



JOIN ON TEAMS
LINK

ONLINE WORKSHOP
25-27 Feb 2026

LOGIC AND BOUNDED RATIONALITY

Organizers
Marcello D'Agostino
marcello.dagostino@unimi.it
Costanza Larese
costanza.larese@unimi.it

PROGRAMME
Central European Time

Wed 25 Feb

- 9:45-10:00 Welcome and opening remarks
- 10:00-10:45 Marcello D'Agostino (University of Milan)
- 10:45-11:30 Federica Conte (University of Milan-Bicocca)
- 11:30-12:15 Carlo Reverberi (University of Milan-Bicocca)

- 14:00-14:45 João Marcos (Federal University of Santa Catarina)
- 14:45-15:30 Carlos Caleiro (Technical University of Lisbon)
- 15:30-16:15 Sérgio Marcelino (Institute of Telecommunications Lisbon)
- 16:15-17:00 Alejandro Solares-Rojas (University of Buenos Aires)

Thu 26 Feb

- 10:00-10:45 Giuseppe Primiero (University of Milan)
- 10:45-11:30 Gerhard Lakemeyer (RWTH Aachen University)
- 11:30-12:15 Nina Gierasimczuk (Technical University of Denmark)

- 14:00-14:45 Francesco Berto (University of St Andrews)
- 14:45-15:30 João Mendes Lopes Neto (University of Brasília)
- 15:30-16:15 Giuseppe Sergioli (University of Cagliari)
- 16:15-17:00 Costanza Larese (University of Milan)

Fri 27 Feb

- 10:00-10:45 Mario Piazza (Scuola Normale Superiore Pisa)
- 10:45-11:30 Sanjay Modgil (King's College London)
- 11:30-12:15 Paolo Baldi (University of Salento)


- 14:00-14:45 Hykel Hosni (University of Milan)
- 14:45-15:30 Marcelo Finger (University of São Paulo)
- 15:30-16:15 Dov Gabbay (King's College London)



Finanziato
dall'Unione europea
NextGenerationEU



Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA



CORTEX
The Cost of Reasoning: Theory
and Experiments
PRIN 2022 project
2022ZLLR3T_001
Italian Ministry for University and
Research



UNIVERSITÀ
DEGLI STUDI
DI MILANO



ReDa
REasoning with DATA
FIS1 Advanced Grant Schema
G53C23000510001
Italian Ministry of University and
Research

**Marcello
D'Agostino**

(University of Milan)

TITLE: An Overview of Depth-Bounded Reasoning

ABSTRACT: Logical systems typically embody an ideal of unbounded inferential power: arbitrarily deep constructions and unrestricted proof search. Yet this ideal conflicts not only with undecidability phenomena, but also with the pervasive intractability of many decidable fragments—even in propositional logic. The problem of reasoning under depth constraints therefore reflects a structural feature of logical complexity, independently of any particular logic.

Depth-Bounded Reasoning (DBR) addresses this tension by restricting the depth at which inference unfolds, while leaving the underlying semantics unchanged. Logical consequence can thus be viewed as a hierarchy of bounded stages, representing progressively less tractable yet increasingly expressive approximations to full validity within a given logical framework.

Over the past three decades, a convergent—though often unsystematic—research trajectory has explored depth-sensitive fragments, bounded search procedures, and resource-limited inference across different logics. These developments share a common heuristic insight: logical consequence can be approached progressively, through controlled increases in inferential depth. It is time to articulate the basic principles and unifying perspective of this program.

A central open issue concerns the notion of depth itself. There may be no single canonical measure, but rather families of context-dependent—and potentially multidimensional—indices that correlate with computational hardness and cognitive difficulty.

Depth-Bounded Reasoning should therefore be understood as a general paradigm: a systematic way of reconciling logical idealization with computational and cognitive limitations, and of studying the graded structure of inferential power across logics.

Federica Conte

(University of Milan-
Bicocca)

TITLE: Insights from the CISIA TOLC-PSI test: testing models of reasoning and performance factors on a large empirical dataset

ABSTRACT: Current psychological models of human reasoning use a small set of assumptions to explain and predict the patterns of success and error in solving deductive reasoning problems. Here, we test those theories on a large dataset derived from the CISIA TOLC-PSI test, which collected answers to a wide range of deductive problems from over 36000 prospective university psychology students.

We compared predictions derived from the computational model of the psychological Theory of Mental Models (MMT) and the Depth-Bounded Boolean Logic (DBBL) against empirical observations from a set of 34 propositional syllogisms of varying complexity, as well as 31 categorical syllogisms and a smaller ($n = 11$) set of dynamic, meta-logical and iterative reasoning problems. We will present and discuss matches between theoretical predictions and empirical observations, and highlight instances in which theoretical predictions under-perform, identifying potential explanations.

Finally, we will present preliminary results examining whether performance on specific deductive reasoning tasks and on the overall TOLC-PSI predicts academic success indicators, including exams passed in the first year, GPA, and time to graduation.

**Carlo
Reverberi**

(University of Milan-
Bicocca)

TITLE: Putting pieces together. An fMRI study on how logical connectives combine propositions

ABSTRACT: Goal achievement is a fundamental human ability often driven by rule-based behavior. These rules typically consist of simple propositions combined by logical connectives (e.g., AND, OR, IF-THEN). Understanding these behaviors requires identifying how the brain employs these connectives to integrate information. This study investigated whether the neural representation of simple rules is compositional, meaning that individual elements—the logical connective, the arguments, and their order—are recoverable from the representation of the whole. We collected fMRI data while participants maintained and evaluated rules in working memory, such as "There is a triangle AND there is a hexagon." In line with previous research, univariate analyses identified an active left frontoparietal network during rule representation. We then employed multivariate pattern analyses, which revealed distinct neural patterns elicited by different rule elements during both the delay and target phases. Intriguingly, the semantic details of propositions and logical connectives were stored independently from other rule components. Taken together, our results provide novel evidence for the compositional nature of human reasoning, suggesting that rule comprehension and neural-level representation rely on the maintenance of discrete, recoverable building blocks.

João Marcos

(Federal University
of Santa Catarina)

TITLE: Depth-bounded cut-based tableaux for logics determined by finite-valued non-deterministic matrices (joint work with Carlos Caleiro)

ABSTRACT: With distinct motivations and objectives, Roman Suszko, Dana Scott, and Newton da Costa proposed in the 1970s that Tarskian logics, regardless of how many truth values they may initially appear to employ, can always be characterized using only two logical values. Over the past two decades, a systematic research program has sought to demonstrate how such bivalent presentations can be achieved for increasingly broad classes of logics while preserving desirable computational properties. In this talk, I will present methods for obtaining bivalent effective semantics and uniform classical-style analytic tableaux for logics characterized by finite-valued nondeterministic semantics. I will further discuss how these results can be integrated with the depth-bounded approach to rationality. This reports on ongoing joint work with Carlos Caleiro.

Carlos Caleiro

(Technical University
of Lisbon)

TITLE: Bounded approximations using analytical Set x Set Calculi (joint work with Sérgio Marcelino)

ABSTRACT: We show that effective bounded approximations can be systematically obtained for the logics presented by any given finite analytical Set x Set calculus. In fact, we can take advantage of analyticity and of the proof-tree notion underlying such a calculus, to define a hierarchy of PTIME approximations converging to the logic of interest by suitably controlling the nesting of genuine multiple-conclusion rules in proof-trees. To make this concrete, we overview the technology that allows one to obtain, automatedly, such analytical Set x Set calculi characterizing the consequence relation corresponding to any given (monadic) finite-valued partial non-deterministic logical matrix (PNmatrix), whose decision problem is always in the class coNP. We illustrate these methods with a few meaningful examples. This talk is based on joint work with Sérgio Marcelino, who will further address the fundamental topic of PTIME consequence relations and their semantics.

Sérgio Marcelino

(Institute of
Telecommunications
Lisbon)

TITLE: PNmatrix semantics for PTIME logics (joint work with Carlos Caleiro)

ABSTRACT: Analytic Set x Set calculi provide a principled way to define a hierarchy of PTIME approximations converging to a target logic. We focus on the base level of this hierarchy, obtained by dropping all branching rules, which guarantees PTIME tractability when we restrict the instantiations of rules to a context given by the analyticity property. In some cases, we can show that this first-level relation is indeed a logic, and that the corresponding Set x Fmla calculi (the usual Hilbert-style linear proofs) are analytical. Although Hilbert calculi are typically considered unsuitable for proof search because they lack the subformula property, in these specific cases they do satisfy it. Such systems are notoriously rare, but we will present examples and strategies to search for these rare logics, by exploring the technology that allows one to obtain analytic calculi characterizing the Set x Set consequence relation corresponding to any (monadic) finite-valued partial non-deterministic matrix (PNmatrix). Semantically, we show that this corresponds to a "fattening" of the departing PNmatrix, where the interpretation of the connectives is relaxed, where some truth-values are added to some entries corresponding to the effect of dropping the branching rules. This talk complements the previous one and is based on joint work with Carlos Caleiro.

Alejandro Solares-Rojas

(University of Buenos Aires)

TITLE: Depth-Bounded Intuitionistic Propositional Reasoning (joint work with M. D'Agostino and R. O. Rodriguez)

ABSTRACT: D'Agostino et al. proposed models for approximating classical propositional reasoning, which provide a measure of each inference depth. The latter corresponds to the number of nested uses of hypothetical case-analysis needed to draw it. Inferences obtainable within a certain depth are associated with a subsystem in a hierarchy increasingly approximating the full logic. Each of these subsystems is tractable whenever the inferences depth is bounded above by a fixed natural number and the set of formulas on which the hypotheses are allowed is suitably parameterized.

We extend the approach to intuitionistic propositional reasoning by replacing the idealized forcing relation of the Kripke semantics with a hierarchy of depth-bounded forcing relations, which measure the available inferential depth. The 0-depth subsystem is characterized by a sort of non-deterministic tables for the connectives specifying local constraints, and a global forcing requirement for the conditional that further constrains the non-deterministic entries of its table. Forcing relations of greater depth are characterized by recursively allowing a form of hypothetical case reasoning over a bounded space of formulas. The proof-theoretic characterization of the 0-depth logic is a system consisting of introduction and elimination rules involving only the use of information actually possessed by the agent. Logics of greater depth are characterized by the iterated applications of a unique branching rule governing the use of hypothetical information and implementing a generalization of the principle of bivalence.

Giuseppe Primiero

(University of Milan)

TITLE: A Proof System with Causal Labels (joint work with Leonardo Ceragioli)

ABSTRACT: Ensuring fairness in machine learning systems requires identifying and formally reasoning about essential resources such as probabilistic dependencies, causal relations, and structural constraints. Acting rationally on the behaviour of such systems is inherently bounded on the ability to express such resources. In this work, we present a proof system for the formal verification of fairness properties in probabilistic classifiers. The typed natural deduction calculus TNDPQ integrates causal labels directly into proofs, enabling principled reasoning about how protected attributes influence probabilistic outcomes. We characterise individual fairness and intersectionality through structural constraints on inference rules and the analysis of conditional independence, and we extend the calculus to model counterfactual reasoning by representing interventions and hypothetical scenarios over causal structures. This allows us to verify whether classifier decisions remain invariant under changes to protected attributes while holding relevant contextual variables fixed. The resulting system provides a sound and expressive post-hoc method for analysing fairness in black-box classifiers, bridging logical proof systems, causal reasoning, and contemporary fairness criteria in machine learning.

**Gerhard
Lakemeyer**

(RWTH Aachen
University)

TITLE: A Logic of Limited Belief Based on Possible Worlds

ABSTRACT: A fundamental problem in knowledge representation is that reasoning about knowledge bases using very expressive languages like first-order logic is generally undecidable. There are essentially two ways to address this problem, either by limiting the language or by devising an appropriate model of limited belief. Most existing work has focused on limiting the language, with description logics being perhaps the most prominent example. In this talk I will present a model of limited belief based on three-valued possible worlds. It is expressive in that knowledge bases can contain arbitrary sentences in first-order logic. Belief is defined in terms levels. At level 0 only the sentences in the knowledge base and simple consequences are believed. Higher levels of belief require more and more reasoning effort to uncover more and more logical consequences. Belief is also eventually complete in the sense that, if a sentence is entailed by a knowledge base, then there is a level where it is believed. Moreover, reasoning at any particular level can be shown to be tractable under reasonable assumptions. Throughout most of the talk I will assume that beliefs are objective, that is, do not mention other beliefs. Time permitting I will also report on recent work where we address the more general case of an agent with limited belief able to introspect on what it believes and does not believe.

**Nina
Gierasimczuk**

(Technical University of
Denmark)

TITLE: Learning about causal variables: an awareness account (joint work with Hermine J. Grosinger and Katrine B. P. Thoft)

ABSTRACT: In this talk I will present a way to enrich the existing dynamic epistemic modal logic of causality with dynamic awareness of causal variables, where agent can become aware of (previously unavailable) causal variables. This approach allows studying important aspects of the development of agents' knowledge about causal structures. We will consider scenarios that feature the expert (teacher) who has full knowledge of the causal dependencies, and the layperson (learner) whose knowledge is limited to only the variables she is aware of. Our models and our logic are dynamic: They allow observations by both the expert and the layperson, and they accommodate gaining awareness of causal variables by the layperson. The logic also allows interventions—counterfactual manipulations of values of variables. In the epistemic causal awareness models, the different epistemic levels of the expert and the layperson correspond to implicit and explicit knowledge, respectively. Those different perspectives can also be seen with respect to the causal graphs, in a 2-layer representation (so-called hybrid causal graph) that makes use of 'shadow' causal variables. Overall, our framework connects the existing lines of research on dynamic epistemic logic of causality and dynamic awareness with learning and pro-activity in AI.

**Francesco
Berto**

(University of St
Andrews)

TITLE: Logical Omniscience and Framing Effects (joint work with Aybüke Özgün)

ABSTRACT: Humans sometimes believe only one of two logically or necessarily equivalent propositions P and Q . That can depend on how such contents are presented to them. But if one has a hyperintensional account of propositions, i.e., one individuating them more finely than as sets of possible worlds, there's an extra possibility: sometimes humans believe only one of such P and Q because they differ in content, in spite of being classically-modally equivalent. We argue that this sort of framing effect is structural and pervasive: it pivots on the distinction between working memory and long term memory. We introduce a class of models featuring topics or subject matters to reason about agents that can be so framed, and a logic that's sound and complete with respect to the class. While doing so, we clarify the sense in which such agent are not logically omniscient – for, we believe, 'logical omniscience' is spoken of in many ways.

**João Mendes
Lopes Neto**

(University of Brasília)

TITLE: A method for automated generation of proof exercises with comparable levels of complexity (joint work with João Marcos and Patrick Terrematte)

ABSTRACT: The automated generation of exercises may benefit educators by significantly reducing the time they spend in manually creating exercises. However, an obstacle in making such automation more present in an educator's professional routine is controlling the level of complexity of mechanically generated exercises. In this work, we present a method for the automated generation of proof exercises with a comparable level of complexity. The inputs of this method are a proof exercise and a set of rules allowing to prove this exercise. The output is a set of proof exercises with a comparable complexity to that of given as input. The scope of exercises we work with are mathematical proof exercises described in first-order languages, covering topics such as those usually taught in Discrete Math undergraduate modules. In order to calculate the level of complexity of these exercises, we base our approach on the effort required to solve them via informal proofs. We argue that such an effort, in turn, may be captured by formal proofs in cut-based tableaux that do not contain logical symbols. The rules used in these proofs are extracted by a mechanizable procedure we provide. We use the analytic character of such rules and the formal structure tableau proofs have to provide a computational procedure of the method in question.

**Giuseppe
Sergioli**

(University of Cagliari)

TITLE: Quantum and Quantum Inspired Machine Learning - from theory to implementations

ABSTRACT: The aim of the talk is to introduce a new quantum-like method for the binary classification applied to classical datasets. Inspired by the quantum Helstrom measurement, this approach enables to define a new binary classifier, called Helstrom Quantum Classifier (HQC). This classifier (inspired by the concept of distinguishability between quantum states) acts on density matrices — called density patterns — that are the quantum encoding of classical patterns of a dataset. We compare the performance of HQC with respect to several standard classifiers over different datasets and we show that HQC outperforms the other classifiers when compared to the Balanced Accuracy and other significant statistical measures. We also show that the performance of our classifier is positively correlated to the increase in the number of “quantum copies” of a pattern and the resulting tensor product thereof. In the last part of the talk we show a large-scale experiment based on the application of HQC to the biomedical imaging context in clonogenic assay evaluation to identify the most discriminative feature, allowing us to enhance cell colony segmentation.

In the final part of the talk, we will also present a method for generalizing the model to the multi-class classification scenario—avoiding an explosion of one-vs-one or one-vs-many binary classifiers—and we will illustrate how this model naturally lends itself to implementation on an actual quantum device.

**Costanza
Larese**

(University of Milan)

TITLE: The Depth-Bounded Approach in the Philosophy of Logic: Normativity and Pluralism

ABSTRACT: Logic is traditionally taken to be normative for rational thought: logical consequence is supposed to determine what agents ought to believe or infer. Yet classical accounts face a tension. Combined with logical monism, they impose excessive demands on finite agents; combined with logical pluralism, those demands remain and incompatible logics threaten the determinacy of rational requirements.

The depth-bounded approach to logic sheds new light on these traditional debates in the philosophy of logic. On this view, a logic is the limit of a hierarchy of tractable, resource-sensitive systems ordered by inferential depth, each modelling reasoning feasible for bounded agents. This framework yields, first, a revised account of logical normativity: depth-bounded normativity. Logical obligations apply only to inferences within an agent’s feasible reach. Classical normativity is reformulated through resource-sensitive bridge principles linking consequence, inferential difficulty, and cognitive capacity. Second, it supports resource-based logical pluralism. Here plurality arises not from rival logics, as in traditional forms of pluralism, but from differences in agents’ inferential resources within a shared framework. Agents at different depths may issue incompatible yet systematically related verdicts about consequence.

The depth-bounded approach thus shows how logic can remain genuinely normative without presupposing logical omniscience, while accommodating a novel form of pluralism.

Mario Piazza

(Scuola Normale
Superiore Pisa)

TITLE: A Proof-Theoretic Framework for Default Deontic Reasoning

ABSTRACT: We introduce a non-modal, proof-theoretic framework for default deontic reasoning for agents operating under incomplete information and potentially conflicting norms. The framework is based on controlled sequents, which regulate the application of defaults and norms via annotated control sets and enable local soundness checks. We prove strong completeness with respect to credulous consequence for default theories and normative systems. Our platform also supports flexible strategies for resolving deontic conflicts and paradoxes.

Sanjay Modgil

(King's College
London)

TITLE: Bounded Non-Monotonic Reasoning: The Dialogical Turn

ABSTRACT: In this talk, I argue that recent developments in AI motivate a systematic dialogical formalisation of non-monotonic logics. I begin by revisiting joint work with Marcello D'Agostino, where we develop dialectical models of single-agent non-monotonic reasoning that yield rational outcomes under resource bounds. I then outline how this framework can be generalised to genuinely dialogical settings, yielding formal accounts of distributed non-monotonic reasoning. On this view, rational outcomes arise from the regulated exchange of speech acts governed by protocols that encode norms of commitment and entitlement within the practice of giving and asking for reasons (R.B. Brandom). The result is a conception of non-monotonic consequence grounded in norm-governed interaction, suitable for modelling resource-bounded, multi-agent epistemic environments.

Paolo Baldi

(University of Salento)

TITLE: Depth-bounded qualitative probabilities**ABSTRACT:** Qualitative probabilities are relational structures which are representable by probability measures. Such structures play a pivotal role in the foundation of subjective probability and decision theory, since they are considered to be a more realistic model of an agent's belief.

As for quantitative probability, however, classical reasoning with qualitative probability is still computationally intractable.

In this talk, we will introduce tractable qualitative structures approximating qualitative probability, on the basis of the tractable approximation of classical logic provided by the Depth-Bounded Boolean logics, introduced by D'Agostino, Finger and Gabbay (2013).

Hykel Hosni

(University of Milan)

TITLE: Depth-bounded Belief Functions in retrospect**ABSTRACT:** I welcome the opportunity to celebrate the first milestone of the Depth-Bounded Logic programme to revisit the motivation and results of [Baldi and Hosni 2020](#) in light of data-driven reasoning. There we showed that Dempster–Shafer belief functions are recovered as the limit of depth-bounded belief functions as logical depth grows. Moreover, we identified the conditions under which (additive) probability functions emerge as the limit of a sequence of (non-additive) depth-bounded belief functions. The purpose of my presentation is to elaborate on the role of depth-bounded logic in posing a sharp conceptual distinction between evidence, i.e. the uncertainty-resolving information an agent has, and their degrees of belief, i.e. what the agent uses in reasoning and decision-making. In particular I will speculate on the applicability of this distinction to the problem of analysing expert disagreement in scientific reasoning.

Marcelo Finger

(University of São Paulo)

TITLE: Satisfiability for Lukasiewicz Logic and Its Approximations with Applications on Proving Properties about Neural Networks (joint work with Sandro Preto)

ABSTRACT: We start by presenting the satisfiability problem for Łukasiewicz infinitely-valued logic (L_{∞} -SAT), with possible applications to representing piecewise linear functions and proving properties about neural networks. We discuss how L_{∞} -SAT can be obtained via Mixed Integer Programming, which, in fact, highlights the NP-completeness of the problem. We then use this gives rise to a polyhedral semantics for a parameterised family of tractable logics whis approximate L_{∞} such that each logic in the family is associated to a polynomial-time linear program.y. In this way, we provide a tractable decision problem for each intermediate logic in the path to obtaining L_{∞} -SAT. We highlight some properties of the logic system derived from polyhedral semantics and the details of an algorithm for the approximation process.

Dov Gabbay

(King's College London)

TITLE: Active Inference, Formal Argumentation and Temporal Logic

ABSTRACT: This paper argues for a systematic dialogue between two largely independent research communities: active inference, which models sentient behaviour in neuroscience, robotics, and AI; and logic and formal argumentation, which model structured human reasoning. Although their application domains overlap, they rely on different formal tools and conceptual frameworks.

Through simple but structurally rich token examples of goal-directed agents acting in dynamically changing environments, we show that each community faces complementary limitations. Active inference offers powerful probabilistic and generative models, yet typically focuses on statistical or generic cases. Logic and argumentation provide precise structural representations of reasoning, but require new logical mechanisms to adequately capture temporality, domain variation, resource sensitivity, and the integration of metalevel effects into object-level reasoning.

We suggest that developments such as labelled deductive systems, reactive semantics, and fibred logics provide promising tools for this task. More generally, we propose a shared research program in which active inference motivates new logical formalisms, and logical analysis helps clarify and extend the formal foundations of active inference.
