

Abstracts

- Diego Agustin Ambrossio: A Non-Monotonic Logic for Distributed Access Control (with Marcos Cramer and Pieter Van Hertum)

We define Distributed Access Control Logic (D-ACL), whose main difference from state-of-the-art says-based access control logics is that it is non-monotonic, thus allowing access denials to be modelled straightforwardly in the formalism. Additionally D-ACL allows for access rights and other properties relevant to access control to be defined inductively, which increases its expressive power compared to other first-order access control logics. The semantics of D-ACL is based on the well-founded semantics for autoepistemic logic. By considering some application scenarios, we compare D-ACL with other saysbased access control logics as well as with other semantics of autoepistemic logic. Furthermore, we define a query-driven decision procedure for the propositional fragment of D-ACL, which allows to determine access rights while minimizing the information flow between principals in order to reduce privacy concerns.

- Guillaume Aucher: Multi-Agent Only Knowing on Planet Kripke

The idea of only knowing is a natural and intuitive notion to precisely capture the beliefs of a knowledge base. However, an extension to the many agent case, as would be needed in many applications, has been shown to be far from straightforward. For example, previous Kripke frame-based accounts appeal to proof-theoretic constructions like canonical models, while more recent works in the area abandoned Kripke semantics entirely. We propose a new account based on Moss characteristic formulas, formulated for the usual Kripke semantics. This is shown to come with other benefits: the logic admits a group version of only knowing, and an operator for assessing the epistemic entrenchment of what an agent or a group only knows is definable. Finally, the multi-agent only knowing operator is shown to be expressible with the cover modality of classical modal logic, which then allows us to obtain a completeness result for a fragment of the logic.

- Richard Booth: Trust-sensitive Belief Revision (with Aaron Hunter)

Belief revision is concerned with incorporating new information into a pre-existing set of beliefs. When the new information comes from

another agent, we must first determine if that agent should be trusted. In this paper, we define trust as a pre-processing step before revision. We emphasize that trust in an agent is often restricted to a particular domain of expertise. We demonstrate that this form of trust can be captured by associating a state partition with each agent, then relativizing all reports to this partition before revising. We position the resulting family of trust-sensitive revision operators within the class of selective revision operators of Ferme and Hansson, and we examine its properties. In particular, we show how trust-sensitive revision is manipulable, in the sense that agents can sometimes have incentive to pass on misleading information. When multiple reporting agents are involved, we use a distance function over states to represent differing degrees of trust; this ensures that the most trusted reports will be believed.

- Jan Broersen: TBA
- Patrice Caire: Collaborative Explanation and Response in Assisted Living Environments Enhanced with Humanoid Robots

An ageing population with increased social care needs has provided recent impetus for research into assisted living technologies, as the need for different approaches to providing supportive environments for senior citizens becomes paramount. Ambient intelligence (AmI) systems are already contributing to this endeavour. A key feature of future AmI systems will be the ability to identify causes and explanations for changes to the environment, in order to react appropriately. We identify some of the challenges that arise in this respect, and argue that an iterative and distributed approach to explanation generation is required, interleaved with directed data gathering. We further argue that this can be realised by developing and combining state-of-the-art techniques in automated distributed reasoning, activity recognition, robotics, and knowledge-based control.

- Martin Caminada: Some reflections on obtaining a lectureship in the UK

In this talk I will share some personal experiences of academic life in the United Kingdom. A brief overview is presented on how research and teaching in the UK differs from the Continent. Also, an analysis

is given on why the UK currently offers quite some job opportunities to overseas researchers.

- Giovanni Casini: On the Entailment Problem for a Logic of Typicality (with Richard Booth, Thomas Meyer and Ivan Varzinczak)

Propositional Typicality Logic (PTL) is a recently proposed logic, obtained by enriching classical propositional logic with a typicality operator. In spite of the non-monotonic features introduced by the semantics adopted for the typicality operator, the obvious Tarskian definition of entailment for PTL remains monotonic and is therefore not appropriate. We investigate different (semantic) versions of entailment for PTL, based on the notion of Rational Closure as defined by Lehmann and Magidor for KLM-style conditionals, and constructed using minimality. Our first relevant result is an impossibility theorem showing that a set of proposed postulates that at first all seem appropriate for a notion of entailment with regard to typicality cannot be satisfied simultaneously. Closer inspection reveals that this result is best interpreted as an argument for advocating the development of more than one type of PTL entailment. In the spirit of this interpretation, we define two primary forms of entailment for PTL and discuss their advantages and disadvantages.

- Marcos Cramer: Cognitive Aspects of Formal Argumentation

In this talk, I will give a brief overview over the research foreseen for my new research project, Cognitive Aspects of Formal Argumentation. In this collaboration with cognitive scientists, we will empirically evaluate assumptions made by formal argumentation theory (e.g. reinstatement and more specific assumptions underlying certain argumentation semantics) by confronting human subjects with carefully chosen arguments and analysing their reactions. In order to minimize the influence of world knowledge on the subjects' evaluation of the arguments, we use arguments from the abstract domains of metalinguistic and mathematical reasoning in this empirical study. In particular, we will confront the subjects with different arguments about possible resolutions of logical paradoxes like the Liar paradox. As a theoretical preparatory work for the empirical study, we will build a formal model of the arguments that have been put forward for and against different resolutions of such

paradoxes in the philosophical literature. For this we plan to combine eelja and Straer's Explanatory Argumentation Frameworks with higher-order argumentation and the meta-argumentation methodology.

- Mehdi Dastani: Monitoring Interaction in Organizations Using Norm Enforcement

In an organisational setting such as an online marketplace, an entity called the organisation or institution defines interaction protocols, monitors agent interaction, and intervenes to enforce the interaction protocols. In this presentation, I explain how interaction protocols can be modeled by explicitly-represented norms, how the enforcement of protocols can be operationalised by means of norm enforcement, and how to analyse the protocol by a logical analysis of the norms.

- Juergen Dix: An agent-based simulation platform and its use for software projects

We describe Maserati, an agent based simulation platform that was developed to scale up, allowing to deploy and simulate hundreds of thousands of agents. We also motivate a particular application of our platform, namely the quality control of software projects during their whole life cycle.

- Dragan Doder: Decreasing and Resolving Inconsistency by Minimal Knowledge Base Changes (with Srdjan Vesic)

This paper studies different techniques for measuring and decreasing inconsistency of a knowledge base. We define an operation that allows to decrease inconsistency of a knowledge base while losing a minimal amount of information. We also propose two different ways to compare knowledge bases. The first is a partial order that we define on the set of knowledge bases. We study this relation and identify its link with a particular class of inconsistency measures. We also study the links between the partial order we introduce and information measures. The second way we propose to compare knowledge bases is to define a class of metrics that give us a distance between knowledge bases. They are based on symmetric set difference of models of pairs of formulae from the two sets in question. We then use those metrics to define a new class of inconsistency measures and prove the links between those metrics.

- Dov Gabbay: Talmudic Norms Approach to the Paradox of the Heap: A Position Paper (with E. David, S. David and U. Schild)

This paper offers a Talmudic norms solution to the paradox of the heap. The claim is that the paradox arises because philosophers use the wrong language to discuss it and the appropriate language is that of an extended blocks world language, together with the Talmudic normative theory of mixing (Talmudic calculus of Sorites) and the principle that a property of any mixture (or indeed any object) is also how it was constructed. We seek a correlation between Talmudic positions on mixtures and philosophical positions on Sorites. The Talmud is very practical and cannot allow for any theoretically unresolved paradox to get in the way, and so it has a lot to offer to philosophy.

- Aldo Gangemi: Tackling higher-order and contextual phenomena in Open Knowledge Extraction

Open Knowledge Extraction aims at the extraction of formal knowledge from text in an unsupervised, open domain way. The current methods and tools have a decent accuracy with respect to the basic structure of sentences, but a mostly unexplored world appears when multiple sentences/texts, higher-order constructions, explicit or implicit contexts are taken into account. Cognitive, social, and normative aspects of text interpretation are unfortunately mostly made of this dark matter. We will present some of the problems, a few existing solutions, and the difficult relation between logical and stochastic approaches.

- Joris Hulstijn: Towards Trusted Trade Lanes (with Wout Hofman, Gerwin Zomer and Yao-Hua Tan)

Customs administrations are exploring system-based approaches to regulatory supervision, which take the entire set of controls in a process into account. In addition to Trusted Traders, which are recognized by a certification process, customs are considering to identify so called Trusted Trade Lanes: companies that collaborate in a trade lane in a reliable manner. If trade lanes are demonstrated to be trustworthy, customs can direct their efforts elsewhere and reduce inspections. In this discussion paper we explore the concept of a trusted trade lane. We define essential characteristics and develop various scenarios in which trade lanes may demonstrate to the authorities that they conform to

these requirements. We review three information architectures to improve supply chain visibility and drive supply chain innovation, in line with current regulatory developments.

- Wojtek Jamroga: On Module Checking and Strategies (with Aniello Murano)

Two verification problems are very close in spirit: module checking of the temporal logic CTL and model checking of the game logic ATL. The latter appears to be a natural multi-agent extension of the former, and it was commonly believed that model checking of ATL subsumes module checking of CTL in a straightforward way.

We show that, on the contrary, the two problems are fundamentally different. The way in which behavior of the environment is understood in module checking cannot be equivalently characterized in ATL. Moreover, if one wants to embed module checking in ATL then its semantics must be extended with two essential features, namely nondeterministic strategies and long-term commitment to strategies.

- Manuela Jungmann: Recursive Interaction Design

In this short talk I will give a brief introduction to the nature of interaction design, a branch of the interdisciplinary field of Human-computer Interaction. I will then introduce one of my designs. Upon completion the design was publicly exhibited. During the exhibitions I collected data from the design to recursively study user behaviour. Recursively in the sense that it was not an evaluation of the design, but rather the aim was to gain insights into the minute dimensions of participant interaction of a successful design. My findings revealed unintended design aspects, which I was able to learn from for future designs.

- Xavier Parent: On Cumulative Transitivity from Deontic Perspective
- Gabriella Pigozzi: Voting and deliberation
- Livio Robaldo: Natural Language Processing and Natural Language Semantics for representing/mining knowledge from Legal Text
- Francois Schwarzentruher: Overview about Epistemic planning

In this talk, we will sum up different results in epistemic planning. We first recall the framework of dynamic epistemic logic and the epistemic

planning problem. We then give some complexity and decidability results for different variants of the epistemic planning problem.

- Marija Slavkovik: Beyond murder and mayhem: machine ethic challenges in decision making

Should the autonomous car be programmed to kill its passenger for the purpose of saving ten people standing by the road? Machine ethics is a new field in AI concerned with enabling autonomous machines with ethical behaviour. Beyond the embodied autonomous systems with capabilities to physically hurt life are the comparatively much less visible but much more prevalent autonomous web based services that are increasingly more intelligent and increasingly more in control of the decisions we think we freely make. Opinions expressed in social media regularly converge into seemingly public collective opinions that have real influence on policy-making as well as individual persons lives. This talk reviews the machine ethics challenges faced by various forms of intentional and unintentional collective decision making systems embedded in web based media.

- Xin Sun: Logic and Games for Normative Agents

In this talk we study how to characterize normative agents in games. We adopt a proposition control game + input/output logic approach. Norms create the normative status of strategies. Agents preference in proposition control games are changed by the normative status of strategies. We distinguish four ethical types of agents: moral, amoral, negatively impartial and positively impartial. Agents of different normative types use different input/output logic for normative reasoning and different procedures to change their preference. Such preference change induces normative proposition control games and notions like normative Nash equilibrium are then introduced. We study some complexity issues related to normative reasoning/status and normative Nash equilibrium.

- Masoud Tabatabaei: Information Security as Strategic (In)effectivity

Security of information flow is commonly understood as preventing any information leakage, regardless of how grave or harmless consequences the leakage can have. Even in models where each piece of information is classified as either sensitive or insensitive, the classification is hardwired

and given as a parameter of the analysis, rather than derived from more fundamental features of the system. In this work, we suggest that information security is not a goal in itself, but rather a means of preventing potential attackers from compromising the correct behavior of the system. To formalize this, we first show how two information flows can be compared by looking at the adversary's ability to harm the system. Then, we propose that the information flow in a system is effectively information-secure if it does not allow for more harm than its idealized variant based on the classical notion of noninterference.

- Leon van der Torre: TBA
- Srdjan Vesic: Ranking Arguments With Compensation-Based Semantics (with Leila Amgoud, Jonathan Ben-Naim and Dragan Doder)

In almost all existing semantics in argumentation, a strong attack has a lethal effect on its target that a set of several weak attacks may not have. This paper investigates the case where several weak attacks may compensate one strong attack. It defines a broad class of ranking semantics, called alpha-OBBS, which satisfy compensation. alpha-OBBS assign a burden number to each argument and order the arguments with respect to those numbers. We study formal properties of alpha-OBBS, implement an algorithm that calculates the ranking, and perform experiments that show that the approach computes the ranking very quickly. Moreover, an approximation of the ranking can be provided at any time.

- Serena Villata: TBA
- Emil Weydert: Nonmonotonic reasoning for set theory

The multiverse paradigm in set theory does not only reflect philosophical preferences, or set up a new playground for mathematical investigation, but it also offers a powerful methodological tool for investigating the conceptual foundations of set theory by guiding the search for and the evaluation of new set-theoretic axioms or facts. The prototypical example is Friedman's Hyperuniverse Program (HUP). We sketch a general abstract inferential framework whose (defeasible, inductive) inference methods are meant to identify or validate new axioms. We consider nonmonotonic consequence relations — , parametrized by specifications of the multiverse and set-theoretic desiderata, which associate

with any suitable $ZFC + X$ new - not necessarily classically derivable - candidate truths. The desiderata could, for instance, consist of consistency conditions or maximization demands w.r.t. preorders over universes. There are a number of possible inductive strategies, but conceptual considerations at the level of set theory may fail to decide among them. The idea is hence to also assess the inferential level and to use rationality postulates for nonmonotonic inference, heavily investigated within AI for modeling commonsense reasoning, to classify and evaluate such procedures. This is a highly non-trivial task because of the special characteristics of axiom induction.

- Zhe Yu: Analysis of natural language argumentation in trial cases where does uncertainty come from?
- Marc van Zee: Intention Reconsideration as Metareasoning

We first present a theoretical framework for the intention reconsideration problem in MDPs, in the same spirit as much other work on metareasoning. This involves the construction of a meta-level MDP in which the two actions are think or act. We then consider Kinny and Georgeffs framework as a special case, reproducing their results, and comparing their agents to an angelic agent who decides optimally when to think or act. Interestingly, even the very simple agents Kinny and Georgeff considered behave nearly optimally in certain environments. However, no agent performs optimally across environments. Our results suggest that meta-meta-reasoning may indeed be called for in this setting, so that an agent might tune its reconsideration strategy flexibly to different environments.