SEMINAR · KNOWLEDGE REPRESENTATION AND REASONING

Talk Outlines and Bibliography

Twenty self-contained introductory talks of 45 minutes each

About this document

There are 20 self-contained talks of 45 minutes each, ordered by the year of each topic's founding contribution. The only prerequisites are propositional and first-order logic, theoretical computer science and undergraduate mathematics.

- **Outline:** a suggested sequence of blocks with timings that add up to 45 minutes. "Core" marks the formal result the talk must get across.
- **Bibliography:** every entry was checked online. Each link goes to an open-access copy where one exists, otherwise to the publisher's page or DOI. Print-only works (e.g. Hintikka 1962) were left out. ACM papers count as open because the ACM Digital Library has been fully open access since January 2026.
- **Access legend:** OPEN = freely accessible; LIBRARY = publisher copy that needs a university subscription (e.g. via SLUB Dresden).
- **Start here** marks the one or two sources a student should read first.

Contents

TALK 01

Deontic Logic

CORE

Standard deontic logic (SDL) as the modal logic KD, the paradoxes that break it, and input/output logic as a response.

Outline 45 min

1. **Motivation** 5 min Norms in law, contracts, compliance checking and machine ethics. Obligations are about what *should* hold, not what does hold.
2. **Standard deontic logic** 10 min Operators O, P, F with P = ¬O¬ and F = O¬. Axioms K and D (Op ⟶ Pp) plus necessitation. Kripke semantics with a serial ideality relation; state the completeness theorem.
3. **Paradoxes** 10 min Ross's paradox: Op ⟶ O(p ∨ q). Chisholm's contrary-to-duty scenario (go and help; tell if you go; don't tell if you don't go; you don't go). Show that every natural SDL rendering is either inconsistent or makes a premise redundant.
4. **Input/output logic** 15 min Norms as pairs (a, x) rather than formulas; out(G, A) = Cn(G(Cn(A))). Derivation rules SI, WO, AND, OR, CT. The four systems out_1–out_4 (simple-minded, basic, simple-minded reusable, basic reusable) with their soundness and completeness. Constrained output as the answer to contrary-to-duty cases.
5. **Outlook** 5 min Dyadic deontic logic, normative multi-agent systems, legal informatics.

Bibliography

OPEN START HERE Paul McNamara and Frederik Van De Putte. Deontic Logic. *Stanford Encyclopedia of Philosophy*, first published 2006, substantive revision 2021.

OPEN START HERE David Makinson and Leendert van der Torre. What is Input/Output Logic? In B. Löwe, W. Malzkorn and T. Räsch (eds.), *Foundations of the Formal Sciences II*, Trends in Logic 17, Kluwer, 163–174, 2003.

LIBRARY Georg Henrik von Wright. Deontic Logic. *Mind* 60(237):1–15, 1951. DOI 10.1093/mind/LX.237.1.

LIBRARY Roderick M. Chisholm. Contrary-to-Duty Imperatives and Deontic Logic. *Analysis* 24(2):33–36, 1963. DOI 10.1093/analys/24.2.33.

LIBRARY David Makinson and Leendert van der Torre. Input/Output Logics. *Journal of Philosophical Logic* 29(4):383–408, 2000. DOI 10.1023/A:1004748624537. Open preprint.

TALK 02

Epistemic Logic and Public Announcements

CORE

Multi-agent S5, common knowledge as reachability, and public announcement logic (PAL) with its reduction axioms.

Outline 45 min

1. **Motivation** 5 min State the muddy children puzzle and leave it open. Reasoning about what others know matters in protocols, games and multi-agent systems.
2. **Multi-agent epistemic logic** 12 min Operators K_a. Kripke models with one equivalence relation per agent. Axioms K, T, 4, 5 and the frame property each one corresponds to. Worked example: a three-card game.
3. **Group knowledge** 8 min Everybody knows (E_G) and common knowledge (C_G). C_G as reachability along the union of the agents' relations. Fixpoint property $C\varphi \leftrightarrow E(\varphi \wedge C\varphi)$; mention the coordinated attack impossibility.
4. **Public announcements** 15 min Semantics of $[!\psi]\varphi$ as restriction of the model to the ψ-worlds. Reduction axioms, the key one being $[!\psi]K_a\varphi \leftrightarrow (\psi \rightarrow K_a[!\psi]\varphi)$. Consequence: without common knowledge, PAL is exactly as expressive as epistemic logic, but the translation can blow up exponentially. Unsuccessful updates (Moore sentences such as "p, but you don't know p"). Solve muddy children formally.
5. **Outlook** 5 min With common knowledge PAL is strictly more expressive; action models (Baltag–Moss–Solecki); epistemic planning.

Bibliography

OPEN START HERE Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek and Barteld Kooi. An Introduction to Logics of Knowledge and Belief. Chapter 1 in *Handbook of Epistemic Logic*, College Publications, 1–51, 2015. arXiv:1503.00806.

OPEN START HERE Alexandru Baltag and Bryan Renne. Dynamic Epistemic Logic. *Stanford Encyclopedia of Philosophy*, first published 2016.

OPEN Rasmus Rendsvig, John Symons and Yanjing Wang. Epistemic Logic. *Stanford Encyclopedia of Philosophy*, first published 2019, substantive revision 2023.

LIBRARY Jan Plaza. Logics of Public Communications. *Synthese* 158(2):165–179, 2007. DOI 10.1007/s11229-007-9168-7. Reprint of the 1989 paper in the ISMIS poster session proceedings (ORNL/DSRD-24), 201–216.

LIBRARY Ronald Fagin, Joseph Y. Halpern, Yoram Moses and Moshe Y. Vardi. Reasoning About Knowledge. MIT Press, 1995. DOI 10.7551/mitpress/5803.001.0001.

LIBRARY Hans van Ditmarsch, Wiebe van der Hoek and Barteld Kooi. Dynamic Epistemic Logic. Synthese Library 337, Springer, 2007. DOI 10.1007/978-1-4020-5839-4.

OPEN Carsten Lutz. Complexity and Succinctness of Public Announcement Logic. *Proc. AAMAS 2006*, ACM, 137–144, 2006. DOI 10.1145/1160633.1160657.

TALK 03

Situation Calculus and the Frame Problem

CORE
Reiter's successor state axioms as a solution to the frame problem, basic action theories, and the regression theorem.

Outline 45 min

1. **Motivation** 5 min An agent must predict what its actions change and what they leave alone. Blocks-world example; McCarthy and Hayes's statement of the frame problem.
2. **The language** 10 min Sorted first-order logic with actions, situations, S0 and do(a, s); relational fluents; the Poss predicate; precondition axioms.
3. **The frame problem** 8 min Effect axioms alone do not determine what stays the same. Explicit frame axioms are needed for roughly every action–fluent pair. Mention the ramification and qualification problems.
4. **Reiter's solution** 12 min Positive and negative effect axioms in normal form. The explanation closure (causal completeness) assumption. The resulting successor state axiom $F(x, do(a, s)) \leftrightarrow \gamma^{+}(x, a, s) \vee (F(x, s) \wedge \neg\gamma^{-}(x, a, s))$. Unique names axioms for actions. A basic action theory consists of foundational axioms, successor state, precondition and unique names axioms, and the initial database.
5. **Regression** 8 min The regression operator. Theorem: a regressable query φ is entailed by the action theory iff its regression is entailed by the initial database plus the unique names axioms. Worked projection example; planning as deduction.
6. **Outlook** 2 min Golog; knowledge and sensing in the situation calculus.

Bibliography

LIBRARY START HERE Fangzhen Lin. Situation Calculus. Chapter 16 in F. van Harmelen, V. Lifschitz and B. Porter (eds.), *Handbook of Knowledge Representation*, Elsevier, 649–669, 2008. DOI 10.1016/S1574-6526(07)03016-7.

OPEN START HERE · FOR THE PROBLEM Murray Shanahan. The Frame Problem. *Stanford Encyclopedia of Philosophy*, first published 2004, substantive revision 2016. Conceptual only: no successor state axioms or regression.

OPEN John McCarthy and Patrick J. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer and D. Michie (eds.), *Machine Intelligence 4*, Edinburgh University Press, 463–502, 1969.

LIBRARY Raymond Reiter. The Frame Problem in the Situation Calculus: A Simple Solution (Sometimes) and a Completeness Result for Goal Regression. In V. Lifschitz (ed.), *Artificial and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, Academic Press, 359–380, 1991. DOI 10.1016/B978-0-12-450010-5.50026-8.

LIBRARY Raymond Reiter. Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems. MIT Press, 2001. DOI 10.7551/mitpress/4074.001.0001. The free companion site has the introduction, errata and Golog code.

TALK 04

Circumscription

CORE

The second-order formula Circ[T; P; Z], its characterisation by minimal models, when it collapses to first-order logic, and its Π_2^p-completeness.

Outline 45 min

1. **Motivation** 5 min The closed-world assumption in databases. “Birds normally fly” via an abnormality predicate. The qualification problem.
2. **Minimal models** 8 min The preorder $\leq_{P;Z}$ on models with minimised predicates P, varying Z and the rest fixed. A propositional example. Why the naive CWA is inconsistent with disjunctions.
3. **Second-order definition** 10 min $\mathrm{Circ}[T; P; Z] = T(P, Z) \land \neg\exists p, z\ (T(p, z) \land p < P)$. Theorem: its models are exactly the (P; Z)-minimal models of T. Demonstrate nonmonotonicity.
4. **Computing circumscription** 10 min Lifschitz’s reduction to first-order logic for solitary and, more generally, separable formulas; predicate completion as a special case. A case that is not first-order: minimising a transitive relation that contains E yields the transitive closure of E.
5. **Complexity and use** 10 min Propositional circumscriptive inference is Π_2^p-complete (Eiter–Gottlob). Abnormality theories and prioritised circumscription; the Yale shooting problem as a cautionary example.
6. **Wrap-up** 2 min

Bibliography

OPEN START HERE Christian Strasser and G. Aldo Antonelli. Non-monotonic Logic. *Stanford Encyclopedia of Philosophy*, first published 2001, substantive revision 2024. Circumscription is covered in §3.1.

LIBRARY START HERE Vladimir Lifschitz. Circumscription. In D. M. Gabbay, C. J. Hogger and J. A. Robinson (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning*, Oxford University Press, 298–352, 1994.

OPEN John McCarthy. Circumscription — A Form of Non-Monotonic Reasoning. *Artificial Intelligence* 13(1–2):27–39, 1980. DOI 10.1016/0004-3702(80)90011-9. The author PDF adds an addendum.

OPEN John McCarthy. Applications of Circumscription to Formalizing Common-Sense Knowledge. *Artificial Intelligence* 28(1): 89–116, 1986. DOI 10.1016/0004-3702(86)90032-9.

OPEN Vladimir Lifschitz. Computing Circumscription. *Proc. IJCAI-85*, Vol. 1, 121–127, 1985.

OPEN Thomas Eiter and Georg Gottlob. Propositional Circumscription and Extended Closed-World Reasoning are Π_2^p-Complete. *Theoretical Computer Science* 114(2):231–245, 1993. DOI 10.1016/0304-3975(93)90073-3.

OPEN Patrick Doherty, Witold Łukaszewicz and Andrzej Szałas. Computing Circumscription Revisited: Preliminary Report. *Proc. IJCAI-95*, 1502–1508, 1995. Optional: the DLS elimination algorithm.

TALK 05

Reiter's Default Logic

CORE

Extensions as fixpoints of the Γ operator, existence and semi-monotonicity for normal defaults, and the Σ_2^p/Π_2^p-completeness of credulous and sceptical reasoning.

Outline 45 min

1. **Motivation** 5 min Rules with exceptions ("birds typically fly"). Why classical logic cannot retract conclusions.
2. **Default theories** 8 min Theories $\langle W, D\rangle$ with defaults $\alpha : \beta / \gamma$. Examples: Tweety and the Nixon diamond.
3. **Extensions** 12 min $\Gamma(E)$ is the least deductively closed set that contains W and applies every default whose prerequisite it contains and whose justification is consistent with E. E is an extension iff $\Gamma(E) = E$. The quasi-inductive characterisation $E = \bigcup E_i$. Examples with zero, one and two extensions; $\{\top : \neg p / p\}$ has none.
4. **Properties** 8 min Normal defaults always yield an extension and are semi-monotonic (proof idea). Credulous vs. sceptical consequence.
5. **Complexity** 7 min Credulous reasoning is Σ_2^p-complete and sceptical reasoning Π_2^p-complete (Gottlob). Intuition: guess the generating defaults, then verify with an NP oracle.
6. **Outlook** 5 min Variants (justified and constrained default logic); the link to answer set programming (talk 8).

Bibliography

OPEN START HERE Grigoris Antoniou. A Tutorial on Default Logics. *ACM Computing Surveys* 31(4):337–359, 1999. DOI 10.1145/344588.344602.

OPEN START HERE Christian Strasser and G. Aldo Antonelli. Non-monotonic Logic. *Stanford Encyclopedia of Philosophy*, substantive revision 2024.

LIBRARY Raymond Reiter. A Logic for Default Reasoning. *Artificial Intelligence* 13(1–2):81–132, 1980. DOI 10.1016/0004-3702(80)90014-4.

LIBRARY Georg Gottlob. Complexity Results for Nonmonotonic Logics. *Journal of Logic and Computation* 2(3):397–425, 1992. DOI 10.1093/logcom/2.3.397.

LIBRARY Grigoris Antoniou. Nonmonotonic Reasoning. MIT Press, 1997. ISBN 978-0-262-01157-0.

LIBRARY V. Wiktor Marek and Mirosław Truszczyński. Nonmonotonic Logic: Context-Dependent Reasoning. Springer, 1993. DOI 10.1007/978-3-662-02906-0.

TALK 06

Formal Concept Analysis

CORE

Derivation operators as a Galois connection, the basic theorem on concept lattices, and the Duquenne–Guigues base of attribute implications.

Outline 45 min

1. **Motivation** 5 min Concepts as pairs of extent and intent. A small running example context (e.g. animals × properties).
2. **Contexts and concepts** 10 min A formal context (G, M, I) and the derivation operators $A \mapsto A'$. Galois connection properties ($A \subseteq A''$, $A' = A'''$); the two closure operators. A formal concept is a pair (A, B) with $A' = B$ and $B' = A$; compute all concepts of the example.
3. **Concept lattices** 10 min The order on concepts; formulas for infimum and supremum. The basic theorem: every concept lattice is a complete lattice, and every complete lattice arises as one (proof sketch of one direction). How to read a line diagram.
4. **Implications** 12 min Attribute implications and their validity; Armstrong rules; the link to propositional Horn logic. Pseudo-intents and the Duquenne–Guigues (stem) base, which no implication base can undercut in size. The NextClosure algorithm in lectic order, briefly.
5. **Attribute exploration and applications** 6 min Interactive knowledge acquisition with an expert; applications in data analysis and ontology engineering.
6. **Wrap-up** 2 min

Bibliography

OPEN START HERE Bernhard Ganter, Sebastian Rudolph and Gerd Stumme. Explaining Data with Formal Concept Analysis. In M. Krötzsch and D. Stepanova (eds.), *Reasoning Web: Explainable Artificial Intelligence*, LNCS 11810, Springer, 153–195, 2019. DOI 10.1007/978-3-030-31423-1_5.

OPEN Uta Priss. Formal Concept Analysis in Information Science. *Annual Review of Information Science and Technology* 40(1): 521–543, 2006. DOI 10.1002/aris.1440400120. Informal overview.

LIBRARY Rudolf Wille. Restructuring Lattice Theory: An Approach Based on Hierarchies of Concepts. In I. Rival (ed.), *Ordered Sets*, NATO ASI Series 83, Reidel, 445–470, 1982. DOI 10.1007/978-94-009-7798-3_15. Reprinted in S. Ferré and S. Rudolph (eds.), *ICFCA 2009*, LNAI 5548, 314–339, 2009.

LIBRARY Bernhard Ganter and Rudolf Wille. Formal Concept Analysis: Mathematical Foundations. 2nd edition, Springer, 2024. DOI 10.1007/978-3-031-63422-2. First edition 1999, DOI 10.1007/978-3-642-59830-2.

OPEN Jean-Louis Guigues and Vincent Duquenne. Familles minimales d'implications informatives résultant d'un tableau de données binaires. *Mathématiques et Sciences Humaines* 95:5–18, 1986.

LIBRARY Bernhard Ganter and Sergei Obiedkov. Conceptual Exploration. Springer, 2016. DOI 10.1007/978-3-662-49291-8.

TALK 07

Allen's Interval Algebra

CORE

Qualitative constraint networks over the 13 base relations, path consistency, NP-completeness of consistency, and the tractable ORD-Horn subclass.

Outline 45 min

1. **Motivation** 5 min Qualitative temporal information in text, planning and scheduling ("breakfast before the meeting, the meeting overlaps lunch").
2. **Base relations** 8 min The 13 jointly exhaustive, pairwise disjoint relations via endpoint orderings. Disjunctive relations: $2^{13} = 8192$ relations. Operations: union, intersection, converse, composition.
3. **Constraint networks** 10 min Qualitative constraint networks and the consistency problem. The composition table and the path-consistency (algebraic closure) algorithm in $O(n^3)$; worked example.
4. **Complexity** 10 min Path consistency is incomplete in general. Consistency is NP-hard (Vilain–Kautz) and in NP by guessing an endpoint ordering, hence NP-complete.
5. **Tractability** 10 min Pointisable relations. The ORD-Horn subclass (868 relations, defined by Horn-like clauses over endpoint relations): path consistency decides it, and it is the unique maximal tractable subclass that contains all base relations. Mention that there are exactly 18 maximal tractable subalgebras (Krokhin–Jeavons–Jonsson).
6. **Outlook** 2 min Other qualitative calculi such as RCC8.

Bibliography

OPEN START HERE Jochen Renz and Bernhard Nebel. Qualitative Spatial Reasoning Using Constraint Calculi. In M. Aiello, I. Pratt-Hartmann and J. van Benthem (eds.), *Handbook of Spatial Logics*, Springer, 161–215, 2007. DOI 10.1007/978-1-4020-5587-4_4.

OPEN START HERE Frank Dylla, Jae Hee Lee, Till Mossakowski, Thomas Schneider, André van Delden, Jasper van de Ven and Diedrich Wolter. A Survey of Qualitative Spatial and Temporal Calculi: Algebraic and Computational Properties. *ACM Computing Surveys* 50(1):1–39, 2017. DOI 10.1145/3038927.

OPEN James F. Allen. Maintaining Knowledge about Temporal Intervals. *Communications of the ACM* 26(11):832–843, 1983. DOI 10.1145/182.358434.

OPEN Marc B. Vilain and Henry A. Kautz. Constraint Propagation Algorithms for Temporal Reasoning. *Proc. AAAI-86*, 377–382, 1986.

OPEN Bernhard Nebel and Hans-Jürgen Bürckert. Reasoning about Temporal Relations: A Maximal Tractable Subclass of Allen's Interval Algebra. *Journal of the ACM* 42(1):43–66, 1995. DOI 10.1145/200836.200848. Conference version: *Proc. AAAI-94*, 356–361.

OPEN Andrei Krokhin, Peter Jeavons and Peter Jonsson. Reasoning about Temporal Relations: The Tractable Subalgebras of Allen's Interval Algebra. *Journal of the ACM* 50(5):591–640, 2003. DOI 10.1145/876638.876639.

TALK 08

Answer Set Programming

CORE

The Gelfond–Lifschitz reduct and stable models, guess-and-check modelling, NP-completeness of stable model existence, and the link to default logic.

Outline 45 min

1. **Motivation** 5 min Declarative problem solving: describe solutions, let a solver find them. Puzzle: what should p ⟵ not q and q ⟵ not p mean?
2. **Definite programs** 7 min Rules and Herbrand models; the least model as the least fixpoint of T_P.
3. **Stable models** 12 min The reduct P^M; M is stable iff M is the least model of P^M. Examples with zero, one and two stable models. Stable vs. supported models: the loop p ⟵ p.
4. **Modelling** 10 min Generate–define–test. Constraints and choice rules. A 3-colouring encoding in clingo syntax; grounding in one slide.
5. **Complexity and relations** 8 min Existence of a stable model is NP-complete for normal programs and Σ_2^P-complete for disjunctive ones. Programs with classical negation are a special case of Reiter's default theories, with answer sets matching extensions (Gelfond–Lifschitz 1991).
6. **Outlook** 3 min Solvers and applications; live demo in clingo if time allows.

Bibliography

OPEN START HERE Gerhard Brewka, Thomas Eiter and Mirosław Truszczyński. Answer Set Programming at a Glance. *Communications of the ACM* 54(12):92–103, 2011. DOI 10.1145/2043174.2043195.

OPEN START HERE Vladimir Lifschitz. What Is Answer Set Programming? *Proc. AAAI 2008*, 1594–1597, 2008.

OPEN Michael Gelfond and Vladimir Lifschitz. The Stable Model Semantics for Logic Programming. In R. A. Kowalski and K. A. Bowen (eds.), *Logic Programming: Proc. 5th International Conference and Symposium (ICLP/SLP 1988)*, Vol. 2, MIT Press, 1070–1080, 1988.

OPEN Michael Gelfond and Vladimir Lifschitz. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing* 9(3–4):365–385, 1991. DOI 10.1007/BF03037169.

LIBRARY Vladimir Lifschitz. Answer Set Programming. Springer, 2019. DOI 10.1007/978-3-030-24658-7. An open draft is on the author's teaching page.

LIBRARY Martin Gebser, Roland Kaminski, Benjamin Kaufmann and Torsten Schaub. Answer Set Solving in Practice. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool, 2012. DOI 10.2200/S00457ED1V01Y201211AIM019.

OPEN Potassco. The Potassco Guide (clingo documentation).

LIBRARY Thomas Eiter and Georg Gottlob. On the Computational Cost of Disjunctive Logic Programming: Propositional Case. *Annals of Mathematics and Artificial Intelligence* 15:289–323, 1995. DOI 10.1007/BF01536399.

TALK 09

Abduction

CORE

Logic-based abduction problems, minimality criteria for explanations, and the Eiter–Gottlob complexity landscape.

Outline 45 min

1. **Motivation** 5 min Peirce's triad of deduction, induction and abduction. Diagnosis (the car won't start), medical reasoning, plan recognition.
2. **Logic-based abduction** 10 min A problem $\langle T, H, O\rangle$ of theory, hypotheses and observations. An explanation is $E \subseteq H$ with $T \cup E$ consistent and $T \cup E \models O$. Poole's Theorist framework; worked examples.
3. **Preference criteria** 8 min Subset-minimal, cardinality-minimal, prioritised and penalised explanations. Relevance (a hypothesis occurs in some minimal explanation) and necessity (in all of them).
4. **Complexity** 12 min Existence of an explanation: Σ^p_2-complete for propositional theories, NP-complete for Horn, polynomial for definite Horn. Relevance: Σ^p_2-complete in general, NP-complete for Horn. Necessity: Π^p_2-complete in general, coNP-complete for Horn. Proof idea of NP-hardness for Horn.
5. **Variants and applications** 7 min Abductive logic programming with integrity constraints. Abductive vs. consistency-based diagnosis.
6. **Wrap-up** 3 min Contrast with induction (talk 11).

Bibliography

OPEN START HERE David L. Poole and Alan K. Mackworth. Abduction. Section 5.8 in *Artificial Intelligence: Foundations of Computational Agents*, 3rd edition, Cambridge University Press, 2023. Free online edition.

OPEN START HERE Igor Douven. Abduction. *Stanford Encyclopedia of Philosophy*, first published 2011, substantive revision 2025.

OPEN David Poole. Explanation and Prediction: An Architecture for Default and Abductive Reasoning. *Computational Intelligence* 5(2):97–110, 1989. DOI 10.1111/j.1467-8640.1989.tb00319.x.

OPEN Thomas Eiter and Georg Gottlob. The Complexity of Logic-Based Abduction. *Journal of the ACM* 42(1):3–42, 1995. DOI 10.1145/200836.200838.

OPEN Bart Selman and Hector J. Levesque. Abductive and Default Reasoning: A Computational Core. *Proc. AAAI-90*, 343–348, 1990.

LIBRARY Antonis C. Kakas, Robert A. Kowalski and Francesca Toni. Abductive Logic Programming. *Journal of Logic and Computation* 2(6):719–770, 1992. DOI 10.1093/logcom/2.6.719.

LIBRARY Luca Console, Daniele Theseider Dupré and Pietro Torasso. On the Relationship between Abduction and Deduction. *Journal of Logic and Computation* 1(5):661–690, 1991. DOI 10.1093/logcom/1.5.661.

LIBRARY Gabriele Paul. Approaches to Abductive Reasoning: An Overview. *Artificial Intelligence Review* 7(2):109–152, 1993. DOI 10.1007/BF00849080.

TALK 10

Preferential and Rational Consequence (KLM)

CORE

The rules of System P, preferential models and the representation theorem, rational monotonicity, and rational closure.

Outline 45 min

1. **Motivation** 5 min Instead of one more formalism, ask which properties any reasonable nonmonotonic consequence relation |\~ should have. Running example: birds, penguins, flying.
2. **Preferential semantics** 8 min Shoham's idea: entail what holds in all *preferred* models. Link back to circumscription (talk 4).
3. **Systems C and P** 10 min Reflexivity, Left Logical Equivalence, Right Weakening, Cut, Cautious Monotonicity, plus Or for System P. Derived rules such as And. What fails: monotonicity, transitivity, contraposition.
4. **Preferential models** 10 min States, labelling function and a smooth strict partial order; the induced relation |\~. Representation theorem: preferential consequence relations are exactly those defined by preferential models (proof idea).
5. **Rational monotonicity and rational closure** 10 min The rule RM and ranked models. Conditional knowledge bases; the exceptionality ranking and the rational closure construction, which needs polynomially many classical entailment checks. Penguin example; connection to Pearl's System Z.
6. **Wrap-up** 2 min

Bibliography

OPEN START HERE Christian Strasser and G. Aldo Antonelli. Non-monotonic Logic. *Stanford Encyclopedia of Philosophy*, substantive revision 2024. Sections on System P, preferential models and rational closure.

OPEN START HERE Adam Kaliski. An Overview of KLM-Style Defeasible Entailment. MSc thesis, University of Cape Town, 2020. Tutorial-style, includes a rational closure algorithm.

OPEN Sarit Kraus, Daniel Lehmann and Menachem Magidor. Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artificial Intelligence* 44(1–2):167–207, 1990. DOI 10.1016/0004-3702(90)90101-5.

OPEN Daniel Lehmann and Menachem Magidor. What Does a Conditional Knowledge Base Entail? *Artificial Intelligence* 55(1):1–60, 1992. DOI 10.1016/0004-3702(92)90041-U.

OPEN Yoav Shoham. Nonmonotonic Logics: Meaning and Utility. *Proc. IJCAI-87*, 388–393, 1987.

OPEN Judea Pearl. System Z: A Natural Ordering of Defaults with Tractable Applications to Nonmonotonic Reasoning. In R. Parikh (ed.), *Proc. TARK-90*, Morgan Kaufmann, 121–135, 1990.

LIBRARY David Makinson. How to Go Nonmonotonic. In D. M. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, 2nd edition, Vol. 12, Springer, 175–278, 2005. DOI 10.1007/1-4020-3092-4_3.

TALK 11

Inductive Logic Programming

CORE
The learning-from-entailment setting, θ-subsumption as a generality order, Plotkin's least general generalisation, and top-down vs. bottom-up search.

Outline 45 min

1. **Motivation** 5 min Learning interpretable relational rules from few examples plus background knowledge. Running example: learning grandparent/2.
2. **Learning setting** 8 min Given B, E^+ and E^-, find H with $B \cup H \models e$ for all $e \in E^+$ and $B \cup H \not\models e$ for all $e \in E^-$. Hypothesis space of definite clauses; language bias and mode declarations.
3. **Generality orders** 12 min θ-subsumption and why deciding it is NP-complete. It is strictly weaker than implication, with the gap arising only for self-resolving clauses. Plotkin's lgg of atoms and clauses (anti-unification), worked example; every finite set of clauses has an lgg. Relative lgg.
4. **Search** 10 min Top-down refinement operators (FOIL-style) vs. bottom-up generalisation. Inverse entailment and the bottom clause in Progol, as a sketch.
5. **Modern view** 7 min Meta-interpretive learning, ASP-based systems, neuro-symbolic approaches; limitations.
6. **Wrap-up** 3 min Contrast with abduction (talk 9).

Bibliography

OPEN START HERE Andrew Cropper and Sebastijan Dumančić. Inductive Logic Programming at 30: A New Introduction. *Journal of Artificial Intelligence Research* 74:765–850, 2022. DOI 10.1613/jair.1.13507.

OPEN Gordon D. Plotkin. A Note on Inductive Generalization. In *Machine Intelligence 5*, Edinburgh University Press, 153–163, 1970. Scan on the author's homepage.

OPEN Stephen Muggleton. Inductive Logic Programming. *New Generation Computing* 8(4):295–318, 1991. DOI 10.1007/BF03037089.

OPEN Stephen Muggleton and Luc De Raedt. Inductive Logic Programming: Theory and Methods. *Journal of Logic Programming* 19–20:629–679, 1994. DOI 10.1016/0743-1066(94)90035-3.

OPEN Stephen Muggleton. Inverse Entailment and Progol. *New Generation Computing* 13(3–4):245–286, 1995. DOI 10.1007/BF03037227.

LIBRARY J. Ross Quinlan. Learning Logical Definitions from Relations. *Machine Learning* 5(3):239–266, 1990. DOI 10.1007/BF00117105.

LIBRARY Shan-Hwei Nienhuys-Cheng and Ronald de Wolf. Foundations of Inductive Logic Programming. LNAI 1228, Springer, 1997. DOI 10.1007/3-540-62927-0.

TALK 12

Forgetting

CORE

Propositional forgetting and its characterisations, formula–variable independence, and first-order forgetting as second-order quantification that need not be first-order definable.

Outline 45 min

1. **Motivation** 5 min Restricting a knowledge base to a relevant vocabulary: reuse, privacy, efficiency. Forgetting at the knowledge level vs. deleting formulas.
2. **Propositional forgetting** 12 min Model-theoretic definition (models that agree except on p). Syntactic form forget(φ, p) ≡ φ[p/⊤] ∨ φ[p/⊥], i.e. ∃p.φ. It is the strongest consequence of φ not mentioning p. Iterating over a set of variables and the size blow-up; worked example.
3. **Independence** 8 min Formula–variable and formula–literal independence; literal forgetting; how independence relates to forgetting (Lang–Liberatore–Marquis).
4. **First-order forgetting** 12 min Lin–Reiter: forgetting a ground atom and a predicate; forget(T; P) ≡ ∃R.T[P/R] in second-order logic, which is not always first-order definable. Link to uniform interpolation and to second-order quantifier elimination (Ackermann's lemma).
5. **Beyond classical logic** 6 min Forgetting in logic programs and description logics; for ALC, deciding whether a uniform interpolant exists is 2EXPTIME-complete (Lutz–Wolter).
6. **Wrap-up** 2 min

Bibliography

OPEN START HERE Thomas Eiter and Gabriele Kern-Isberner. A Brief Survey on Forgetting from a Knowledge Representation and Reasoning Perspective. *KI – Künstliche Intelligenz* 33(1):9–33, 2019. DOI 10.1007/s13218-018-0564-6.

OPEN Fangzhen Lin and Raymond Reiter. Forget It! In *Relevance: Papers from the 1994 AAAI Fall Symposium*, Technical Report FS-94-02, AAAI Press, 141–146, 1994.

OPEN Jérôme Lang, Paolo Liberatore and Pierre Marquis. Propositional Independence: Formula-Variable Independence and Forgetting. *Journal of Artificial Intelligence Research* 18:391–443, 2003. DOI 10.1613/jair.1113.

OPEN James P. Delgrande. A Knowledge Level Account of Forgetting. *Journal of Artificial Intelligence Research* 60:1165–1213, 2017. DOI 10.1613/jair.5530.

OPEN Carsten Lutz and Frank Wolter. Foundations for Uniform Interpolation and Forgetting in Expressive Description Logics. *Proc. IJCAI 2011*, 989–995, 2011.

TALK 13

Probabilistic Logic Programming

CORE

Sato's distribution semantics for probabilistic facts, and exact inference by reduction to weighted model counting.

Outline 45 min

1. **Motivation** 5 min Uncertain relational knowledge, e.g. link discovery in probabilistic biological networks. Why recursion over relations is awkward in Bayesian networks.
2. **Background** 5 min Definite programs and their least Herbrand model.
3. **Distribution semantics** 12 min Probabilistic facts p::f; total choices and the probability of a world; P(q) is the total probability of the total choices whose least model entails q. Worked example (burglary/alarm or coin tosses). Annotated disjunctions; Sato's measure-theoretic extension to infinitely many facts, briefly.
4. **Inference** 12 min Explanations as a DNF; computing its probability is the disjoint-sum problem, which is #P-hard. Compile to d-DNNF or SDD and do weighted model counting; conditioning on evidence.
5. **Learning and relatives** 8 min Parameter learning by EM (Sato's PRISM). Relation to Poole's probabilistic Horn abduction; contrast with Markov logic.
6. **Demo and wrap-up** 3 min Run a query in the online ProbLog editor.

Bibliography

OPEN START HERE Luc De Raedt and Angelika Kimmig. Probabilistic (Logic) Programming Concepts. *Machine Learning* 100(1):5–47, 2015. DOI 10.1007/s10994-015-5494-z.

OPEN START HERE DTAI, KU Leuven. ProbLog homepage, with an online editor, a first tutorial lesson and the documentation.

OPEN Taisuke Sato. A Statistical Learning Method for Logic Programs with Distribution Semantics. *Proc. 12th International Conference on Logic Programming (ICLP 1995)*, MIT Press, 715–729, 1995. Scanned extended version on the author's page.

OPEN Luc De Raedt, Angelika Kimmig and Hannu Toivonen. ProbLog: A Probabilistic Prolog and Its Application in Link Discovery. *Proc. IJCAI 2007*, 2468–2473, 2007.

OPEN Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens and Luc De Raedt. Inference and Learning in Probabilistic Logic Programs Using Weighted Boolean Formulas. *Theory and Practice of Logic Programming* 15(3):358–401, 2015. DOI 10.1017/S1471068414000076.

OPEN David Poole. Probabilistic Horn Abduction and Bayesian Networks. *Artificial Intelligence* 64(1):81–129, 1993. DOI 10.1016/0004-3702(93)90061-F.

LIBRARY Fabrizio Riguzzi. Foundations of Probabilistic Logic Programming: Languages, Semantics, Inference and Learning. 2nd edition, River Publishers, 2023. DOI 10.1201/9781003427421.

TALK 14

CDCL SAT Solving

CORE
Conflict analysis on the implication graph, first-UIP clause learning with non-chronological backjumping, and CDCL as a resolution proof system.

Outline 45 min

1. **Motivation** 5 min SAT is NP-complete, yet industrial instances with millions of clauses are solved routinely. SAT solvers power ASP, planning and verification.
2. **From DP to DPLL** 8 min CNF; unit propagation and pure literals; DP (variable elimination by resolution) vs. DPLL (backtracking search). Trail and decision levels.
3. **Conflict analysis** 15 min The implication graph; conflicts and cuts; unique implication points (from GRASP). The first-UIP learned clause (worked example) and non-chronological backjumping. Each learned clause is derivable by resolution from the input and previously learned clauses.
4. **Engineering** 7 min Two watched literals, VSIDS branching, restarts, clause deletion (Chaff, MiniSat).
5. **Proof-complexity view** 7 min Resolution as a proof system. Nondeterministic CDCL with restarts polynomially simulates general resolution, so resolution lower bounds (e.g. pigeonhole) apply to CDCL.
6. **Wrap-up** 3 min Incremental SAT, MaxSAT, model counting (link to talk 18).

Bibliography

OPEN START HERE Carla P. Gomes, Henry Kautz, Ashish Sabharwal and Bart Selman. Satisfiability Solvers. Chapter 2 in F. van Harmelen, V. Lifschitz and B. Porter (eds.), *Handbook of Knowledge Representation*, Elsevier, 89–134, 2008. DOI 10.1016/S1574-6526(07)03002-7.

LIBRARY START HERE João Marques-Silva, Inês Lynce and Sharad Malik. Conflict-Driven Clause Learning SAT Solvers. Chapter 4 in A. Biere, M. Heule, H. van Maaren and T. Walsh (eds.), *Handbook of Satisfiability*, 2nd edition, FAIA 336, IOS Press, 133–182, 2021. DOI 10.3233/FAIA200987.

OPEN Martin Davis and Hilary Putnam. A Computing Procedure for Quantification Theory. *Journal of the ACM* 7(3):201–215, 1960. DOI 10.1145/321033.321034.

OPEN Martin Davis, George Logemann and Donald Loveland. A Machine Program for Theorem-Proving. *Communications of the ACM* 5(7):394–397, 1962. DOI 10.1145/368273.368557.

OPEN João P. Marques-Silva and Karem A. Sakallah. GRASP: A Search Algorithm for Propositional Satisfiability. *IEEE Transactions on Computers* 48(5):506–521, 1999. DOI 10.1109/12.769433.

OPEN Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang and Sharad Malik. Chaff: Engineering an Efficient SAT Solver. *Proc. 38th Design Automation Conference (DAC 2001)*, ACM, 530–535, 2001. DOI 10.1145/378239.379017.

OPEN Lintao Zhang, Conor F. Madigan, Matthew W. Moskewicz and Sharad Malik. Efficient Conflict Driven Learning in a Boolean Satisfiability Solver. *Proc. ICCAD 2001*, IEEE, 279–285, 2001. DOI 10.1109/ICCAD.2001.968634. Establishes first-UIP as the best learning scheme.

OPEN Paul Beame, Henry Kautz and Ashish Sabharwal. Towards Understanding and Harnessing the Potential of Clause Learning. *Journal of Artificial Intelligence Research* 22:319–351, 2004. DOI 10.1613/jair.1410.

OPEN Knot Pipatsrisawat and Adnan Darwiche. On the Power of Clause-Learning SAT Solvers as Resolution Engines. *Artificial Intelligence* 175(2):512–525, 2011. DOI 10.1016/j.artint.2010.10.002.

OPEN Sam Buss and Jakob Nordström. Proof Complexity and SAT Solving. Chapter 7 in *Handbook of Satisfiability*, 2nd edition, FAIA 336, IOS Press, 233–350, 2021. DOI 10.3233/FAIA200990. Optional.

TALK 15

Inconsistent Query Answering

CORE

Subset repairs, the AR, IAR and brave semantics, and the data complexity of consistent query answering, up to the FO / L / coNP trichotomy.

Outline 45 min

1. **Motivation** 5 min Integrated data violates keys (an employee with two salaries). Cleaning throws information away; instead, answer queries robustly.
2. **Repairs** 10 min Integrity constraints: keys, functional dependencies, denial constraints. Subset repairs as maximal consistent subsets; the conflict hypergraph. Databases with exponentially many repairs.
3. **Semantics** 8 min AR (true in every repair, the consistent answers), brave (true in some repair), IAR (true in the intersection of all repairs). IAR answers ⊆ AR answers ⊆ brave answers.
4. **Complexity** 12 min Data complexity. A single key constraint and a fixed conjunctive query can already make AR answering coNP-complete. FO-rewritable classes (Fuxman–Miller). For self-join-free conjunctive queries under primary keys each query is FO-rewritable, L-complete or coNP-complete (Koutris–Wijsen).
5. **With ontologies** 8 min Constraints given by a lightweight ontology (DL-Lite): AR is coNP-complete, while IAR is FO-rewritable. Keep the logic minimal, since description logics are covered in the parallel lecture.
6. **Wrap-up** 2 min

Bibliography

OPEN START HERE Jef Wijsen. Foundations of Query Answering on Inconsistent Databases. *SIGMOD Record* 48(3):6–16, 2019. DOI 10.1145/3377391.3377393.

OPEN START HERE · ONTOLOGIES Meghyn Bienvenu and Camille Bourgaux. Inconsistency-Tolerant Querying of Description Logic Knowledge Bases. In *Reasoning Web 2016*, LNCS 9885, Springer, 156–202, 2017. DOI 10.1007/978-3-319-49493-7_5.

OPEN Marcelo Arenas, Leopoldo Bertossi and Jan Chomicki. Consistent Query Answers in Inconsistent Databases. *Proc. PODS 1999*, ACM, 68–79, 1999. DOI 10.1145/303976.303983.

OPEN Jan Chomicki and Jerzy Marcinkowski. Minimal-Change Integrity Maintenance Using Tuple Deletions. *Information and Computation* 197(1–2):90–121, 2005. DOI 10.1016/j.ic.2004.04.007.

OPEN Ariel Fuxman and Renée J. Miller. First-Order Query Rewriting for Inconsistent Databases. *Journal of Computer and System Sciences* 73(4):610–635, 2007. DOI 10.1016/j.jcss.2006.10.013.

OPEN Paraschos Koutris and Jef Wijsen. Consistent Query Answering for Self-Join-Free Conjunctive Queries under Primary Key Constraints. *ACM Transactions on Database Systems* 42(2):1–45, 2017. DOI 10.1145/3068334.

OPEN Paraschos Koutris and Jef Wijsen. Consistent Query Answering for Primary Keys in Logspace. *Proc. ICDT 2019*, LIPIcs 127, 23:1–23:19, 2019. DOI 10.4230/LIPIcs.ICDT.2019.23.

OPEN Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi and Domenico Fabio Savo. Inconsistency-Tolerant Semantics for Description Logics. *Proc. RR 2010*, LNCS 6333, Springer, 103–117, 2010. DOI 10.1007/978-3-642-15918-3_9.

LIBRARY Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi and Domenico Fabio Savo. Query Rewriting for Inconsistent DL-Lite Ontologies. *Proc. RR 2011*, LNCS 6902, Springer, 155–169, 2011. DOI 10.1007/978-3-642-23580-1_12.

OPEN Leopoldo Bertossi. Database Repairs and Consistent Query Answering: Origins and Further Developments. *Proc. PODS 2019*, ACM, 48–58, 2019. DOI 10.1145/3294052.3322190.

LIBRARY Leopoldo Bertossi. Database Repairing and Consistent Query Answering. Synthesis Lectures on Data Management, Morgan & Claypool, 2011. DOI 10.1007/978-3-031-01883-1.

TALK 16

CP-Nets

CORE

Ceteris paribus semantics via improving flips, the linear-time forward sweep for acyclic nets, and the complexity of dominance testing.

Outline 45 min

1. **Motivation** 5 min Preferences over combinatorial domains (a menu, a product configuration) have 2^n outcomes. People state preferences locally and conditionally.
2. **Ceteris paribus preferences** 8 min Conditional preference tables over parent variables; a CP-net as a directed graph plus tables. Semantics: the induced order is the transitive closure of improving flips. Running dinner example (soup, wine, main course).
3. **Consistency and optimisation** 10 min Every acyclic CP-net is consistent. The forward sweep finds the unique optimal outcome in linear time. Optimisation under hard constraints, briefly.
4. **Dominance** 12 min Dominance as the existence of an improving flipping sequence. Polynomial for binary-valued trees and polytrees, NP-complete for directed-path singly connected nets, PSPACE-complete in general (Goldsmith et al.). Ordering queries are easier than dominance.
5. **Extensions** 7 min TCP-nets and utility-based alternatives; elicitation and learning of CP-nets.
6. **Wrap-up** 3 min Link to preference aggregation (talk 19).

Bibliography

OPEN START HERE Ronen I. Brafman and Carmel Domshlak. Preference Handling — An Introductory Tutorial. *AI Magazine* 30(1):58–86, 2009. DOI 10.1609/aimag.v30i1.2114.

OPEN START HERE Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos and David Poole. CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. *Journal of Artificial Intelligence Research* 21:135–191, 2004. DOI 10.1613/jair.1234.

OPEN Craig Boutilier, Ronen I. Brafman, Holger H. Hoos and David Poole. Reasoning with Conditional Ceteris Paribus Preference Statements. *Proc. UAI 1999*, 71–80, 1999.

OPEN Judy Goldsmith, Jérôme Lang, Mirosław Truszczyński and Nic Wilson. The Computational Complexity of Dominance and Consistency in CP-Nets. *Journal of Artificial Intelligence Research* 33:403–432, 2008. DOI 10.1613/jair.2627.

LIBRARY Carmel Domshlak, Eyke Hüllermeier, Souhila Kaci and Henri Prade. Preferences in AI: An Overview. *Artificial Intelligence* 175(7–8):1037–1052, 2011. DOI 10.1016/j.artint.2011.03.004.

LIBRARY Francesca Rossi, Kristen Brent Venable and Toby Walsh. A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool, 2011. DOI 10.1007/978-3-031-01556-4.

TALK 17

Actual Causality

CORE

Structural equation models with interventions, the Halpern–Pearl definition (conditions AC1–AC3), and how its three versions differ in complexity.

Outline 45 min

1. **Motivation** 5 min Legal and moral responsibility, explaining failures, explainable AI. Suzy and Billy throw rocks at a bottle: the naive “but-for” test fails under preemption and overdetermination.
2. **Structural equation models** 10 min Exogenous and endogenous variables, structural equations, recursive models, contexts. Interventions $X \leftarrow x$ and the semantics of $[X \leftarrow x]\varphi$; causal networks.
3. **The HP definition** 15 min The modified (2015) definition: AC1 (actuality), AC2 (counterfactual dependence while a witness set W is held at its actual values), AC3 (minimality). Worked examples: forest fire (conjunctive and disjunctive), rock throwing with the hit variables SH and BH, double prevention.
4. **Versions and complexity** 8 min Original (2001), updated (2005) and modified (2015) definitions. Deciding causality is Σ_2^p-complete, D_2^p-complete and D^p-complete respectively.
5. **Responsibility and blame** 5 min Degree of responsibility $1/(k+1)$, where k is the least number of changes that make the outcome counterfactually depend on the candidate (Chockler–Halpern).
6. **Wrap-up** 2 min

Bibliography

OPEN START HERE Joseph Y. Halpern. Actual Causality. MIT Press, 2016. DOI 10.7551/mitpress/10809.001.0001. Open access (CC BY-NC-ND 4.0).

OPEN START HERE Christopher Hitchcock. Causal Models. *Stanford Encyclopedia of Philosophy*, first published 2018. Section 3.3 covers actual causation.

OPEN Joseph Y. Halpern and Judea Pearl. Causes and Explanations: A Structural-Model Approach. Part I: Causes. *British Journal for the Philosophy of Science* 56(4):843–887, 2005. DOI 10.1093/bjps/axi147.

OPEN Joseph Y. Halpern. A Modification of the Halpern-Pearl Definition of Causality. *Proc. IJCAI 2015*, 3022–3033, 2015.

OPEN Hana Chockler and Joseph Y. Halpern. Responsibility and Blame: A Structural-Model Approach. *Journal of Artificial Intelligence Research* 22:93–115, 2004. DOI 10.1613/jair.1391.

OPEN Gadi Aleksandrowicz, Hana Chockler, Joseph Y. Halpern and Alexander Ivrii. The Computational Complexity of Structure-Based Causality. *Journal of Artificial Intelligence Research* 58:431–451, 2017. DOI 10.1613/jair.5229.

LIBRARY Thomas Eiter and Thomas Lukasiewicz. Complexity Results for Structure-Based Causality. *Artificial Intelligence* 142(1): 53–89, 2002. DOI 10.1016/S0004-3702(02)00271-0.

LIBRARY Judea Pearl. Causality: Models, Reasoning, and Inference. 2nd edition, Cambridge University Press, 2009. DOI 10.1017/CBO9780511803161. Chapter 10: “The Actual Cause”.

TALK 18

Knowledge Compilation Map

CORE

The NNF language family (DNNF, d-DNNF, OBDD, …), which queries and transformations each supports in polynomial time, and the succinctness hierarchy.

Outline 45 min

1. **Motivation** 5 min Split reasoning into an expensive offline compilation and cheap online queries. Examples: product configuration, probabilistic inference.
2. **Languages** 12 min Negation normal form and its properties: decomposability (DNNF), determinism, smoothness; d-DNNF. FBDD and OBDD (with a fixed variable order) as special cases; CNF, DNF, prime implicates and prime implicants for comparison. OBDDs are canonical.
3. **Queries and transformations** 12 min Consistency (CO), validity (VA), clausal entailment (CE), implicants (IM), equivalence (EQ), sentential entailment (SE), model counting (CT), model enumeration (ME); conditioning, forgetting and the Boolean operations. Show the table for DNNF, d-DNNF and OBDD. Proofs on the board: CO on DNNF, CT by bottom-up counting on smooth d-DNNF.
4. **Succinctness** 8 min Definition. Each added property (determinism, decision, ordering) makes a language less succinct; some separations hold only unless the polynomial hierarchy collapses. The succinctness–tractability trade-off.
5. **Compilers and applications** 6 min Compilers such as c2d; sentential decision diagrams (SDDs); weighted model counting for probabilistic inference (link to talk 13).
6. **Wrap-up** 2 min

Bibliography

OPEN START HERE Adnan Darwiche and Pierre Marquis. A Knowledge Compilation Map. *Journal of Artificial Intelligence Research* 17:229–264, 2002. DOI 10.1613/jair.989.

OPEN START HERE Adnan Darwiche. Three Modern Roles for Logic in AI. *Proc. PODS 2020*, ACM, 229–243, 2020. DOI 10.1145/3375395.3389131.

OPEN Adnan Darwiche. Decomposable Negation Normal Form. *Journal of the ACM* 48(4):608–647, 2001. DOI 10.1145/502090.502091. Conference precursor: Compiling Knowledge into Decomposable Negation Normal Form, *Proc. IJCAI 1999*, 284–289.

OPEN Bart Selman and Henry Kautz. Knowledge Compilation and Theory Approximation. *Journal of the ACM* 43(2):193–224, 1996. DOI 10.1145/226643.226644.

OPEN Randal E. Bryant. Graph-Based Algorithms for Boolean Function Manipulation. *IEEE Transactions on Computers* C-35(8): 677–691, 1986. DOI 10.1109/TC.1986.1676819.

OPEN Adnan Darwiche. SDD: A New Canonical Representation of Propositional Knowledge Bases. *Proc. IJCAI 2011*, 819–826, 2011.

LIBRARY Marco Cadoli and Francesco M. Donini. A Survey on Knowledge Compilation. *AI Communications* 10(3–4):137–150, 1997. DOI 10.3233/EAI-1997-133.

TALK 19

Judgment Aggregation

CORE

Agendas and aggregation rules, the List–Pettit impossibility theorem with proof, and the median property that characterises when majority rule is consistent.

Outline 45 min

1. **Motivation** 5 min Three judges decide a breach-of-contract case: majority on the premises and majority on the conclusion disagree (the doctrinal paradox). The same issue arises in committees and multi-agent systems.
2. **Framework** 8 min An agenda closed under complement; complete and consistent judgment sets; profiles and aggregation functions. Majority rule; premise-based vs. conclusion-based procedures.
3. **Axioms and impossibility** 15 min Universal domain, collective rationality, anonymity, systematicity (neutrality plus independence). The List–Pettit theorem for the agenda $\{p, q, p \wedge q\}$ and negations, with the full proof. Which assumption to relax, and what each relaxation buys.
4. **Agenda characterisations** 8 min Majority rule is consistent for an agenda iff every minimally inconsistent subset has size at most 2 (the median property; Nehring–Puppe; odd number of agents). Mention the Arrow-style theorem of Dietrich–List.
5. **Computational aspects** 6 min Rules such as Kemeny aggregation and the complexity of computing outcomes (Endriss–Grandi–Porello).
6. **Wrap-up** 3 min Relation to voting and to belief merging.

Bibliography

OPEN START HERE Ulle Endriss. Judgment Aggregation. Chapter 17 in F. Brandt, V. Conitzer, U. Endriss, J. Lang and A. D. Procaccia (eds.), *Handbook of Computational Social Choice*, Cambridge University Press, 399–426, 2016. DOI 10.1017/CBO9781107446984.018. The whole handbook is free from an editor's site.

OPEN START HERE Christian List. The Theory of Judgment Aggregation: An Introductory Review. *Synthese* 187(1):179–207, 2012. DOI 10.1007/s11229-011-0025-3.

OPEN Lewis A. Kornhauser and Lawrence G. Sager. Unpacking the Court. *Yale Law Journal* 96(1):82–117, 1986.

OPEN Christian List and Philip Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy* 18(1):89–110, 2002. DOI 10.1017/S0266267102001098.

LIBRARY Franz Dietrich and Christian List. Arrow's Theorem in Judgment Aggregation. *Social Choice and Welfare* 29(1):19–33, 2007. DOI 10.1007/s00355-006-0196-x.

LIBRARY Klaus Nehring and Clemens Puppe. The Structure of Strategy-Proof Social Choice — Part I: General Characterization and Possibility Results on Median Spaces. *Journal of Economic Theory* 135(1):269–305, 2007. DOI 10.1016/j.jet.2006.04.008.

OPEN Ulle Endriss, Umberto Grandi and Daniele Porello. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research* 45:481–514, 2012. DOI 10.1613/jair.3708.

LIBRARY Davide Grossi and Gabriella Pigozzi. Judgment Aggregation: A Primer. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool, 2014. DOI 10.1007/978-3-031-01568-7.

TALK 20

LTL on Finite Traces

CORE

Finite-trace semantics of LTL, translation to automata, PSPACE-completeness of satisfiability, and 2EXPTIME-completeness of synthesis.

Outline 45 min

1. **Motivation** 5 min Many specifications describe finite behaviours: temporally extended goals in planning, declarative business processes (Declare), non-Markovian rewards in reinforcement learning.
2. **LTL recap** 6 min Syntax with X, U, F, G and semantics over infinite traces, briefly.
3. **LTLf** 10 min Semantics over finite traces; strong vs. weak next; the formula "last". Surprises: on finite traces GF p and FG p both say that p holds at the last position. LTLf is as expressive as first-order logic over finite linear orders (star-free regular languages).
4. **Automata** 12 min Linear-size alternating automata, NFAs of exponential size, DFAs of doubly exponential size; construction sketch. LDLf reaches all regular languages (MSO).
5. **Reasoning** 8 min Satisfiability, validity and implication are PSPACE-complete. When the choice between finite and infinite semantics does not matter ("insensitivity to infiniteness"). Reactive synthesis via DFA games is 2EXPTIME-complete.
6. **Wrap-up** 4 min Back to the applications from the start.

Bibliography

OPEN START HERE Giuseppe De Giacomo and Moshe Y. Vardi. Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. *Proc. IJCAI 2013*, 854–860, 2013.

OPEN Giuseppe De Giacomo and Moshe Y. Vardi. Synthesis for LTL and LDL on Finite Traces. *Proc. IJCAI 2015*, 1558–1564, 2015.

OPEN Giuseppe De Giacomo, Riccardo De Masellis and Marco Montali. Reasoning on LTL on Finite Traces: Insensitivity to Infiniteness. *Proc. AAAI 2014*, 1027–1033, 2014. DOI 10.1609/aaai.v28i1.8872.

OPEN Valeria Fionda and Gianluigi Greco. LTL on Finite and Process Traces: Complexity Results and a Practical Reasoner. *Journal of Artificial Intelligence Research* 63:557–623, 2018. DOI 10.1613/jair.1.11256.

LIBRARY Amir Pnueli. The Temporal Logic of Programs. *Proc. 18th Annual Symposium on Foundations of Computer Science (FOCS 1977)*, IEEE, 46–57, 1977. DOI 10.1109/SFCS.1977.32.

OPEN Moshe Y. Vardi and Pierre Wolper. An Automata-Theoretic Approach to Automatic Program Verification (Preliminary Report). *Proc. LICS 1986*, IEEE, 332–344, 1986. Scan on the author's page.

OPEN Jorge A. Baier and Sheila A. McIlraith. Planning with First-Order Temporally Extended Goals Using Heuristic Search. *Proc. AAAI 2006*, 788–795, 2006.

OPEN Maja Pesic, Helen Schonenberg and Wil M. P. van der Aalst. DECLARE: Full Support for Loosely-Structured Processes. *Proc. 11th IEEE International Enterprise Distributed Object Computing Conference (EDOC 2007)*, IEEE, 2007. DOI 10.1109/EDOC.2007.14.

OPEN Alberto Camacho, Rodrigo Toro Icarte, Toryn Q. Klassen, Richard Valenzano and Sheila A. McIlraith. LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. *Proc. IJCAI 2019*, 6065–6073, 2019. DOI 10.24963/ijcai.2019/840.