12:30am	**Jeff Barrett** (UC Irvine) and **Christian Torsell** (UC Irvine) **Learning How to Learn (by Reinforcement)**
12:30pm-2pm	Lunch Break
2pm-2:45pm	**Jürgen Landes** (Munich Center for Mathematical Philosophy) **On the Value of Varied Evidence for Imprecise Probabilities**
2:45pm-3:00pm	Coffee Break
3pm-3:45pm	**Jon Williamson** (University of Manchester) **The Heuristic Use of Conditionalisation**
3:45pm-5pm	Reception

Friday, April 4, 2025

8:30am-9am	Breakfast
9am-10:30am	Plenary Talk **Sean Walsh** (UCLA) **The Expressive Power of Counterfactuals and Descriptive Set Theory**
10:30am-10:45am	Coffee Break
10:45am-11:30am	**Nate Ackerman** (Harvard), **Cameron Freer** (MIT), **Dan Roy** (Toronto) **Computability of Properties of Stochastic Processes Used for Learning**
11:30am-12:15pm	**Sam Eisenstat** (Machine Intelligence Research Institute) **Logical Uncertainty and Self-Reference in Garrabrant Induction**
12:15pm-2pm	Lunch Break
2pm-2:45pm	**Giovanni Duca** (Northeastern/University of Milan) **Updating on Uncertain Evidence: a correspondence between belief revision and Jeffrey conditioning**
2:45pm-3pm	Coffee Break
3pm-3:45pm	**Milan Mossé** (UC Berkeley) **Reasoning about Confirmation**
3:45pm-5pm	Reception
6:30pm	Conference Dinner

Saturday, April 5, 2025

8:30am-9am	Breakfast
9am-10:30am	Plenary Talk **Kun Zhang** (Carnegie Mellon University) **Causal Learning: Why It Matters, How It Works, and Its Implications for Generative AI**
10:30am-10:45am	Coffee Break
10:45am-11:30am	**Sander Beckers** (Cornell University) **Nondeterministic Causal Models**
11:30am-12:15pm	**Aydin Mohseni** (Carnegie Mellon University) **Naturalizing Causation in Causal Models**
12:15pm-2pm	Lunch Break
2pm-2:45pm	**Siddharth Namachivayam** (Carnegie Mellon University) **Topological Semantics for Common Inductive Knowledge**
2:45pm-3pm	Coffee Break
3pm-3:45pm	**Adam Bjorndahl** (Carnegie Mellon University) **A Logic of Uncertain Interpretation**
3:45pm-4pm	Concluding Remarks
4pm-5pm	Reception

Progic 2025

Center for Formal Epistemology, Carnegie Mellon University

April 3–5, 2025

Abstracts

Dan Roy (University of Toronto)

Title: Admissibility is Bayes Optimality with Infinitesimals

Abstract: We give an exact characterization of admissibility in statistical decision problems in terms of Bayes optimality in a so-called nonstandard extension of the original decision problem, as introduced by Duanmu and Roy. Unlike the consideration of improper priors or other generalized notions of Bayes optimality, the nonstandard extension is distinguished, in part, by having priors that can assign "infinitesimal" mass in a sense that can be made rigorous using results from nonstandard analysis. With these additional priors, we find that, informally speaking, a decision procedure δ_0 is admissible in the original statistical decision problem if and only if, in the nonstandard extension of the problem, the nonstandard extension of δ_0 is Bayes optimal among the extensions of standard decision procedures with respect to a nonstandard prior that assigns at least infinitesimal mass to every standard parameter value. We use the above theorem to give further characterizations of admissibility, one related to Blyth's method, one to a condition due to Stein which characterizes admissibility under some regularity assumptions; and finally, a characterization using finitely additive priors in decision problems meeting certain regularity requirements. Our results imply that Blyth's method is a sound and complete method for establishing admissibility. Buoyed by this result, we revisit the univariate two-sample common-mean problem, and show that the Graybill–Deal estimator is admissible among a certain class of unbiased decision procedures.

Joint work with Haosui Duanmu and David Schrittesser.

Link to preprint: https://arxiv.org/abs/2112.14257

Nate Ackerman (Harvard), **Cameron Freer** (MIT), **Rehana Patel** (Wesleyan)

Title: Learning Invariant Measures

Abstract: We consider the class $\mathbb{S}_{\mathfrak{L}}$ of measures on $2^{\mathbb{N}}$ expressible as the weak limit of a sequence of measures obtained by sampling from finite $\mathfrak{L}$-structures (where $\mathfrak{L}$ is a fixed finite relational language). Such measures are the ergodic S_∞-invariant probability measures on the space of $\mathfrak{L}$-structures with underlying set $\mathbb{N}$. We study when a measure in $\mathbb{S}_{\mathfrak{L}}$ can, with high probability, be approximately learned from a fixed number of samples. This depends on the notion of approximation, and we study both a metric capturing weak convergence and the total variation metric.

Note that an arbitrary probability measure on $2^{\mathbb{N}}$ can, with high probability, be approximately learned with respect to weak convergence from initial segments of size k_ϵ from n_ϵ-many samples, where k_ϵ and n_ϵ depend only on the desired accuracy ϵ. We show that any measure in $\mathbb{S}_{\mathfrak{L}}$ can, with high probability, be approximately learned with respect to weak convergence from an initial segment (whose size depends only on ϵ) of a *single* sample. In contrast, we show that there is no way, with high probability, to approximately learn measures in $\mathbb{S}_{\mathfrak{L}}$ with respect to the total variation metric from the set of initial segments of finitely many samples. However, any measure in $\mathbb{S}_{\mathfrak{L}}$ can, with probability 1, be learned from the full contents of a single sample.

Jeff Barrett (UC Irvine) and **Christian Torsell** (UC Irvine)

Title: Learning How to Learn (by Reinforcement)

Abstract: David Hume argued that we lack rational justification for our beliefs regarding matters of fact. In his skeptical solution to the problem of induction, he further argued that such beliefs are learned, not justified. But there is no canonically best way to learn in all contexts. We are concerned here with how an agent might learn how to learn more reliably in the context of a particular type of learning problem. The issue is not one of rationally justifying a particular form of learning. Rather, it is one of how nature has equipped us, and other animals, with the ability to learn how to learn in ways that are well suited to the problems we face as we interact with each other and the world.

Jürgen Landes (Munich Center for Mathematical Philosophy)

Title: On the Value of Varied Evidence for Imprecise Probabilities

Abstract: It has long been considered a truism that we can learn more from a variety of sources than from highly correlated sources. This truism is captured by the Variety of Evidence Thesis. To the surprise of many, this thesis turned out to fail in a number of Bayesian settings. In other words, replication can trump variation. Here we offer two formulations of this thesis with imprecise probabilities. We show that the formulations fail in a number of cases. Overall, we find that whether more or less varied evidence grants a greater reduction in uncertainty subtly depends on parameter values capturing prior uncertainty.

Joint work with Sebastien Destercke.

Jon Williamson (University of Manchester)

Title: The Heuristic Use of Conditionalisation

Abstract: This paper argues that Bayesian conditionalisation should not be viewed as a universal norm of updating, but rather as a heuristic principle that is helpful in some circumstances but fails in others. I go on to show that the heuristic use of conditionalisation can be validated by an approach to inductive logic that appeals to the maximum entropy principle, namely objective Bayesian inductive logic (OBIL). This argument has far-reaching consequences. In particular, Bayesian conditionalization should not be thought of as constitutive of Bayesianism; instead, Bayesianism needs to be grounded in a more general framework, such as OBIL, that can provide an account of the scope of conditionalisation.

Sean Walsh (UCLA)

Title: The Expressive Power of Counterfactuals and Descriptive set Theory

Abstract: We look at the complexity of counterfactuals through the classification systems provided by classical and effective descriptive set theory. This provides a way of understanding the expressive gains made by the traditional semantics for counterfactuals when one takes the distance to be a literal distance provided by a metric (on a completely metrizable space, that is, a Polish space). Since the space of probability measures is a Polish space which metrizes the topology of weak convergence, one can also take this to apply to settings where one considers minimal alternations to one's current credences.

Nate Ackerman (Harvard), **Cameron Freer** (MIT), **Dan Roy** (Toronto)

Title: Computability of Properties of Stochastic Processes Used for Learning

Abstract: Jackson, Kalai, and Smorodinsky 1999 give necessary and sufficient conditions for when a measure, conditioned on its tail σ-algebra, almost surely gives a measure that is both sufficient for prediction and merges with the original. As a first step towards analyzing the computable content of this result, we study the computational strength of the notions of "sufficient for prediction" and "merging of measures". We find that each of these classical notions bifurcates into two distinct computable versions based on the chosen notion of convergence.

Sam Eisenstat (Machine Intelligence Research Institute)

Title: Logical Uncertainty and Self-Reference in Garrabrant Induction

Abstract: Garrabrant induction is a model of inductive inference in which agents are not logically omniscient, and instead express uncertainty about propositions like those of mathematics or about the observational consequences of theories using quantitative degrees of belief. In this paper, we interpret Garrabrant induction as a normative epistemic theory, and examine what it has to say to say about a number of different aspects of epistemology. It makes weaker coherence demands upon agents than typical versions of Bayesian epistemology, in that it does not demand that agents know consequences of their beliefs to be true, but this opens up a space for it also to make new kinds of demands. We demonstrate what Garrabrant induction has to say about problems including epistemic self-reference, old evidence, merging of opinions between different subjective epistemic states, and inductive learning of universally quantified generalizations. We argue that these behaviors are justified; faced with uncertainty about the logical consequences of one's beliefs, one ought to reason as a Garrabrant inductor does.

Giovanni Duca (Northeastern/University of Milan)

Title: Updating on Uncertain Evidence: a correspondence between belief revision and Jeffrey conditioning

Abstract: Hanti Lin and Kevin Kelly introduce a way of relating probability functions and binary beliefs, referred to as odds-threshold method, which is based on comparing the probabilities of the elementary outcomes: one outcome is preferred to another if it is sufficiently more likely that the latter; the agent's belief set is the set of mostly preferred outcomes. As

main contribution, we give a probability-independent definition of the preferential orders generated by the odds-threshold method. This allows for a logical characterization of the belief revision properties induced on binary beliefs by Jeffrey conditioning (JC) on a piece of evidence. Specifically, we show that if JC is assumed to be successful—namely, the evidence is implied by the belief set induced by the posterior probability—then the corresponding belief revision on binary beliefs satisfies all the rules of the nonmonotonic system P together with Disjunctive Rationality and a new rule named Underminer Monotonicity. As JC does not make the evidence certain, we also considered rules for iterated belief revision. In the case of successful JC, the induced iterated belief revision satisfies the rules (C1), (C2) and (C3), proposed by Darwiche and Pearl. In addition, if JC is also assumed to be confirmatory—i.e. the posterior probability of the evidence is strictly greater than its prior—then the rule (C4) is satisfied as well.

Milan Mossé (UC Berkeley)

Title: Reasoning about Confirmation

Abstract: When does evidence confirm a hypothesis? On Hempelian and hypothetico-deductive views, confirmation is a deductive relation between evidence, a hypothesis, and some auxiliary assumptions. On probabilistic views, evidence confirms a hypothesis by raising its probability past its prior likelihood (the so-called "relevance" view) or above a fixed threshold of "firmness." These views provide alternative interpretations of a simple formal language of confirmation, built from Boolean combinations of the statements "H is possible" and "E confirms H". We investigate the difficulty of reasoning about confirmation, understood as the complexity of satisfiability for the resulting logics of confirmation.

Kun Zhang (Carnegie Mellon University)

Title: Causal Learning: Why It Matters, How It Works, and Its Implications for Generative AI

Abstract: Causality is a fundamental notion in science, engineering, and even in machine learning. Specifically, for the purpose of scientific discovery, we aim to develop a platform that utilizes all available measured data as input to generate hypotheses about hidden entities and causal influences. This will continually expand human knowledge and make humans more resourceful. How can we achieve it? Causal discovery (or causal representation learning) involving causally-related hidden variables plays a pivotal role. In this talk, I will report our

recent algorithmic advances in this endeavour, along with their applications to real problems in psychology, machine learning, and computer vision. I will demonstrate how the "modularity" property of causal systems, paired with suitable simplicity assumptions, makes it possible to recover the underlying causal process from observational data. Finally, I will demonstrate how identifiable causal representation learning can naturally benefit generative AI, with image generation, editing, and refinement as an example.

Sander Beckers (Cornell University)

Title: Nondeterministic Causal Models

Abstract: I generalize acyclic deterministic structural causal models to the nondeterministic case and argue that it offers an improved semantics for counterfactuals. The standard, deterministic, semantics developed by Halpern (and based on the initial proposal of Galles & Pearl) assumes that for each assignment of values to parent variables there is a unique assignment to their child variable, and it assumes that the actual world (an assignment of values to all variables of a model) specifies a unique counterfactual world for each intervention. Both assumptions are unrealistic, and therefore I drop both of them in my proposal. I do so by allowing multi-valued functions in the structural equations. In addition, I adjust the semantics so that the solutions to the equations that obtained in the actual world are preserved in any counterfactual world. I provide a sound and complete axiomatization of the resulting logic and compare it to the standard one by Halpern and to more recent proposals that are closer to mine. Finally, I extend these models to the probabilistic case and show that they open up the way to identifying counterfactuals even in Causal Bayesian Networks.

Aydin Mohseni (Carnegie Mellon University)

Title: Naturalizing Causation in Causal Models

Abstract: The ontological status and explanatory role of causation have been perennial puzzles. In recent work, Pearl and Mackenzie (2018) advance the thesis of a causal hierarchy (PCH) and posit the irreducibility of causal claims to merely probabilistic ones. Bareinboim et al. (2022) claim to have proven this irreducibility in the context of structural causal models (SCMs). We challenge this claim and demonstrate a general reduction of interventional propositions to probabilistic ones within the same context of SCMs.

Joint work with Ben Levinstein, Daniel Herrmann, Gerard Rothfus and Bruce Rushing.

Siddharth Namachivayam (Carnegie Mellon University)

Title: Topological Semantics for Common Inductive Knowledge

Abstract: Lewis' original account of common knowledge in Convention describes the process which drives the generation of higher-order expectations as hinging upon agents' inductive standards. Subsequent formal treatments of common knowledge such as those in Aumann 1976 and Geanakoploks & Polemarchakis 1982 were purely deductive and took as primitives agents' information partitions over possible worlds. This paper attempts to draw from insights in learning theory to provide a formal account of common inductive knowledge. We follow Kelly's approach in The Logic of Reliable Inquiry and take as primitives agents' information bases to formulate a Byzantine generals-like learning game where players must eventually coordinate on reporting a proposition P is true within a fixed budget of mind switches. This leads us to analyze a structure we call the multitopological difference lattice which generalizes the topological difference hierarchy first touched upon by Putnam's 1965 paper about trial and error predicates. Namely, we provide a topological semantics for common inductive knowledge which precisely captures the epistemic content of the multitopological difference lattice and, in the deductive case, also coincides with the notion of asynchronous common knowledge recently introduced by Gonczarowski & Moses 2024.

Adam Bjorndahl (Carnegie Mellon University)

Title: A Logic of Uncertain Interpretation

Abstract: We introduce a logical framework for reasoning about "uncertain interpretations" and investigate two key applications: a new semantics for implication capturing a kind of "meaning entailment", and a conservative notion of "evidentially supported" belief that takes the form of a Dempster-Shafer belief function.