Taming Exact Inference in Temporal Probabilistic Relational Models

Colloquium

Marcel Gehrke

Institute of Information Systems
University of Lübeck

December 17, 2021

Artificial Intelligence: An Agent Perspective

Russell and Norvig (1995)

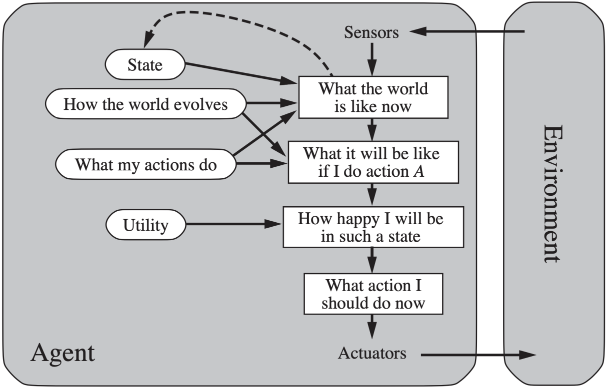

Artificial Intelligence: An Agent Perspective

Russell and Norvig (1995)

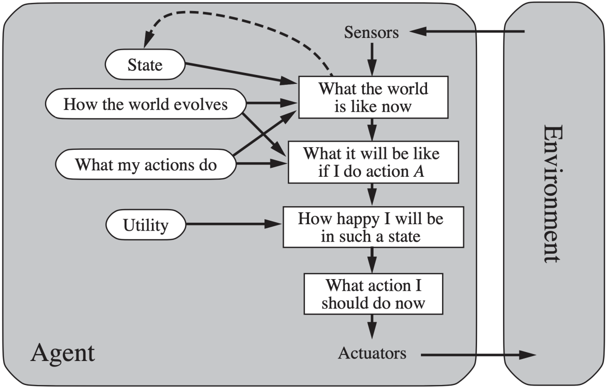

Knowledge representation and reasoning under uncertainty

$\rightarrow$ 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. $P_G = \frac{1}{Z}\prod_{f \in gr(u(G))} f$

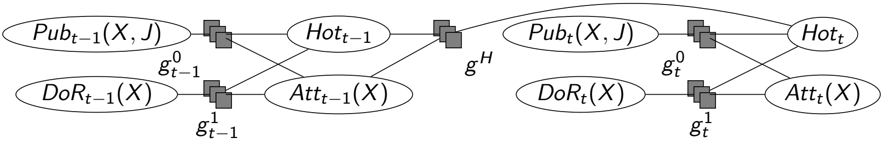

Probabilistic Temporal Relational and Lifted Models

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

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

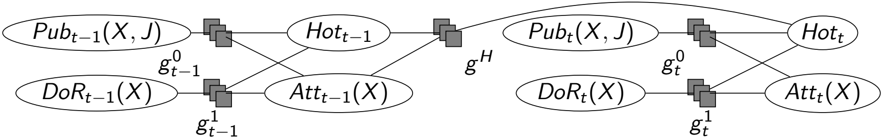

Marginal distribution query: $P(A^i_\pi | E_{0:t})$ w.r.t. the model:

- 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. $P_G = \frac{1}{Z}\prod_{f \in gr(u(G))} f$

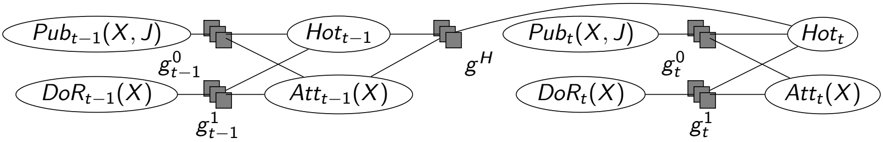

QA: Eliminate all non-query variables
while **avoiding grounding and unrolling** G **as well as building** P_G

QA: Lifted Variable Elimination (LVE)

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

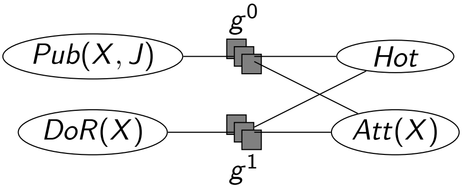

$P(DoR(eve))$

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

QA: Lifted Variable Elimination (LVE)

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

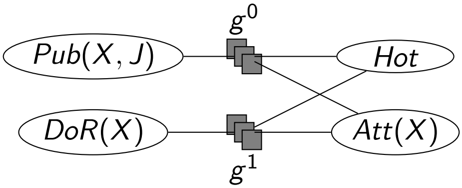

$$P(\textcolor{red}{DoR(eve)}) \propto \sum_{Hot} \left(\sum_{\substack{DoR(X) \\ X \neq eve}} \sum_{Att(X)} g^1 \left(\sum_{Pub(X,J)} g^0 \right)^{|J|} \right)^{|X|_{X \neq eve}}$$

$\sum_V$ 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)

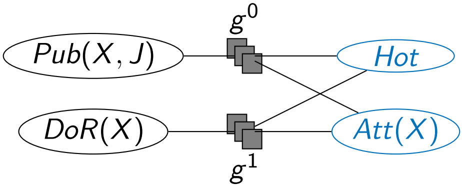

QA based on submodels
ensured to be independent

QA: Lifted Junction Tree Algorithm (LJT)

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

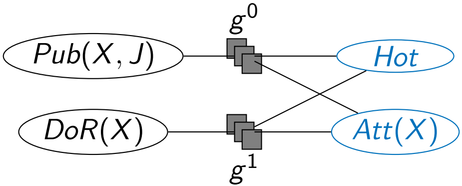

QA based on submodels
ensured to be independent

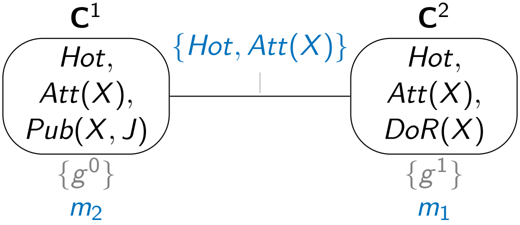

Lifting + Temporal Conditional Independences

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

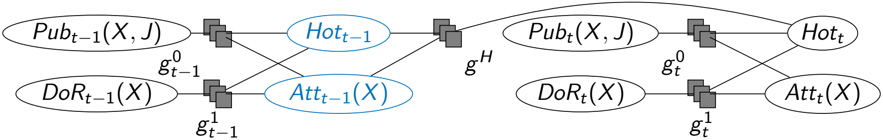

QA based on submodels and time slices ensured to be independent

Lifting + Temporal Conditional Independences

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

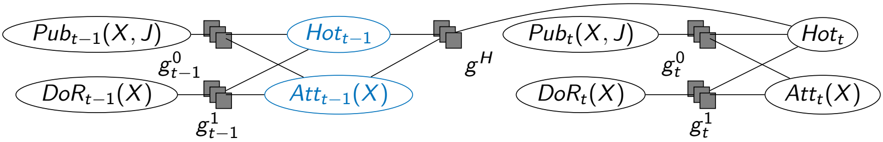

QA based on submodels and time slices
ensured to be independent

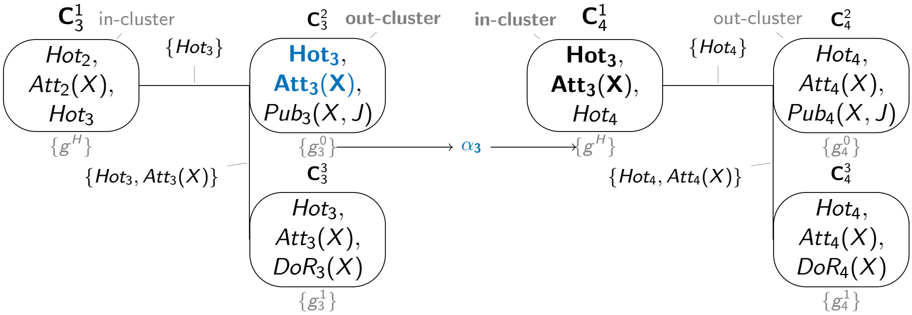

Lifting + Temporal Conditional Independences and Beyond

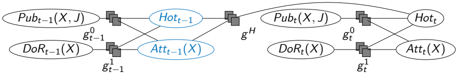

QA based on submodels
and time slices
ensured to be independent

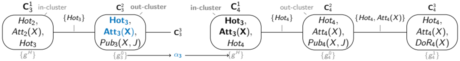

Lifted Dynamic Junction Tree Algorithm (LDJT)

G. et al. (2018d)
Answer multiple temporal queries efficiently
Filtering: $P(DoR_5(eve) | Hot_5 = 1)$

Lifting + Temporal Conditional Independences and Beyond

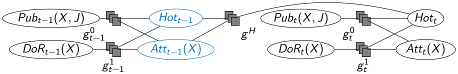

QA based on submodels
and time slices
ensured to be independent

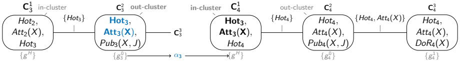

Lifted Dynamic Junction Tree Algorithm (LDJT)

Liftability

G. et al. (2018a,c)

Avoid temporal message-induced groundings

Lifting + Temporal Conditional Independences and Beyond

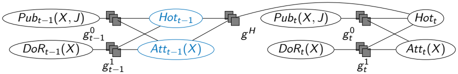

QA based on submodels
and time slices
ensured to be independent

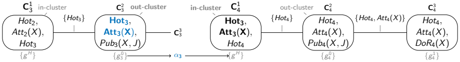

Lifted Dynamic Junction Tree Algorithm (LDJT)

Liftability

Marginal queries

G. et al. (2018b, 2019a)

Hindsight: $P(DoR_1(eve)|Hot_5 = 1)$

Prediction: $P(DoR_5(eve)|Hot_1 = 1)$

Conjunctive: $P(DoR_1(eve), Hot_5)$

Lifting + Temporal Conditional Independences and Beyond

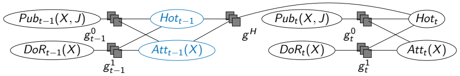

QA based on submodels
and time slices
ensured to be independent

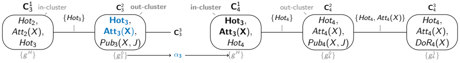

Lifted Dynamic Junction Tree Algorithm (LDJT)

Liftability

Marginal queries

Assignments queries

G. et al. (2019b)
LDJT versions using arg max

Lifting + Temporal Conditional Independences and Beyond

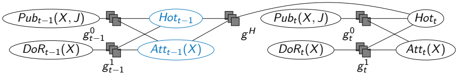

QA based on submodels
and time slices
ensured to be independent

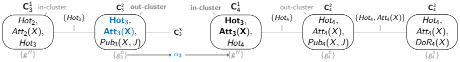

Lifted Dynamic Junction Tree Algorithm (LDJT)

Liftability

Marginal queries

Assignments queries

Complexity & Completeness

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

Lifted Inference Continued...

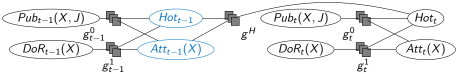

QA based on submodels
and time slices
ensured to be independent

Empirical studies

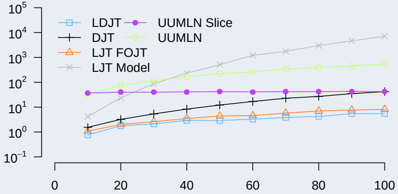

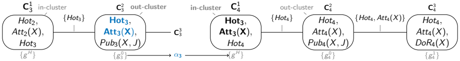

Lifted Inference Continued...

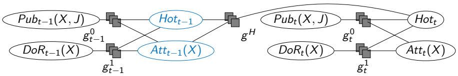

QA based on submodels
and time slices
ensured to be independent

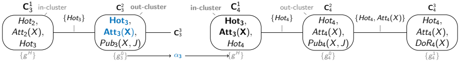

Empirical studies

Uncertain evidence

G. et al. (2019d)

In LVE and LJT

Lifted Inference Continued...

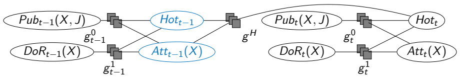

QA based on submodels
and time slices
ensured to be independent

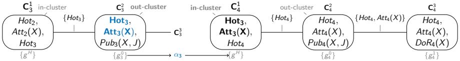

Empirical studies

Uncertain evidence

Decision making

G. et al. (2019e,c)
LJT and LDJT to solve the Maximum Expected Utility problem

Lifted Inference Continued...

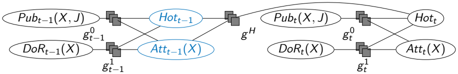

QA based on submodels
and time slices
ensured to be independent

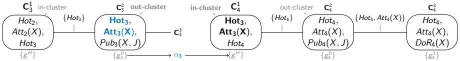

Empirical studies

Uncertain evidence

Decision making

Taming temporal reasoning

G. et al. (2020)

Approximate symmetries over time to retain tractability

Lifted Inference Continued...

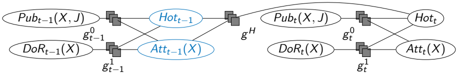

QA based on submodels
and time slices
ensured to be independent

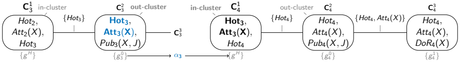

Empirical studies

Uncertain evidence

Decision making

Taming temporal reasoning

Continued by/with colleagues

- lifted decision making cont.
- lifted continuous models
- real life settings
- hybrid models

Conclusion

Taming Temporal Inference

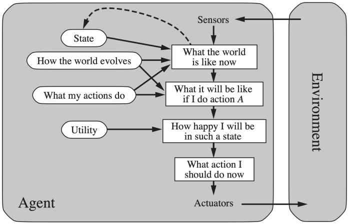

Russell and Norvig (1995)

Knowledge representation and reasoning under uncertainty
→ Statistical Relational AI

Exact Lifted Temporal Inference

- Lifting dynamic junction trees
- Lifting temporal marginal and assignment queries
- Keeping reasoning polynomial

Tractable inference for a variety of queries

Conclusion

Taming Temporal Inference

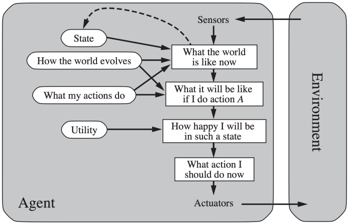

Russell and Norvig (1995)

Knowledge representation and reasoning under uncertainty
→ Statistical Relational AI

Exact Lifted Temporal Inference

- Lifting dynamic junction trees
- Lifting temporal marginal and assignment queries
- Keeping reasoning polynomial

Tractable inference for a variety of queries

Future Work

- Human aware XAI
- Universe/Domain changes

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.

Adnan Darwiche. *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press, 2009.

References II

Marcel Gehrke, Tanya Braun, and Ralf Möller. Towards Preventing Unnecessary Groundings in the Lifted Dynamic Junction Tree Algorithm. In *Proceedings of KI 2018: Advances in Artificial Intelligence*, pages 38–45. Springer, 2018a.

Marcel Gehrke, Tanya Braun, and Ralf Möller. Answering Multiple Conjunctive Queries with the Lifted Dynamic Junction Tree Algorithm. In *Proceedings of the AI 2018: Advances in Artificial Intelligence*, pages 543–555. Springer, 2018b.

Marcel Gehrke, Tanya Braun, and Ralf Möller. Preventing Unnecessary Groundings in the Lifted Dynamic Junction Tree Algorithm. In *Proceedings of the AI 2018: Advances in Artificial Intelligence*, pages 556–562. Springer, 2018c.

References III

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, 2018d.

Marcel Gehrke, Tanya Braun, and Ralf Möller. Relational Forward Backward Algorithm for Multiple Queries. In *Proceedings of the 32nd International Florida Artificial Intelligence Research Society Conference (FLAIRS-32)*, pages 464–469. AAAI Press, 2019a.

Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Temporal Most Probable Explanation. In *Proceedings of the 24th International Conference on Conceptual Structures*, pages 72–85. Springer, 2019b.

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, 2019c.

References IV

Marcel Gehrke, Tanya Braun, and Ralf Möller. Uncertain Evidence for Probabilistic Relational Models. In *Proceedings of the 32nd Canadian Conference on Artificial Intelligence, Canadian AI 2019*, pages 80–93. Springer, 2019d.

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, 2019e.

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, 2020.

References V

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 VI

Stuart J Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 1995.

Nima Taghipour, Jesse Davis, and Hendrik Blockeel. First-order Decomposition Trees. In *NIPS13 Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, pages 1052–1060. Curran Associates Inc., 2013a.

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, 2013b.