Temporal Probabilistic Relational Models and Beyond

Marcel Gehrke

Institute of Information Systems
University of Lübeck

January 30, 2023

Artificial Intelligence: An Agent Perspective

Russell and Norvig (2020)

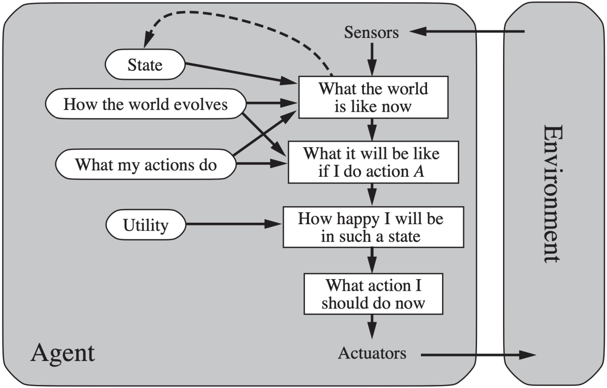

Artificial Intelligence: An Agent Perspective

Russell and Norvig (2020)

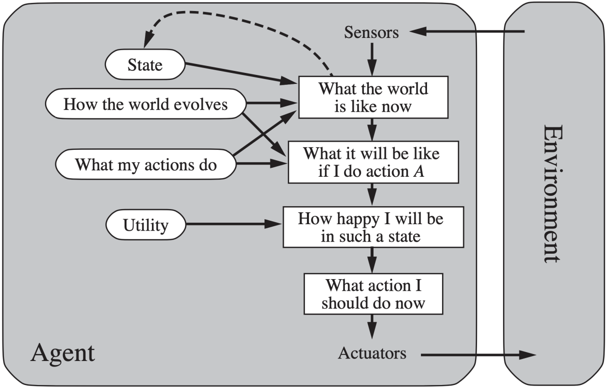

Knowledge representation and reasoning under uncertainty

$\rightarrow$ Statistical Relational AI

Probabilistic Temporal Relational and Lifted Models

Murphy (2002), Poole (2003), Ahmadi et al. (2013)

Parfactor graph G: Compact encoding of full joint d. $P_G = \frac{1}{Z} \prod_{f \in gr(u(G))} f$

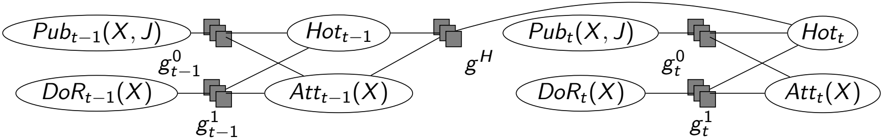

Probabilistic Temporal Relational and Lifted Models

Murphy (2002), Poole (2003), Ahmadi et al. (2013)

Parfactor graph G: Compact encoding of full joint d. $P_G = \frac{1}{Z}\prod_{f \in gr(u(G))} f$

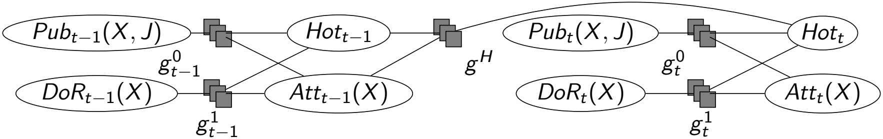

Marginal distribution query: $P(A^i_\pi|E_{0:t})$ w.r.t. the model:

- Prediction: $\pi > t$ (is the topic hot in $\pi - t$ days?)
- Filtering: $\pi = t$ (is the topic hot today?)
- Hindsight: $\pi < t$ (was the topic hot $t - \pi$ days ago?)

Probabilistic Temporal Relational and Lifted Models

Murphy (2002), Poole (2003), Ahmadi et al. (2013)

Parfactor graph G: Compact encoding of full joint d. $P_G = \frac{1}{Z} \prod_{f \in gr(u(G))} f$

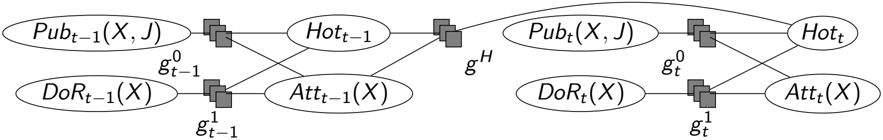

QA: Eliminate all non-query variables
while **avoiding grounding and unrolling** G **as well as building** P_G

QA: Lifted Variable Elimination (LVE)

Poole (2003), de Salvo Braz et al. (2005), Milch et al. (2008), Taghipour et al. (2013)

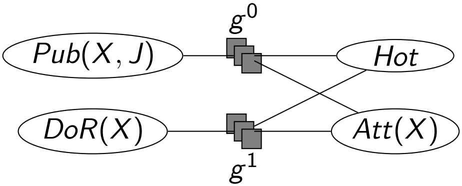

$P(DoR(eve))$

$\sum_V$ indicates a sum over the values of V, $|X|$ a domain size

QA: Lifted Variable Elimination (LVE)

Poole (2003), de Salvo Braz et al. (2005), Milch et al. (2008), Taghipour et al. (2013)

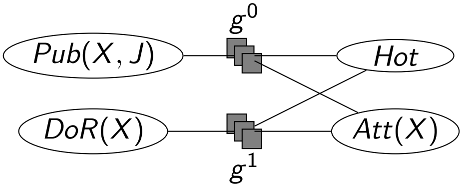

$$P(DoR(eve)) \propto \sum_{Hot} \left(\sum_{\substack{DoR(X)\\X \neq eve}} \sum_{Att(X)} g^1 \left(\sum_{Pub(X,J)} g^0 \right)^{|J|} \right)^{|X|_{X \neq eve}}$$

$\sum_V$ indicates a sum over the values of V, $|X|$ a domain size

QA: Lifted Junction Tree Algorithm (LJT)

Lauritzen and Spiegelhalter (1988), Braun and Möller (2016)

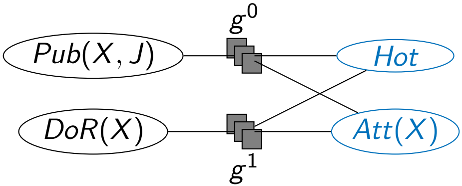

QA based on submodels
ensured to be independent

QA: Lifted Junction Tree Algorithm (LJT)

Lauritzen and Spiegelhalter (1988), Braun and Möller (2016)

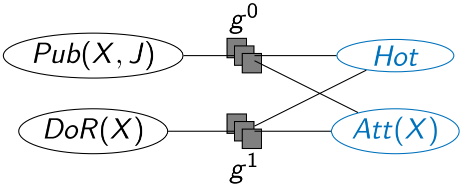

QA based on submodels
ensured to be independent

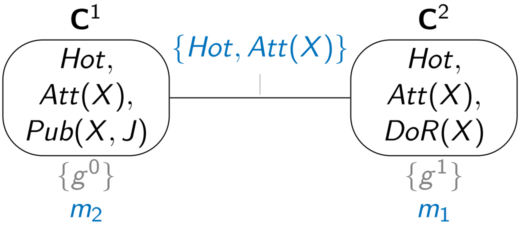

Lifting + Temporal Conditional Independences

Braun and Möller (2016), Murphy (2002), Gehrke et al. (2018)

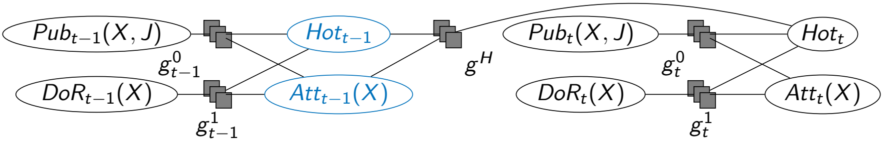

QA based on submodels and time slices
ensured to be independent

Lifting + Temporal Conditional Independences

Braun and Möller (2016), Murphy (2002), Gehrke et al. (2018)

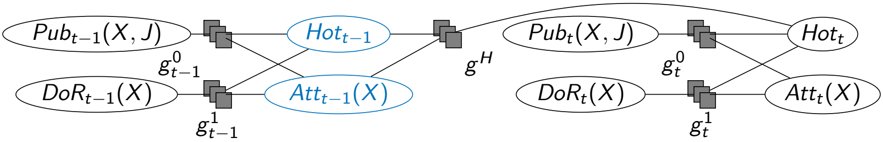

QA based on submodels and time slices ensured to be independent

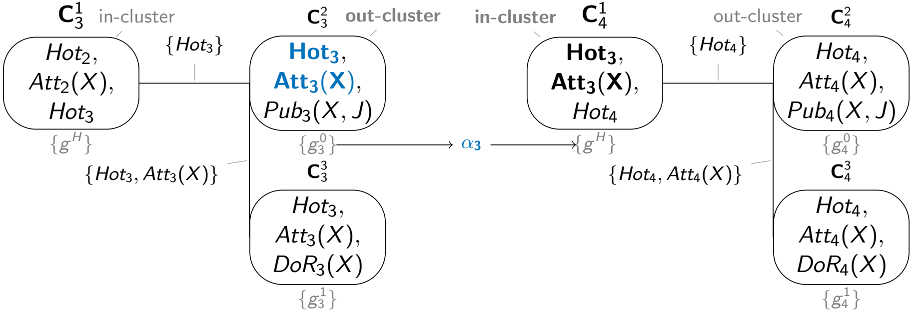

Lifting + Temporal Conditional Independences and Beyond

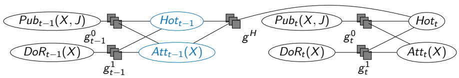

QA based on submodels
and time slices
ensured to be independent

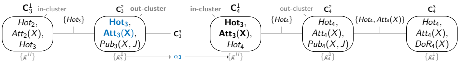

Lifted Dynamic Junction Tree Algorithm (LDJT)

Gehrke et al. (2018)

Answer multiple temporal queries efficiently

Filtering: $P(DoR_5(eve) | Hot_5 = 1)$

Lifting + Temporal Conditional Independences and Beyond

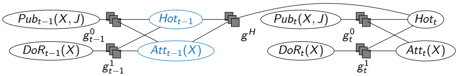

QA based on submodels
and time slices
ensured to be independent

Lifted Dynamic Junction Tree Algorithm (LDJT)

Complexity & Completeness

Polynomial w.r.t. domain size
Linear w.r.t. # time steps
Classes of liftable temporal models

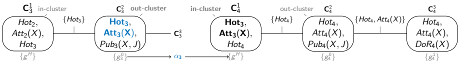

Lifting + Temporal Conditional Independences and Beyond

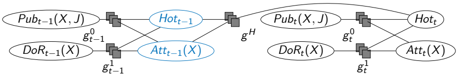

QA based on submodels
and time slices
ensured to be independent

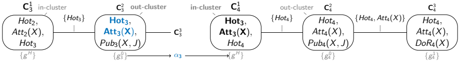

Lifted Dynamic Junction Tree Algorithm (LDJT)

Complexity & Completeness

Decision making

Gehrke et al. (2019b,a)
LJT and LDJT to solve the Maximum Expected Utility problem

Lifting + Temporal Conditional Independences and Beyond

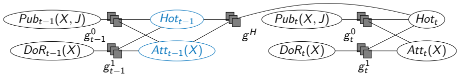

QA based on submodels
and time slices
ensured to be independent

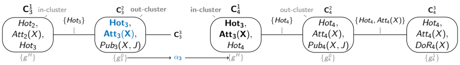

Lifted Dynamic Junction Tree Algorithm (LDJT)

Complexity & Completeness

Decision making

Taming temporal reasoning

Gehrke et al. (2020)

Approximate symmetries over time to retain tractability

Lifted Inference Continued...

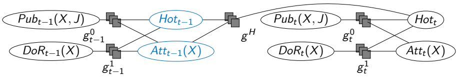

QA based on submodels
and time slices
ensured to be independent

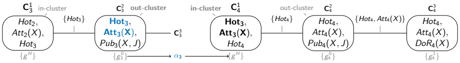

Who did it?

Identifying the most likely source to exhibit an event (combinatorial problem)

Lifted Inference Continued...

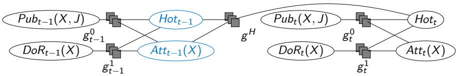

QA based on submodels
and time slices
ensured to be independent

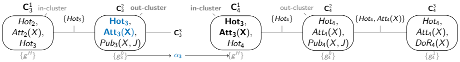

Who did it?

Causality

Started to have a look at lifting causality with Malte

Lifted Inference Continued...

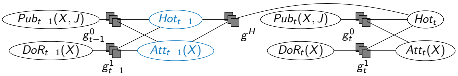

QA based on submodels
and time slices
ensured to be independent

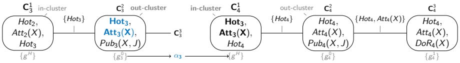

Who did it?

Causality

Preserving Privacy

Its complicated

Lifted Inference Continued...

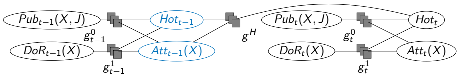

QA based on submodels
and time slices
ensured to be independent

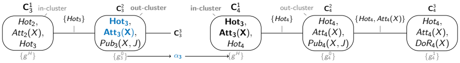

Who did it?

Causality

Preserving Privacy

Text Understanding

Lightweight text understanding using PGMs with Magnus.

References I

Babak Ahmadi, Kristian Kersting, Martin Mladenov, and Sriraam Natarajan. Exploiting Symmetries for Scaling Loopy Belief Propagation and Relational Training. *Machine learning*, 92(1):91–132, 2013.

Tanya Braun and Ralf Möller. Lifted Junction Tree Algorithm. In *Proceedings of KI 2016: Advances in Artificial Intelligence*, pages 30–42. Springer, 2016.

Rodrigo de Salvo Braz, Eyal Amir, and Dan Roth. Lifted First-order Probabilistic Inference. In *IJCAI05 Proceedings of the 19th International Joint Conference on Artificial intelligence*, pages 1319–1325. Morgan Kaufmann Publishers Inc., 2005.

Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Dynamic Junction Tree Algorithm. In *Proceedings of the 23rd International Conference on Conceptual Structures*, pages 55–69. Springer, 2018.

References II

Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Temporal Maximum Expected Utility. In *Proceedings of the 32nd Canadian Conference on Artificial Intelligence, Canadian AI 2019*, pages 380–386. Springer, 2019a.

Marcel Gehrke, Tanya Braun, Ralf Möller, Alexander Waschkau, Christoph Strumann, and Jost Steinhäuser. Lifted Maximum Expected Utility. In *Proceedings of Artificial Intelligence in Health*, pages 131–141. Springer International Publishing, 2019b.

Marcel Gehrke, Ralf Möller, and Tanya Braun. Taming Reasoning in Temporal Probabilistic Relational Models. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, pages 2592–2599. IOS Press, 2020.

References III

Steffen L. Lauritzen and David J Spiegelhalter. Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 50(2):157–224, 1988.

Brian Milch, Luke S. Zettlemoyer, Kristian Kersting, Michael Haimes, and Leslie Pack Kaelbling. Lifted Probabilistic Inference with Counting Formulas. In *AAAI08 Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 2*, pages 1062–1068. AAAI Press, 2008.

Kevin Patrick Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, University of California, Berkeley, 2002.

David Poole. First-order probabilistic inference. In *IJCAI03 Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 985–991. Morgan Kaufmann Publishers Inc., 2003.

References IV

Stuart J Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 2020.

Nima Taghipour, Daan Fierens, Jesse Davis, and Hendrik Blockeel. Lifted Variable Elimination: Decoupling the Operators from the Constraint Language. *Journal of Artificial Intelligence Research*, 47(1):393–439, 2013.