WWU
MÜNSTER

A Glimpse into Statistical Relational AI

The Power of Indistinguishability

Tanya Braun, University of Münster

Agenda

- Statistical Relational Artificial Intelligence
  - Probabilistic relational models
  - Grounding semantics
  - Context
- The Power of Indistinguishability
  - Lifted query answering and tractability
  - Keeping indistinguishability over time
  - Indistinguishability in decision making
- Summary

Statistical Relational Artificial Intelligence (StaRAI)

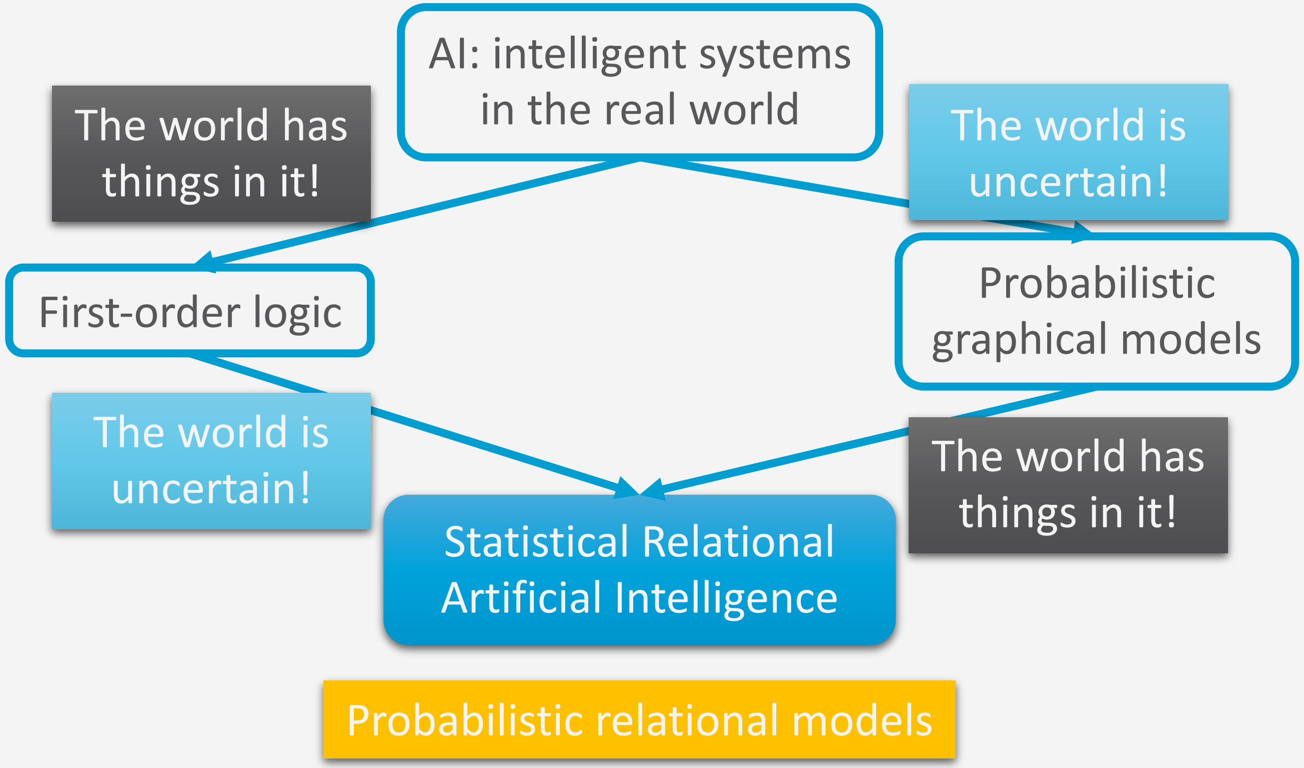

Figure based on Stuart Russell

Application: Epidemics

- Atoms: Parameterised random variables = PRVs
  - With logical variables
    - E.g., X, M
    - Possible values (domain):
      $$dom(X) = \{alice, eve, bob\}$$
      $$dom(M) = \{injection, tablet\}$$
  - With range
    - E.g., Boolean
    - $ran(Travel(X)) = \{true, false\}$
  - Represent sets of *indistinguishable* random variables

$$Nat(D) = natural\ disaster\ D$$
$$Acc(A) = accident\ A$$

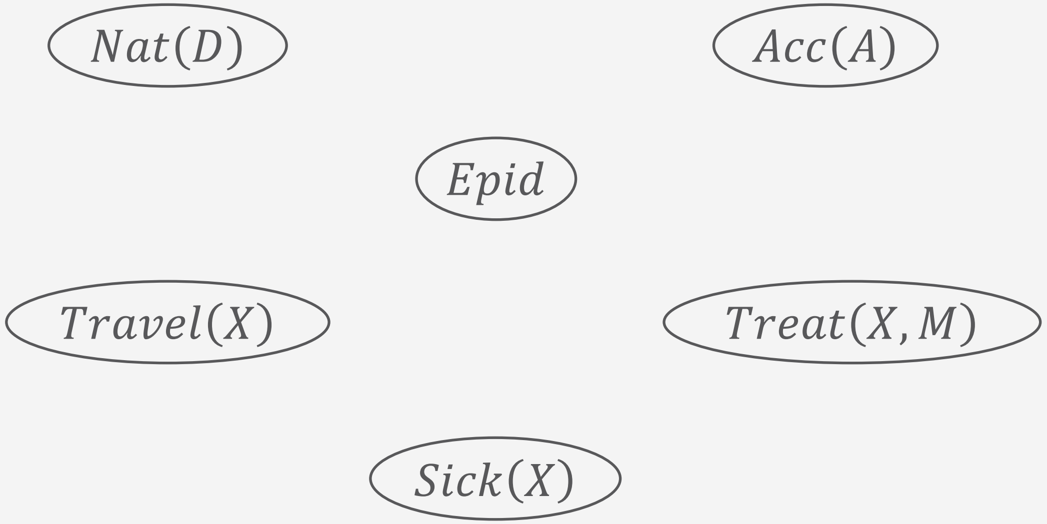

Encoding the Joint Distribution: Factorisation

- Factors with PRVs = parfactors
  - E.g., g_2

$Travel(X)$	$Epid$	$Sick(X)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$false$	$true$	6
$true$	$true$	$false$	2
$true$	$true$	$true$	9

Potentials
- In parfactors, just like in factors, no probability distribution as factors required

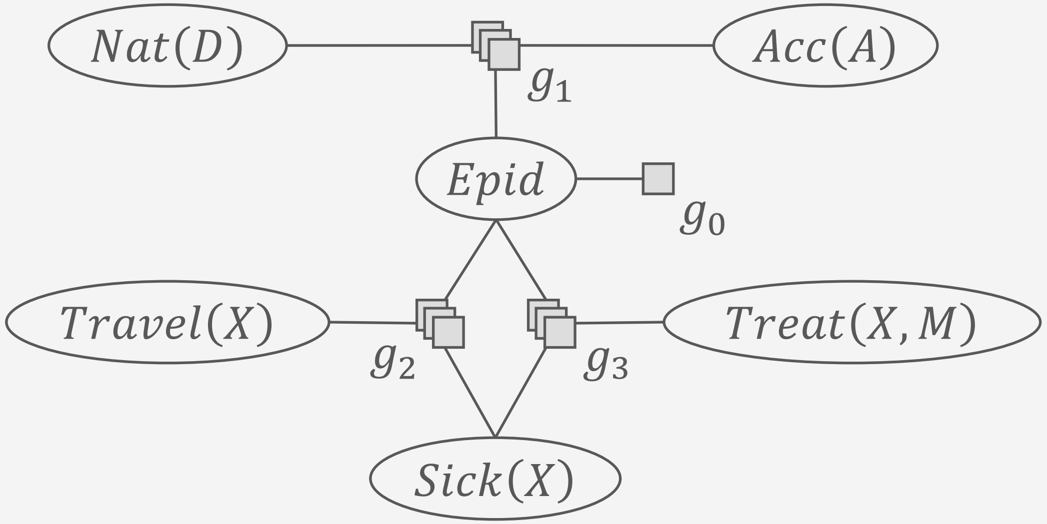

Factors

- Grounding
  - E.g., $gr(g_2) = \{f_2^1, f_2^2, f_2^3\}$

$Travel(X)$	$Epid$	$Sick(X)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$false$	$true$	6
$true$	$true$	$false$	2
$true$	$true$	$true$	9

$Travel(eve)$	$Epid$	$Sick(eve)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$fals$		
$true$	tru		
$true$	tru		

$Travel(bob)$	$Epid$	$Sick(bob)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
		$true$	6
		$false$	4
		$true$	6
		$false$	2
		$true$	9

$Travel(alice)$	$Epid$	$Sick(alice)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$false$	$true$	6
$true$	$true$	$false$	2
$true$	$true$	$true$	9

$reat(X,M)$

Encoding the Joint Distribution

- Set of parfactors = model
  - E.g., $G = \{g_1, g_2, g_3\}$
  - Semantics: Joint probability distribution P_G
    - Build by grounding, multiplying all grounded factors, and normalising the result
    - Grounding semantics [Sato 95, Fuhr 95]

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

$$Z = \sum_{v \in r(rv(gr(G)))} \prod_{f \in gr(G)} f_i(\pi_{rv(f_i)}(v))$$

$\pi_{variables}(v)$ = projection of v onto $variables$

Sparse encoding of joint distribution

$3 \cdot 2^3 = 24$ entries in 3 parfactors, 6 PRVs

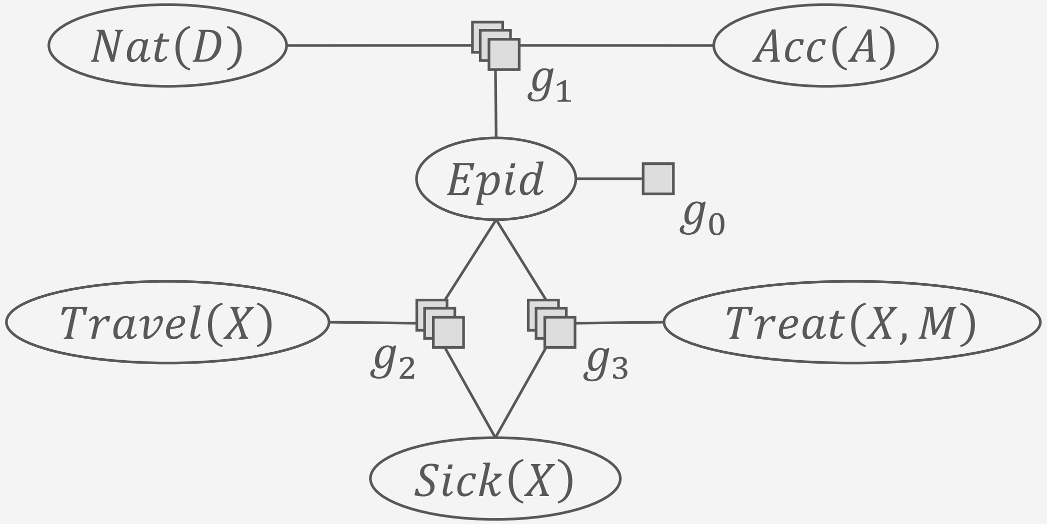

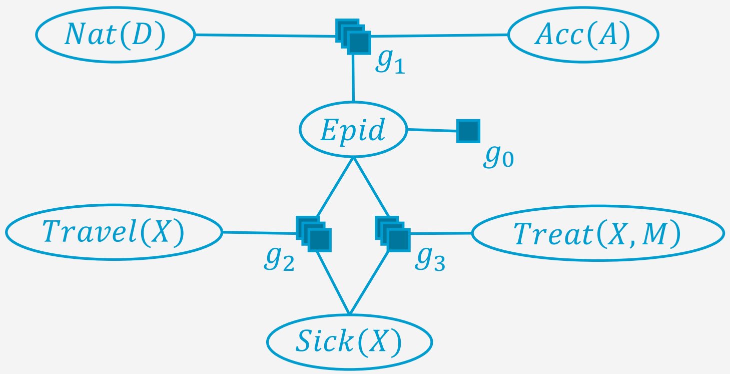

Grounded Model

- Given domains
  - $dom(X) = \{alice, eve, bob\}$
  - $dom(M) = \{m_1, m_2\}$
  - $dom(D) = \{flood, fire\}$
  - $dom(W) = \{virus, war\}$
- Indistinguishability in
  - Graph structure
  - Factors

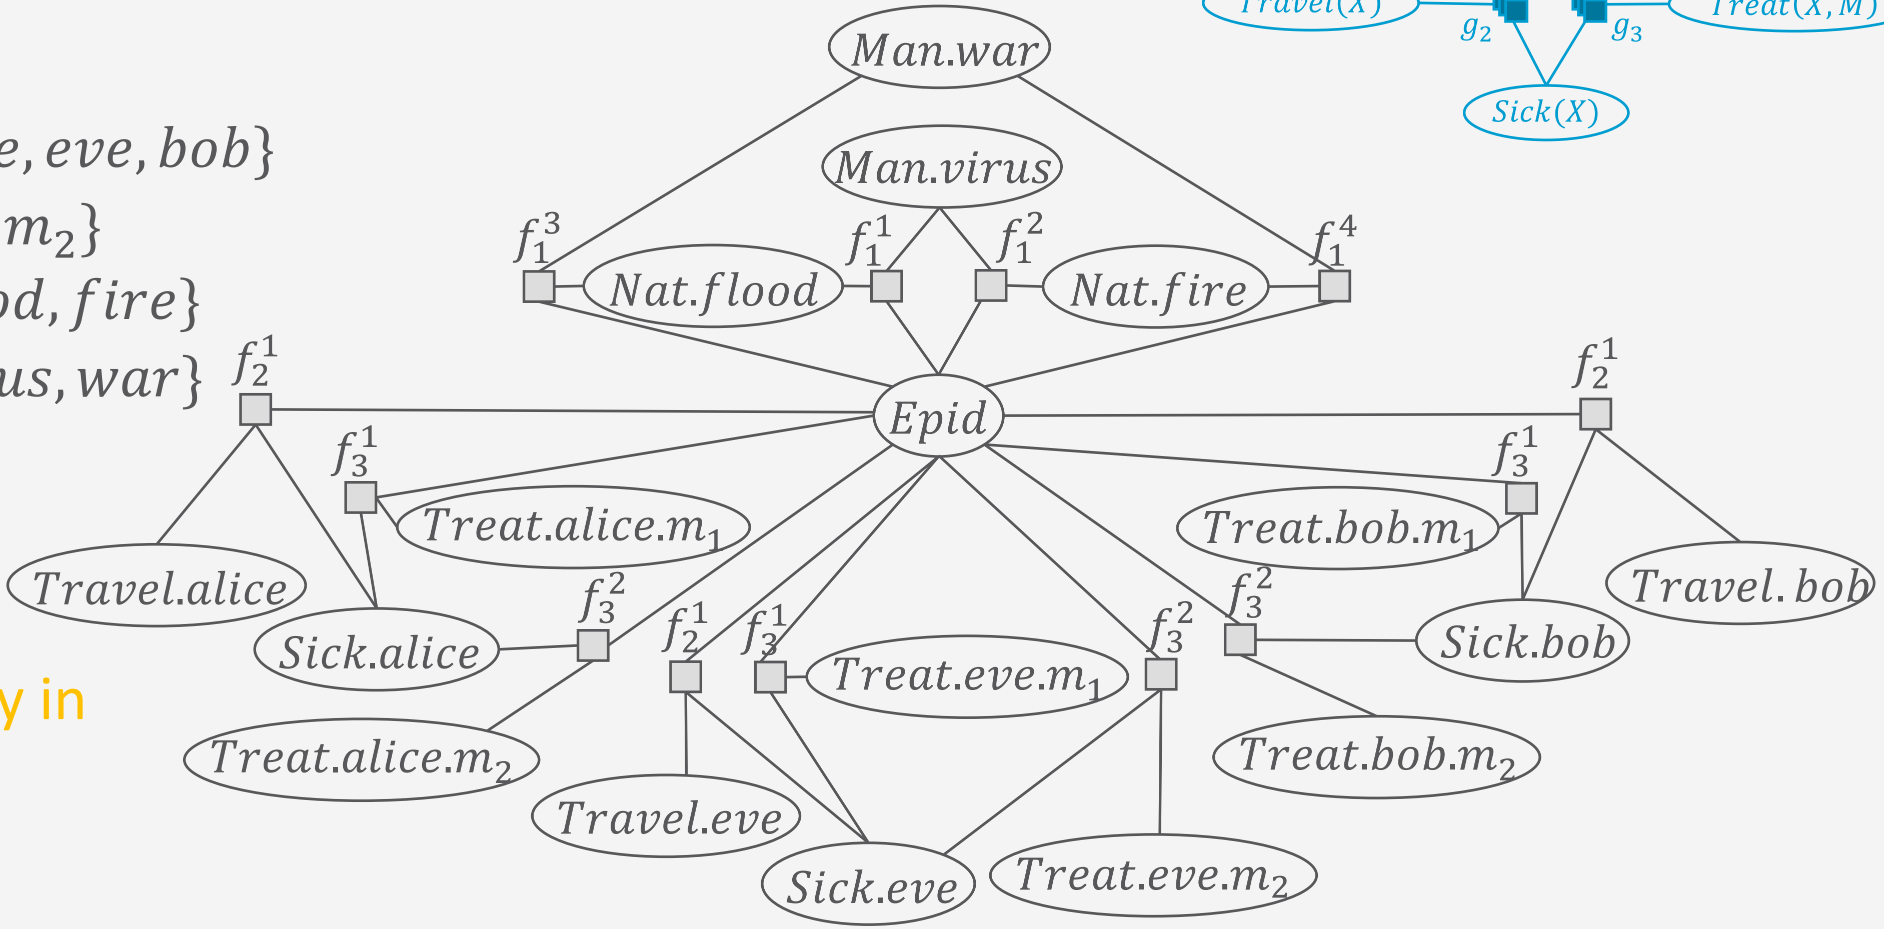

Probabilistic Relational Models and Variants

- Parfactors Models
  [Poole 03, Taghipour et al. 13, B & Möller 16-19, Gehrke, B & Möller 18-19]

- Markov Logic Networks (MLNs) [Richardson & Domingos 06]
  - Use logical formulas to specify potential functions

- Probabilistic Soft Logic (PSL) [Bach et al. 17]
  - Use density functions to specify potential functions

- Based on grounding semantics [Sato 95, Fuhr 95]

The Larger Scope

Statistical Relational Learning & AI

- Study and design
  - intelligent agents
  - that reason about and
  - act in noisy worlds
  - composed of objects and relations among the objects

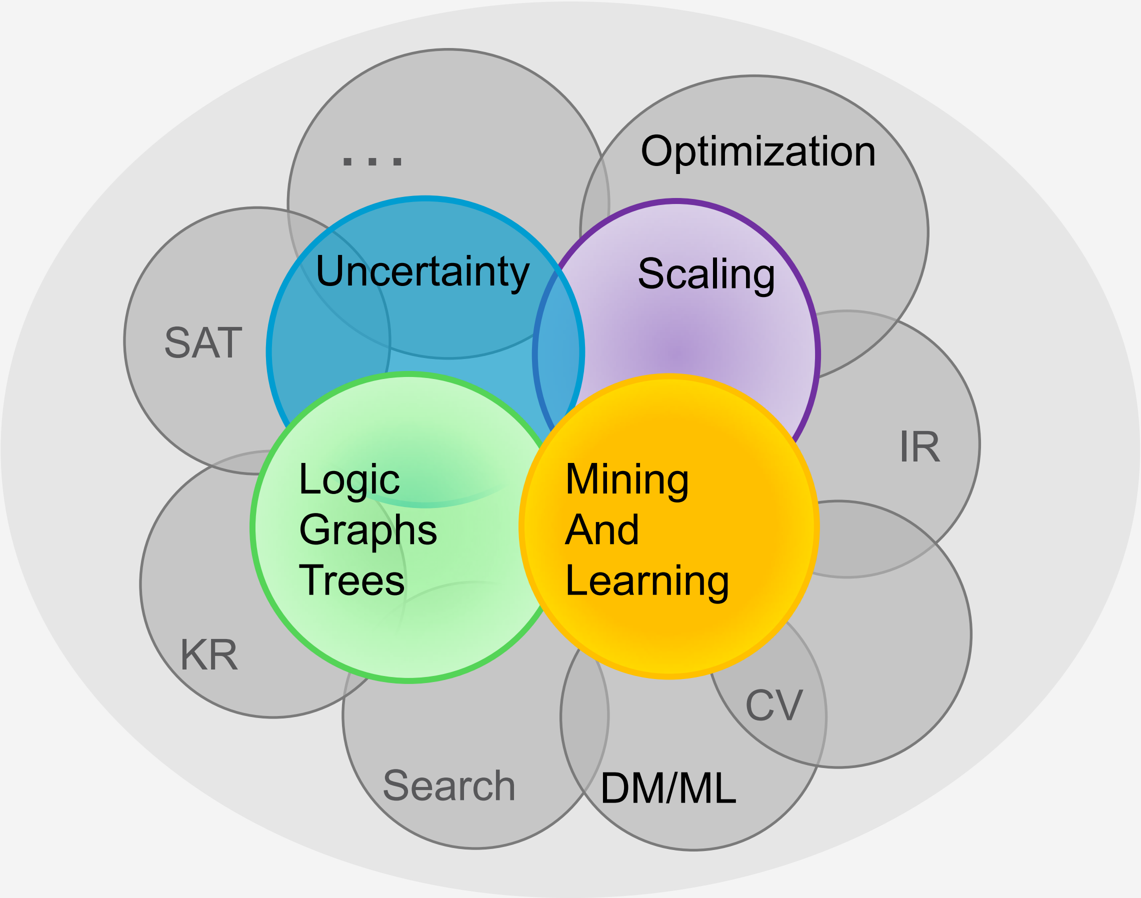

Agenda

- Statistical Relational Artificial Intelligence
  - Probabilistic relational models
  - Grounding semantics
  - Context
- The Power of Indistinguishability
  - Lifted query answering and tractability
  - Keeping indistinguishability over time
  - Indistinguishability in decision making
- Summary

Lifted Query Answering and Tractability

The Power of Indistinguishability

Reasoning on Probabilistic Relational Models

- Inference task: query answering (QA)
- Queries:
  - Marginal distribution
    - $P(Sick(eve))$
    - $P(Travel(eve,)\ Treat(eve, m_1))$
  - Conditional distribution
    - $P(Sick(eve) | Epid)$
    - $P(Epid | Sick(eve) = true)$
  - Assignment queries: $\arg\max_{\boldsymbol{a} \in ran(\boldsymbol{A})} P(\boldsymbol{a} | \boldsymbol{e})$
    - MPE: $\boldsymbol{A} = rv(\boldsymbol{G}) \setminus rv(\boldsymbol{e})$
    - MAP: $\boldsymbol{A} \subseteq rv(\boldsymbol{G}) \setminus rv(\boldsymbol{e})$
      - What is not in $\boldsymbol{A}$ needs to be summed out

Goal: Avoid groundings!
→ *lifted* inference

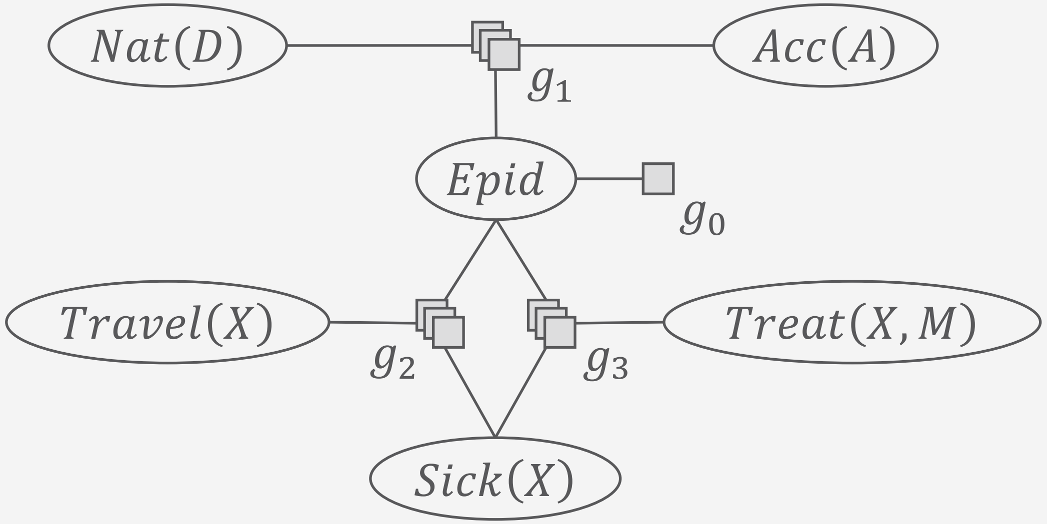

QA: Lifted Variable Elimination (LVE)

- Eliminate all variables not appearing in query
- Lifted summing out
  - Sum out *representative* instance as in propositional variable elimination
  - Exponentiate result for indistinguishable instances
- Correctness: Equivalent ground operation
  - Each instance is summed out
  - Result: factor f that is identical for all instance
  - Multiplying indistinguishable results
    → exponentiation of one representative f

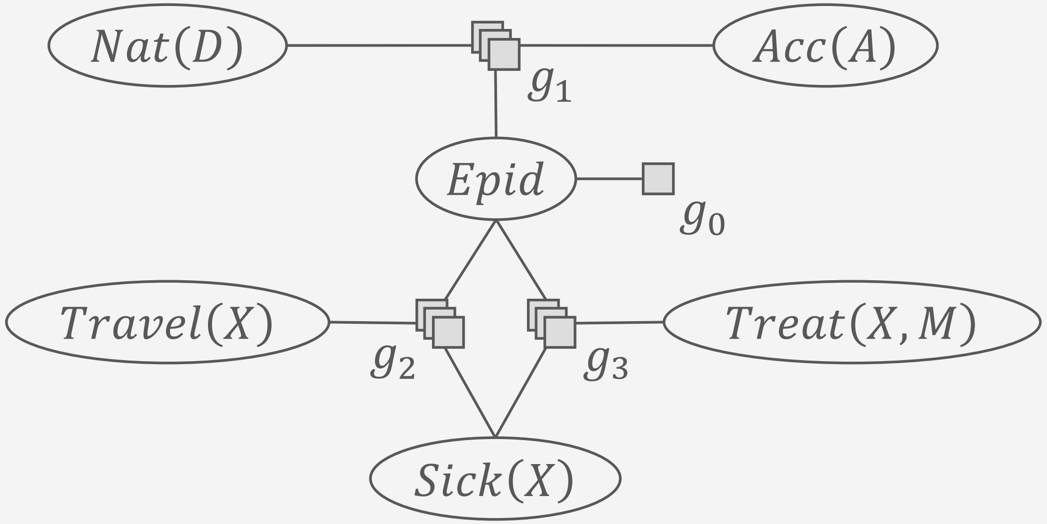

QA: LVE in Detail

- E.g., marginal
  - $P(Travel(eve))$
  - Split atoms $R(\dots, X, \dots)$ w.r.t. eve if eve in $dom(X)$

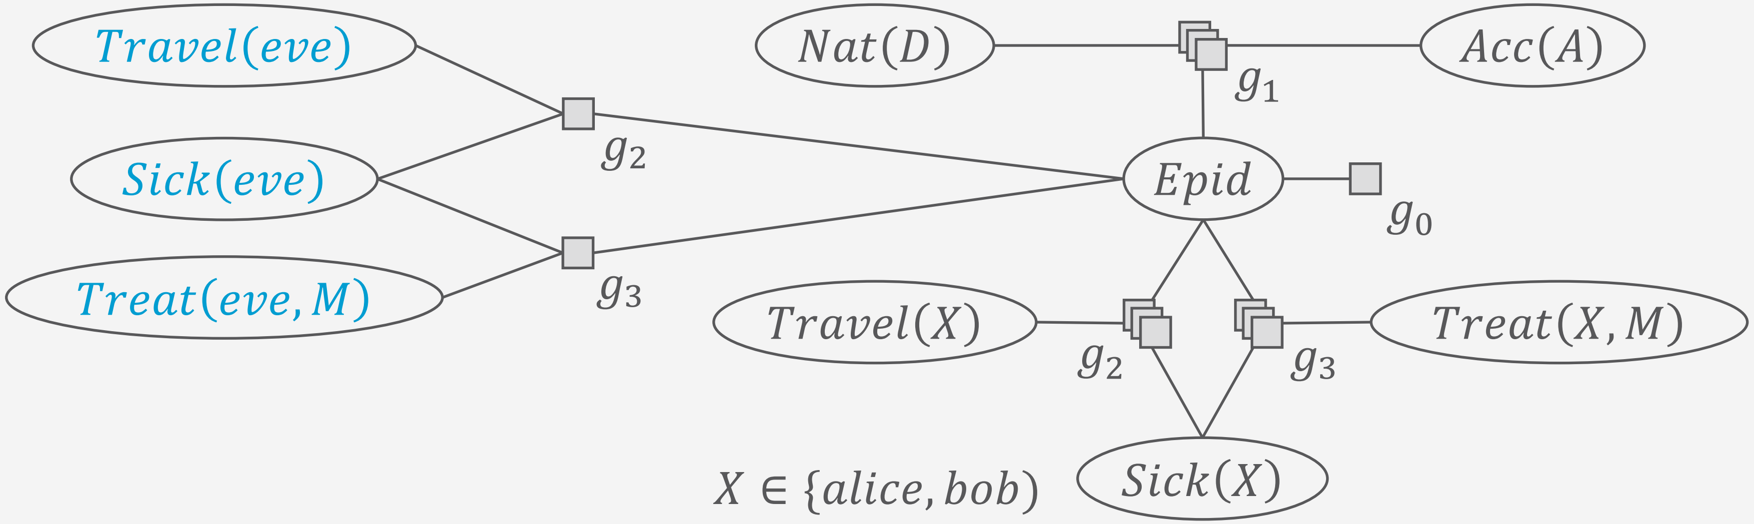

QA: LVE in Detail

- E.g., marginal
  - $P(Travel(eve))$
  - Split atoms $R(\dots, X, \dots)$ w.r.t. eve if eve in $dom(X)$
  - Eliminate all non-query variables

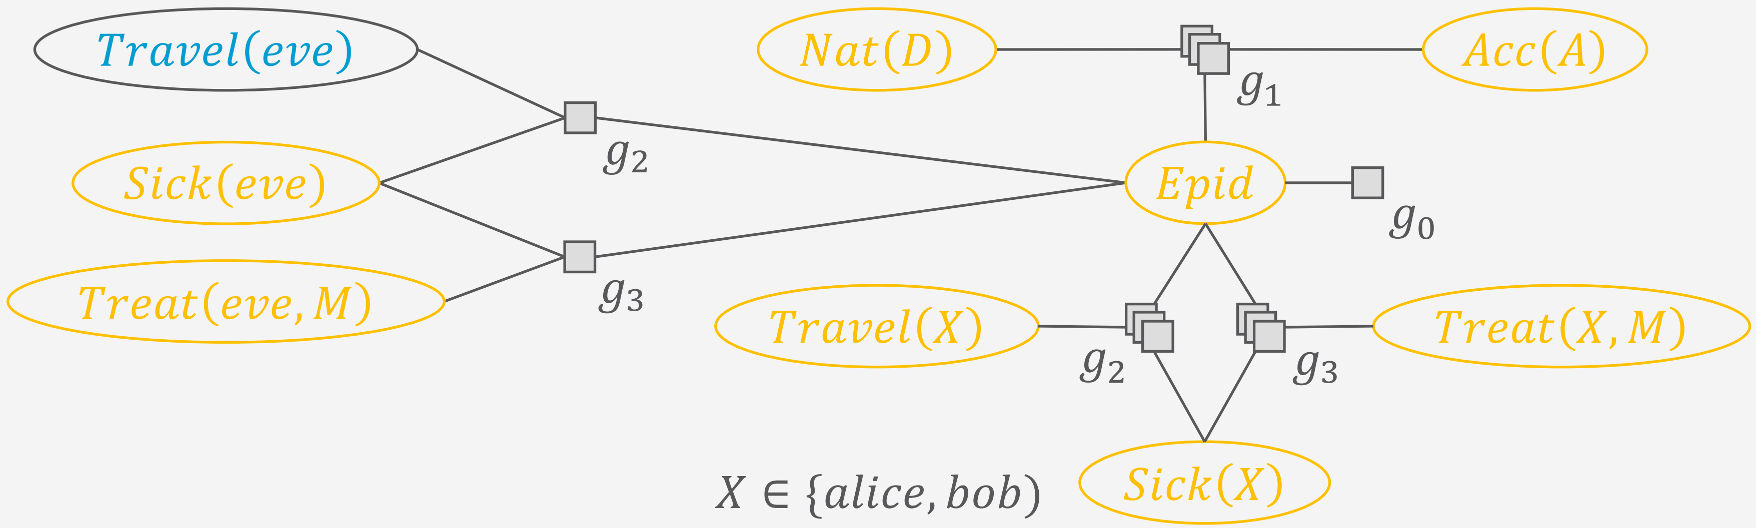

QA: LVE in Detail

- Eliminate $Treat(X, M)$
  - Appears in only one g: g_3
  - Contains all logical variables of g_3: X, M
  - For each X constant: the same number of M constants
  - ✓ Preconditions of lifted summing out fulfilled, lifted summing out possible

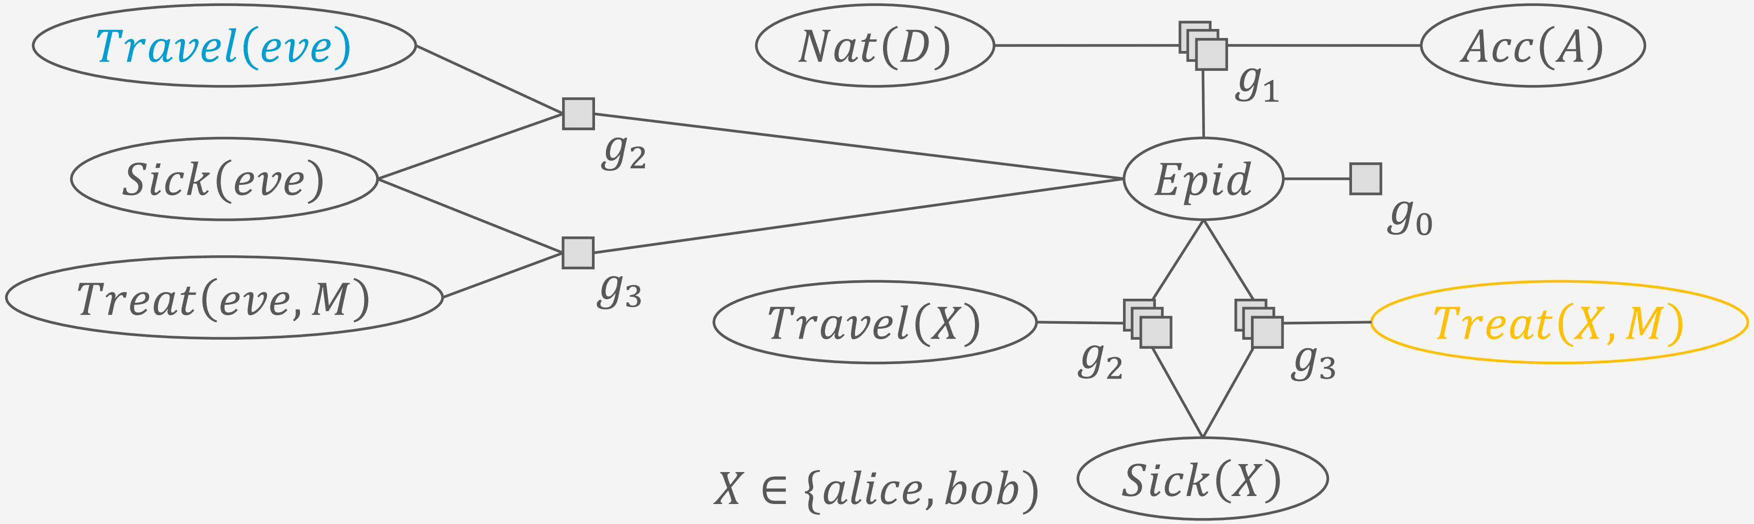

LVE in Detail: Lifted Summing Out

- Eliminate $Treat(X, M)$ by lifted summing out
  1. Sum out representative
  2. Exponentiate for indistinguishable objects

$$\left(\sum_{t \in r(Treat(X,M))} g_3(Epid = e, Sick(X) = s, Treat(X, M) = t)\right)^{\#M|X}$$

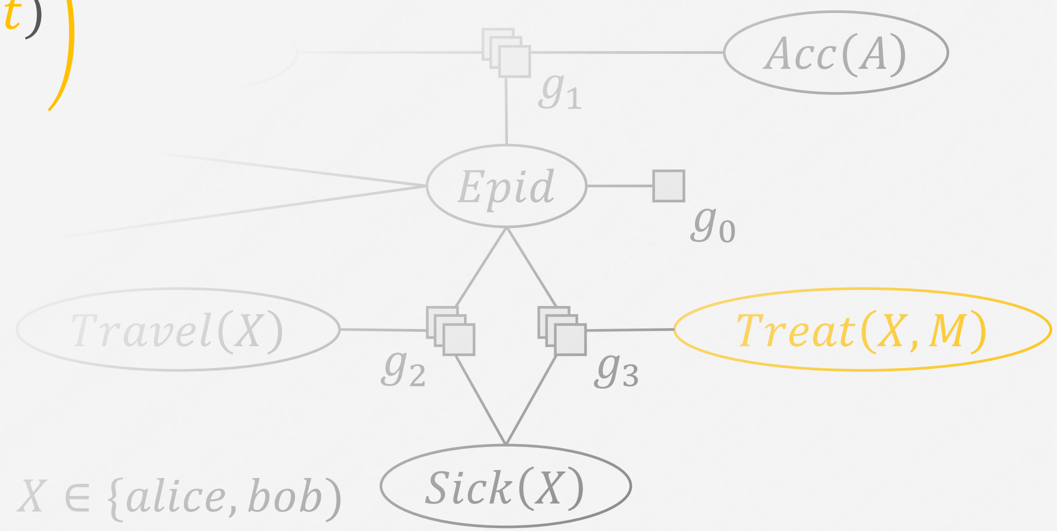

LVE in Detail: Lifted Summing Out

$$\left(\sum_{t \in r(Treat(X,M))} g_3(Epid = e, Sick(X) = s, Treat(X,M) = t)\right)^{\#M|X}$$

$Epid$	$Sick(X)$	$Treat(X,M)$	g_3
$false$	$false$	$false$	9
$false$	$false$	$true$	1
$false$	$true$	$false$	6
$false$	$true$	$true$	3
$true$	$false$	$false$	7
$true$	$false$	$true$	5
$true$	$true$	$false$	4
$true$	$true$	$true$	8

$Epid$	$Sick(X)$	Σ
$false$	$false$	10
$false$	$true$	9
$true$	$false$	12
$true$	$true$	12

$Epid$	$Sick(X)$	^
$false$	$false$	10^2
$false$	$true$	9^2
$true$	$false$	12^2
$true$	$true$	12^2

LVE in Detail: Lifted Summing Out

Only here, domain size comes into play
→ no change in graph / parfactor if domain size changes

- Result after summing out $Treat(X, M)$

$$\left(\sum_{t \in r(Treat(X,M))} g_3(Epid = e, Sick(X) = s, Treat(X,M) = t)\right)^{\#M|X}$$

$Epid$	$Sick(X)$	g_3'
$false$	$false$	100
$false$	$true$	81
$true$	$false$	144
$true$	$true$	144

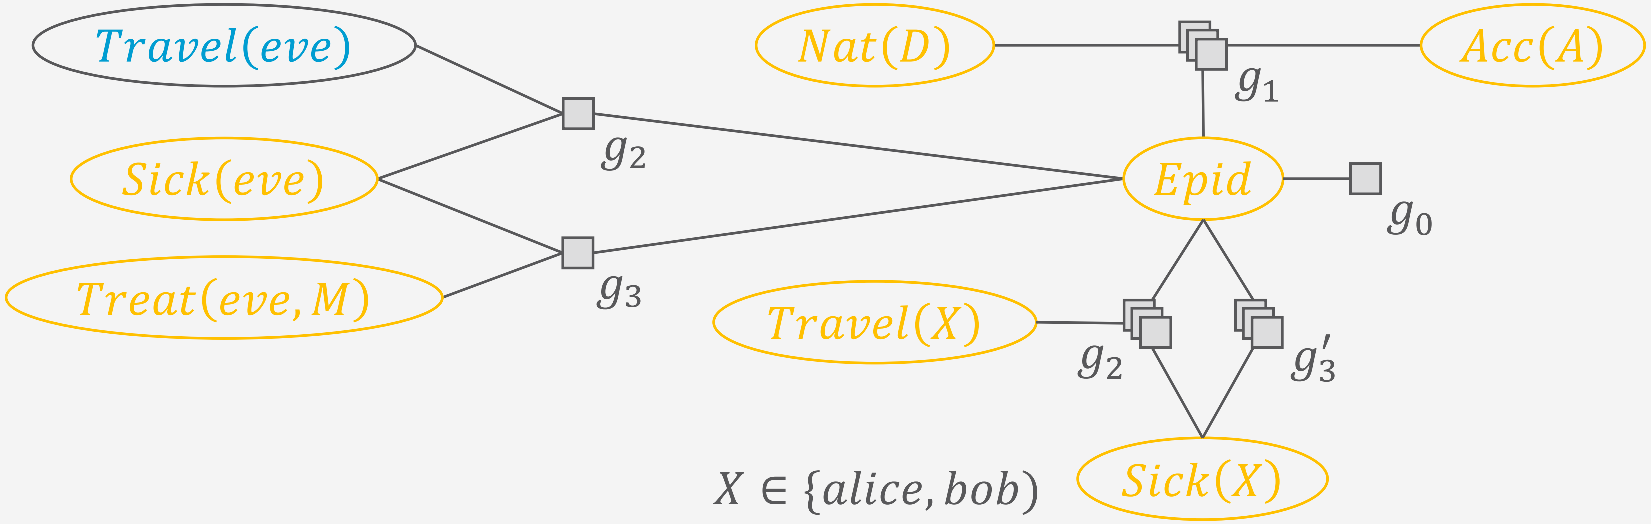

Tractability

- Given a model that allows for lifted calculations
  - I.e., no groundings during solving an instance of the problem
- Solving an instance of the problem is possible in time polynomial in domain sizes

  → The query answering algorithm is domain-lifted
- An query answering problem is tractable
  - when it is solved by an efficient algorithm, running in time polynomial in the number of random variables
- Assume that the number of random variables is characterised by domain sizes
  - Then, solving a query answering problem is tractable under domain-liftability
    - Runtime might still be exponential in other terms
    - More general results by Niepert & Van den Broeck (2014)

Indistinguishable Evidence and Query Terms

Evidence

- Observations for instances of a PRV
  - One of the range values
  - Not available
- Treat as groups per observation
  - Shatter model on the groups
- Example: 10 instances observed true

$Sick(X^T)$	g_e^T
$false$	0
$true$	1

$dom(X^T) = \{x_1, \dots, x_{10}\}$
$dom(X) = \{x_{11}, \dots, x_n\}$

Query Terms

- Indistinguishable instances in query:
  - $P\big(Sick(alice), Sick(eve), Sick(bob)\big)$
  - Result will have local symmetries, e.g., 2 false and 1 true maps to potential of 2
- Parameterised query: $P\big(Sick(X)\big)$
- Use standard LVE
  - Count conversion yields wanted result

$\#_X[Sick(X)]$	g
$[0,3]$	1
$[1,2]$	2
$[2,1]$	3
$[3,0]$	4

Keeping Indistinguishability over Time

The Power of Indistinguishability

Dynamic Probabilistic Relational Models & Temporal Queries

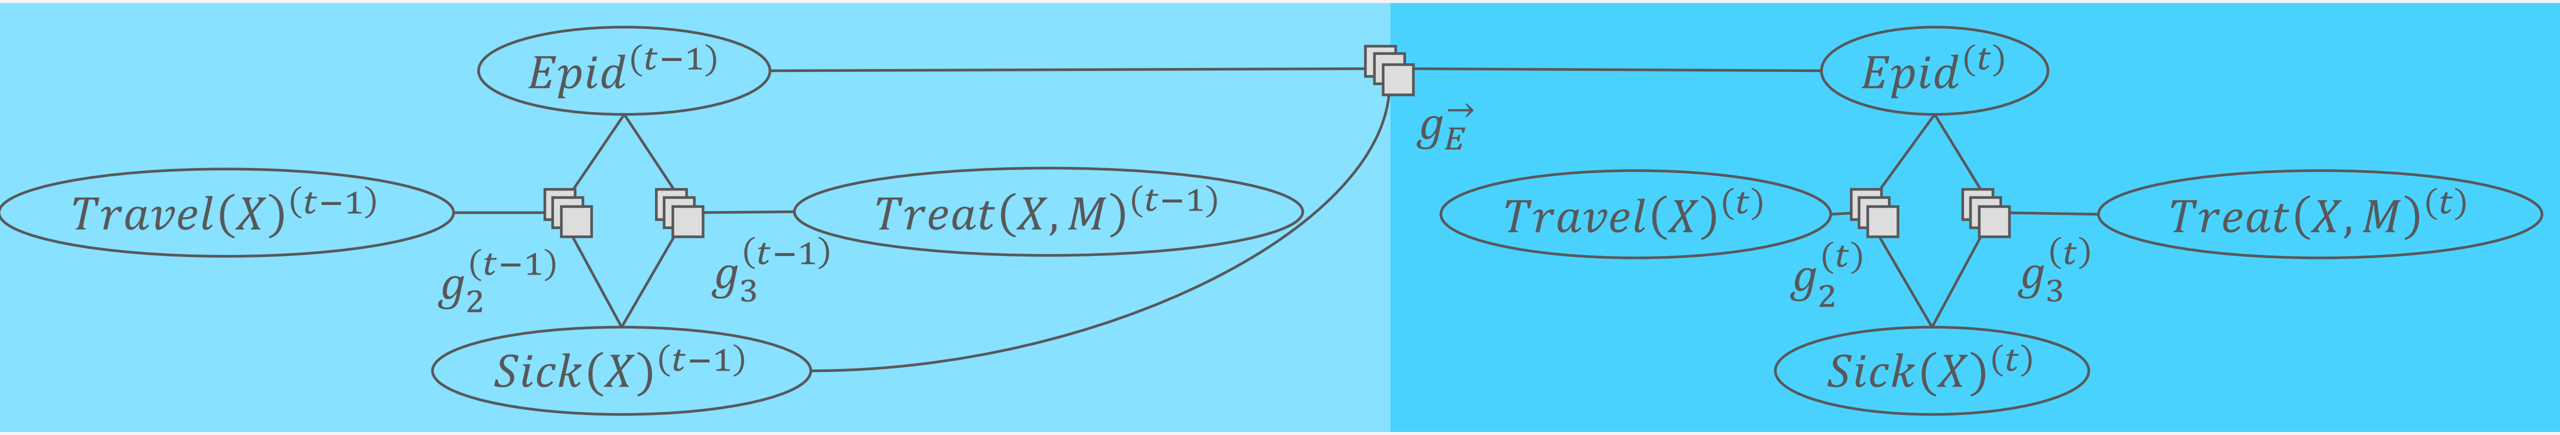

- Marginal distribution queries: $P\left(A_\pi^i \,\middle|\, E_{0:t}\right)$
  - Hindsight: $\pi < t$ (Was there an epidemic $t - \pi$ days ago?)
  - Filtering: $\pi = t$ (Is there currently an epidemic?)
  - Prediction: $\pi > t$ (Will there be an epidemic in $\pi - t$ days?)
- Assignment queries on temporal sequence

Reasoning over Time: Interfaces

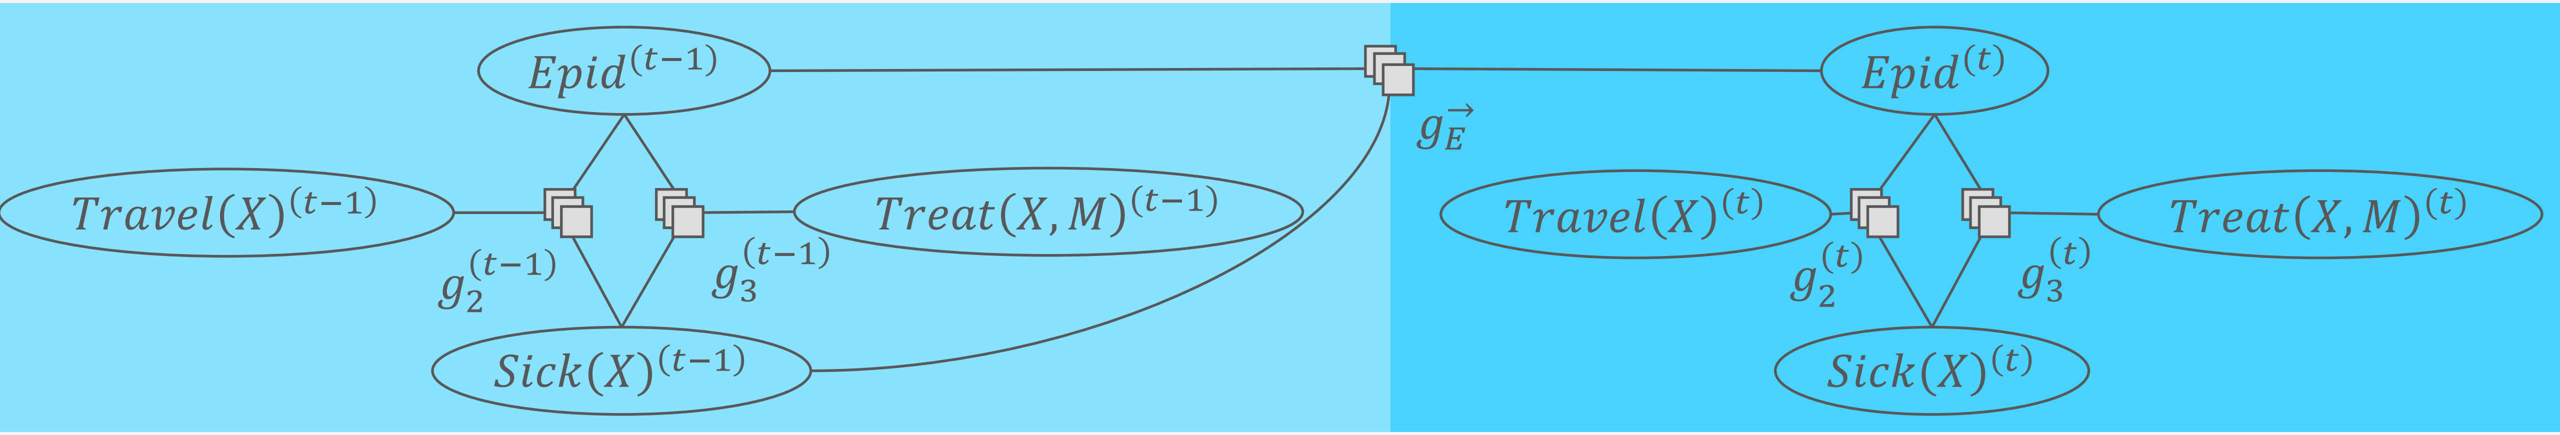

- Main idea: Use temporal conditional independences for efficient temporal QA
  - Normally only a subset of random variables influence next time step → interface variables
  - State description of interface from time slice $t-1$ suffices to perform inference on time slice t
  - → Makes present independent from past / future

Algorithms:
- Propositional: Interface Algorithm [Murphy, 2002]
- Lifted: Lifted Dynamic Junction Tree Algorithm [Gehrke et al, 2018]

Taming Reasoning

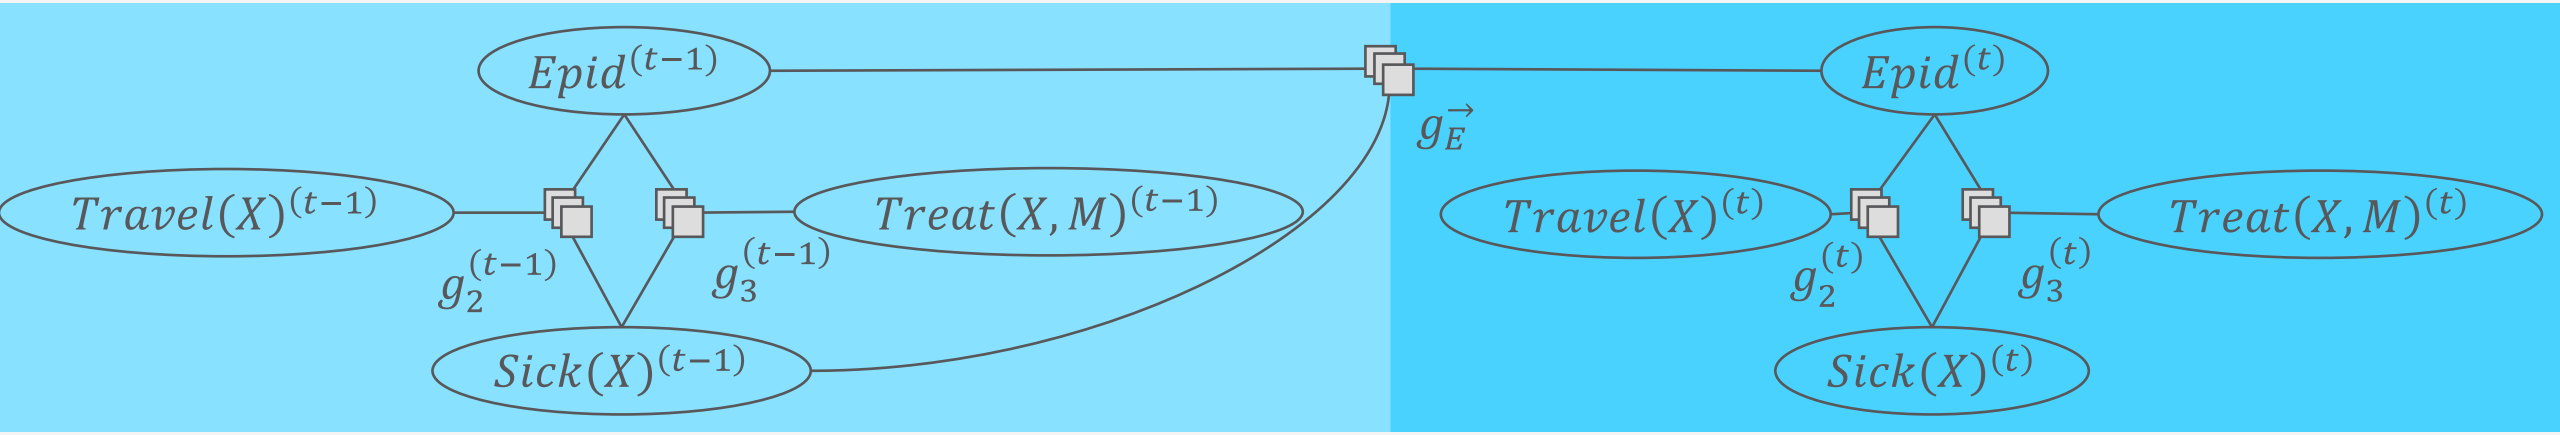

- Evidence can ground a model over time
- Non-symmetric evidence
  - Observe evidence for some instances in one time step
  - Observe evidence for a subset of these instances in another time step
  - Split the logical variable slowly over time

Interface carries over splits, leading to slowly grounding a model over time

Undoing Splits

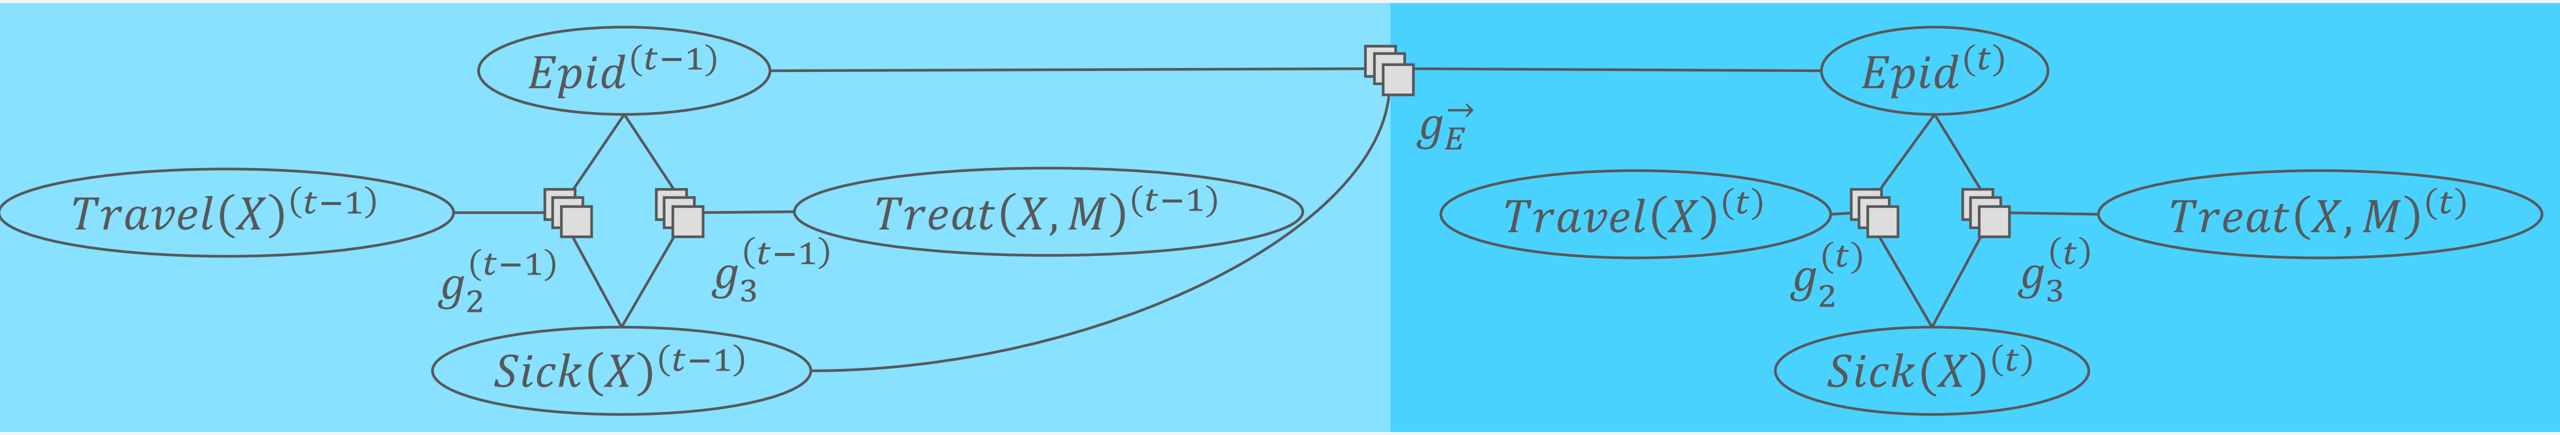

- Need to undo splits to

  keep reasoning polynomial w.r.t. domain sizes
- Where can splits be undone efficiently?
  - When moving from one time step to the next, i.e., in the interface
- How to undo splits?
  - Find approximate symmetries
  - Merge based on groundings
- Is it reasonable to undo splits?
  - Effect of slight differences in evidence?
  - Impact of evidence vs. temporal model

Is It Reasonable to Undo Splits?

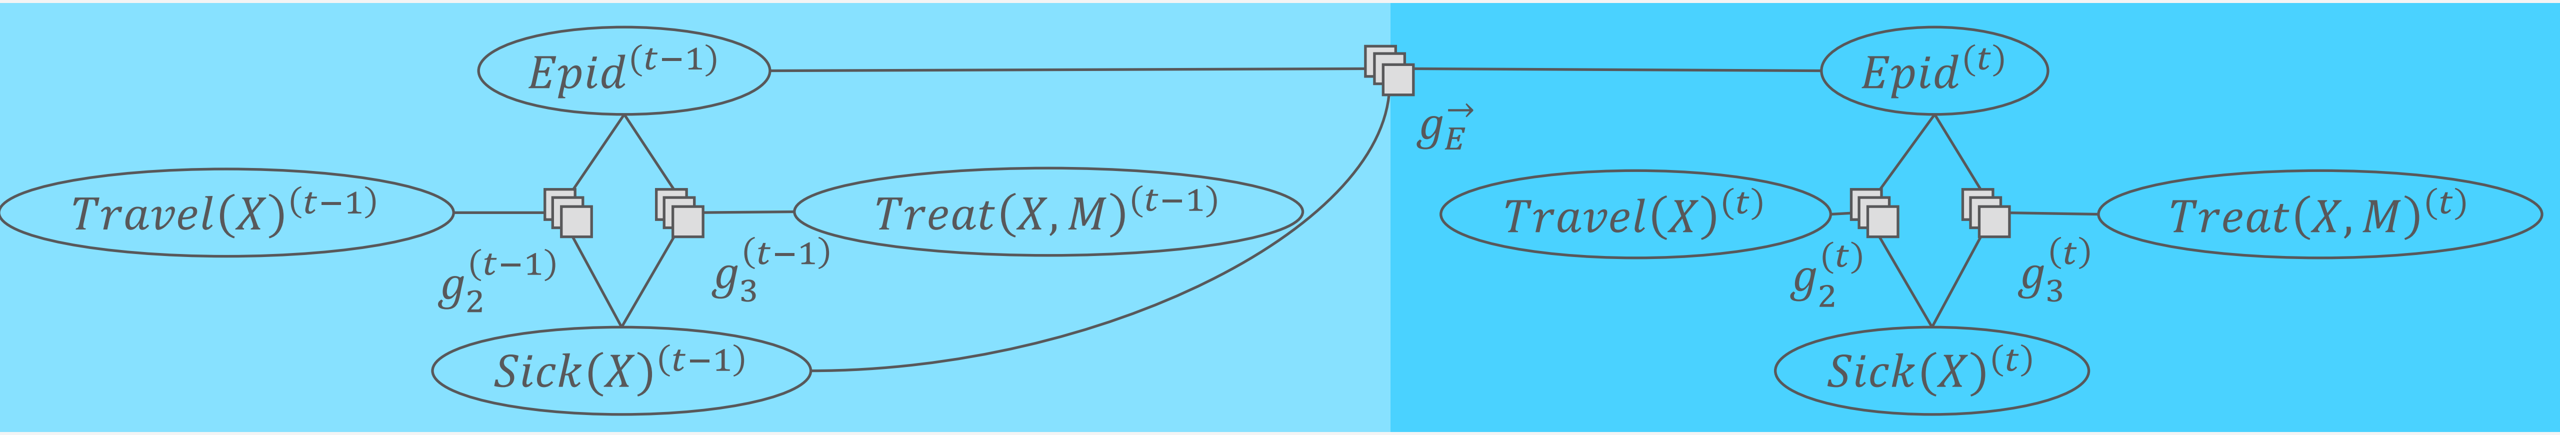

- Approximate forward message
- For each time step the temporal behaviour is multiplied on the forward message
- Indefinitely bounded error due to temporal behaviour

Results

- DBSCAN for Clustering
- ANOVA for checking fitness of clusters
- Right: runtimes
- Below: approximation error

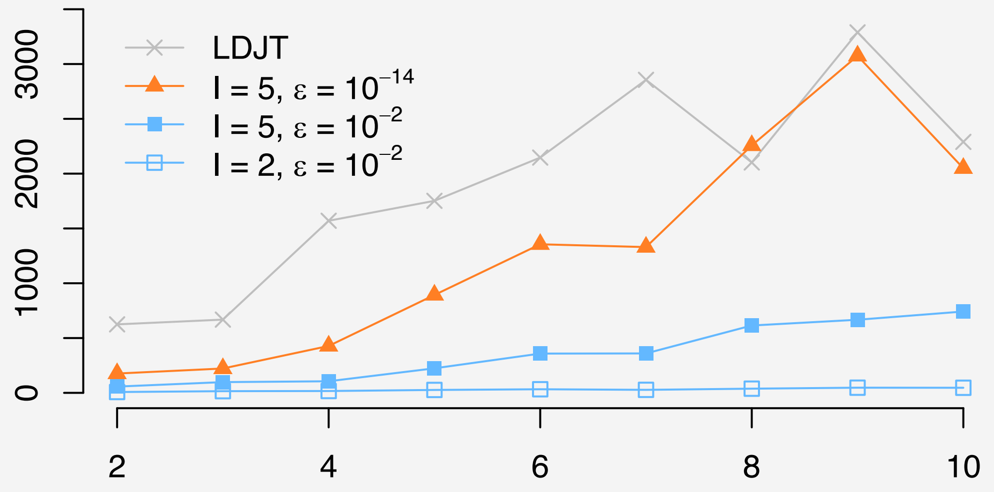

π	Max	Min	Average
0	0.0001537746121	0.0000000001720	0.0000191206488
2	0.0000000851654	0.0000000000001	0.0000000111949
4	0.0000000000478	0	0.0000000000068

Indistinguishability in Decision Making

The Power of Indistinguishability

Indistinguishability for Decision Making

- Online decision making: Graphical models extended by decision and utility nodes
  - Parameterise decisions to make decisions for whole groups of indistinguishable instances: $Treat(X, M)$ (box in graph)
  - PRVs in utility functions to denote identical share in contributed utility U (diamond in graph) : $\phi_U\big(Epid, Sick(X)\big)$
  - (Dynamic) decision parfactor models, Markov logic decision networks

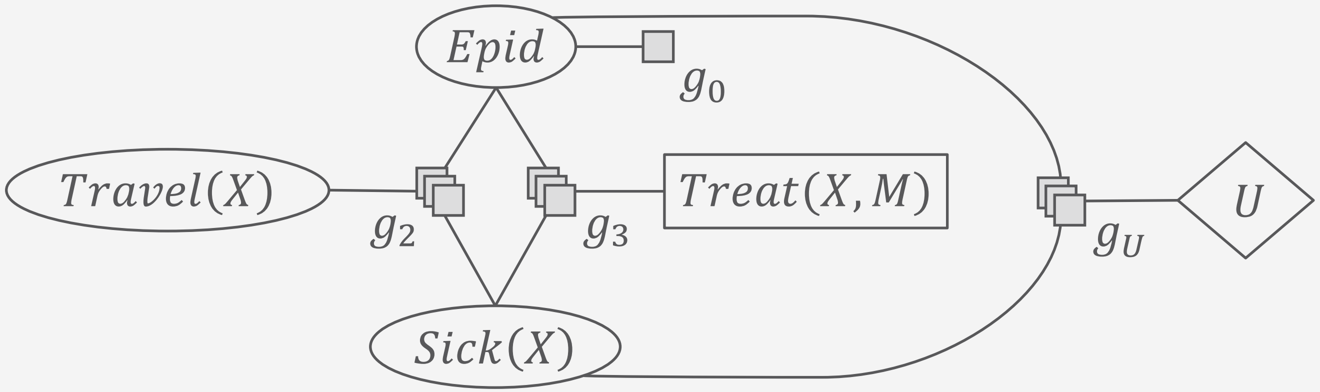

Gehrke et al. 2019b, Gehrke et al. 2019c; Nath & Domingos 2009

Indistinguishability for Decision Making

- Inference task: maximum expected utility (MEU) query
  - *Which actions can be expected to lead to the maximum utility?*
  - Standard inference algorithms more or less still work
    - Iterate through all possible decisions, set decisions as evidence, calculate expected utility, store current maximum
    - Solve an MAP query with decision variables as query terms and the other variables in the model to eliminate

Assign same action to group of indistinguishable instances
- Fewer possible decisions to consider → *tractability*!

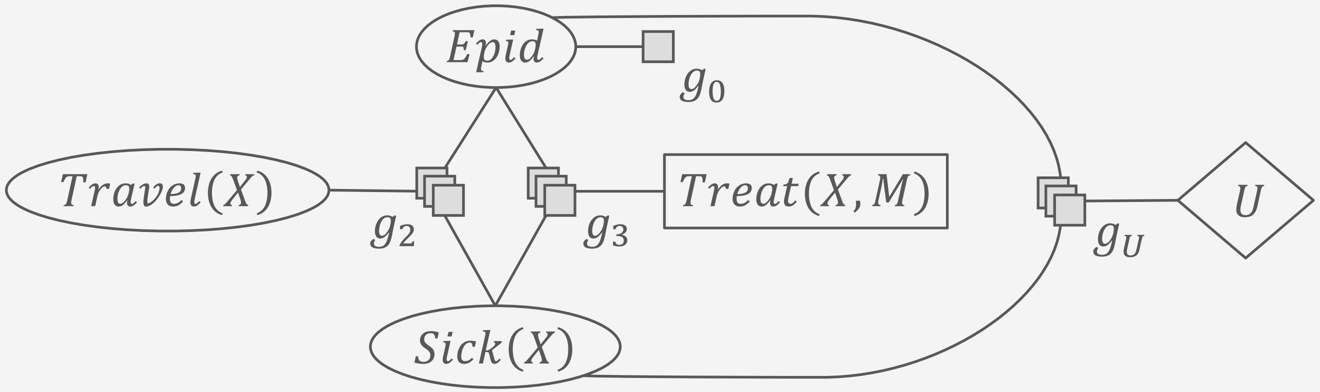

Indistinguishability for Decision Making

- Offline decision making: solve a (partially observable) Markov decision problem (POMDP)
  - First-order / relational MDPs: indistinguishability in the environment [Sanner & Kersting 2012]
    - Based on situation calculus: work with representatives
      - E.g., it is important that a box with medical supplies arrives at a destination but not which one it is in particular (of a set of boxes with medical supplies)
    - Novel propositional situations worth exploring may be instances of a well-known context in the relational setting → *exploitation* promising
      - E.g., household robot learning water-taps
      - Having opened one or two water-taps in a kitchen, one can expect other water-taps in kitchens to work similarly

    ⇒ Priority for exploring water-taps in kitchens in general reduced

    ⇒ Information gathered likely to carry over to water-taps in other places

    ❖ Hard to model in propositional setting: each water-tap is novel

Indistinguishability for Decision Making

- Multi-agent setting: decentralised POMDP [Oliehoek & Amato 2016]
  - Set of agents with
    - Own set of available actions, observations
    - *Shared* state and reward
- Lifting for agents [B et al. 2022]
  - Agents with indistinguishable behaviour → types
    - The same sets of actions, observations available
    - Same strategy / program applies if certain independences hold
  - Groups by types can be treated by representatives
    - Reduce exponential dependence on agent numbers
  - Application: Nanoagent network

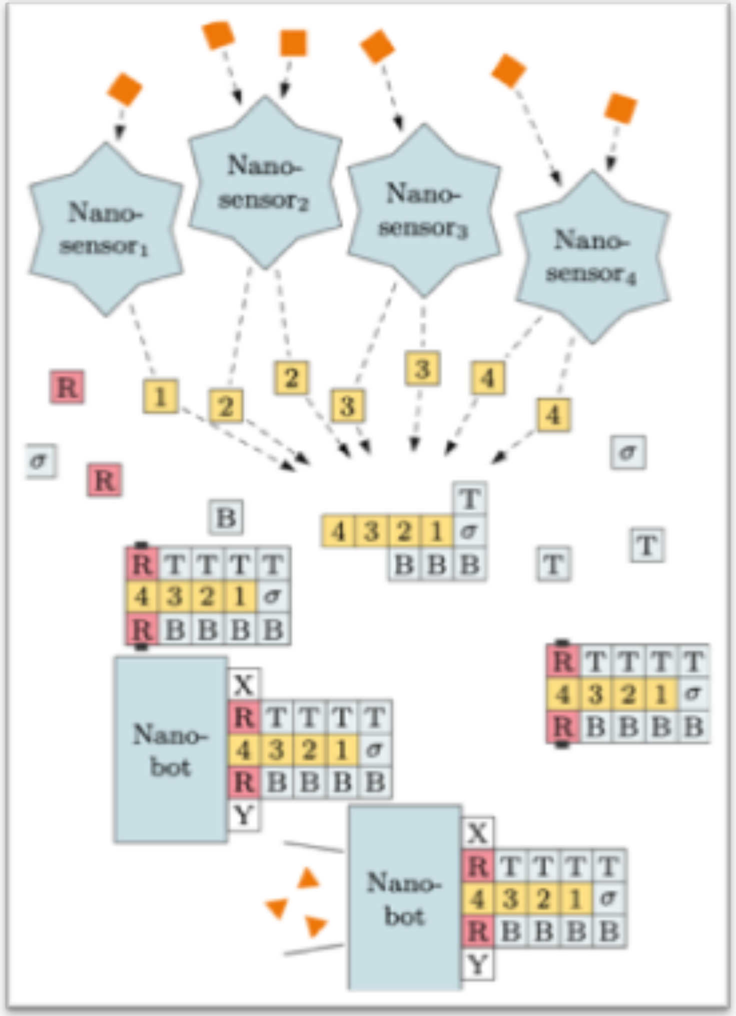

Agenda

- Statistical Relational Artificial Intelligence
  - Probabilistic relational models
  - Grounding semantics
  - Context
- The Power of Indistinguishability
  - Lifted query answering and tractability
  - Keeping indistinguishability over time
  - Indistinguishability in decision making
- Summary

The Finish Line: The Power of Indistinguishability

- Lifted query answering and tractability
  - Use information about indistinguishability to speed up inference
  - Tractability in terms of domain sizes through lifting
  - Handle evidence in groups of indistinguishable observations
  - Count values in histograms for lifted queries
- Keeping indistinguishability over time
  - Merge parfactors with bounded error
- Indistinguishability in decision making
  - Relational environment encoded
  - Agent types

What else is there to do? – Oh, so much…
- Approximating symmetries
- Generalising lifting operators
- More robust learning algorithms
- Privacy
- Ethical behaviour
- Explainability
- …

Thank you!

Bibliography & Further Papers

Ordered topic-wise and then alphabetically

Bibliography – General

- [Ahmadi et al. 13]
  Babak Ahmadi, Kristian Kersting, Martin Mladenov, and Sriraam Natarajan. Exploiting Symmetries for Scaling Loopy Belief Propagation and Relational Training. In Machine Learning. 92(1):91-132, 2013
- [Bach, Broecheler, Huang, Getoor 17]
  Stephen H. Bach, Matthias Broecheler, Bert Huang, Lise Getoor. Hinge-Loss Markov Random Fields and Probabilistic Soft Logic. In: J. Mach. Learn. Res. 18:. 109:1-109:67, 2017
- [Chavira and Darwiche 07]
  Mark Chavira, Adnan Darwiche. Compiling Bayesian Networks Using Variable Elimination. In: Proc. IJCAI 2007: 2443-2449
- [Cohen et al. 17] Cohen, William W., Fan Yang and Kathryn Mazaitis. “TensorLog: Deep Learning Meets Probabilistic DBs.” ArXiv abs/1707.05390 (2017)
- [De Salvo Braz et al. 05]
  Rodrigo de Salvo Braz, Eyal Amir, and Dan Roth. Lifted First-order Probabilistic Inference. IJCAI-05 Proceedings of the 19th International Joint Conference on Artificial Intelligence, 2005
- [De Salvo Braz et al. 06]
  Rodrigo de Salvo Braz, Eyal Amir, and Dan Roth. MPE and Partial Inversion in Lifted Probabilistic Variable Elimination. AAAI-06 Proceedings of the 21st Conference on Artificial Intelligence, 2006
- [Gogate and Domingos 11]
  Vibhav Gogate and Pedro Domingos. Probabilistic Theorem Proving. In: Proc. UAI 2011: 256–265
- [Kersting et al. 09]
  Kristian Kersting, Babak Ahmadi, and Sriraam Natarajan. Counting Belief Propagation. In UAI-09 Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence, 2009
- [Milch et al. 08]
  Brian Milch, Luke S. Zettelmoyer, Kristian Kersting, Michael Haimes, and Leslie Pack Kaelbling. Lifted Probabilistic Inference with Counting Formulas. In AAAI-08 Proceedings of the 23rd AAAI Conference on Artificial Intelligence, 2008
- [Poole 03]
  David Poole. First-order probabilistic inference. IJCAI 2003: 985-991
- [Poole & Mackworth 03]
  David Poole, Alan Mackworth. Artificial Intelligence: Foundations of Computational Agents, 2nd Edition, Cambridge University Press, 2017

Bibliography

- [Richardson & Domingos 06]
  Matthew Richardson, Pedro Domingos. Markov logic networks. In: J. Machine Learning. Band 62. Nr. 1-2. 2006. 107–136
- [Russell & Norvig 16] Stuart Russell, Peter Norvig, Artificial Intelligence: A Modern Approach, Pearson, 2016
- [Sarkhel, Venugopal et al. 14]
  Somdeb Sarkhel, Deepak Venugopal, Parag Singla, Vibhav Gogate:. Lifted MAP Inference for Markov Logic Networks. AISTATS 2014: 859-867
- [Singla and Domingos 08]
  Parag Singla and Pedro Domingos. Lifted First-order Belief Propagation. In AAAI-08 Proceedings of the 23rd AAAI Conference on Artificial Intelligence, 2008
- [Tagipour et al. 13]
  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
- [Taghipour et al. 13a]
  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
- [van den Broeck 13]
  Guy Van den Broeck. Lifted Inference and Learning in Statistical Relational Models, PhD thesis, KU Leuven, 2013

WWU MÜNSTER

Bibliography

- [Braun & Möller 16]
  Tanya Braun and Ralf Möller. Lifted Junction Tree Algorithm. In Proceedings of KI 2016: Advances in Artificial Intelligence, pages 30–42, 2016
- [Braun & Möller 17]
  Tanya Braun and Ralf Möller. Preventing Groundings and Handling Evidence in the Lifted Junction Tree Algorithm. In Proceedings of KI 2017: Advances in Artificial Intelligence, pages 85–98, 2017
- [Braun & Möller 17a]
  Tanya Braun and Ralf Möller. Counting and Conjunctive Queries in the Lifted Junction Tree Algorithm. In Postproceedings of the 5th International Workshop on Graph Structures for Knowledge Representation and Reasoning, 2017
- [Braun & Möller 18]
  Tanya Braun and Ralf Möller. Adaptive Inference on Probabilistic Relational Models. In Proceedings of the 31st Australasian Joint Conference on Artificial Intelligence, 2018
- [Braun & Möller 18a]
  Tanya Braun and Ralf Möller. Parameterised Queries and Lifted Query Answering. In IJCAI-18 Proceedings of the 27th International Joint Conference on Artificial Intelligence, 2018
- [Braun & Möller 18b]
  Tanya Braun and Ralf Möller. Lifted Most Probable Explanation. In Proceedings of the International Conference on Conceptual Structures, 2018
- [Braun & Möller 18c]
  Tanya Braun and Ralf Möller. Fusing First-order Knowledge Compilation and the Lifted Junction Tree Algorithm. In Proceedings of KI 2018: Advances in Artificial Intelligence, 2018
- [Braun & Möller 19]
  Tanya Braun, Ralf Möller: Exploring Unknown Universes in Probabilistic Relational Models, in: Proceedings of AI 2019: Advances in Artificial Intelligence, 2019

Bibliography

- [Gehrke et al. 18]
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Dynamic Junction Tree Algorithm. In Proceedings of the International Conference on Conceptual Structures, 2018
- [Gehrke et al. 18b]
  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, 2018
- [Gehrke et al. 18c]
  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, 2018
- [Gehrke et al. 19]
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Relational Forward Backward Algorithm for Multiple Queries. In FLAIRS-32 Proceedings of the 32nd International Florida Artificial Intelligence Research Society Conference, 2019
- [Gehrke et al. 19b]
  Marcel Gehrke, Tanya Braun, Ralf Möller, Alexander Waschkau, Christoph Strumann, and Jost Steinhäuser. Lifted Maximum Expected Utility. In Artificial Intelligence in Health, 2019
- [Gehrke et al. 19c]
  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, 2019
- [Gehrke et al. 19d]
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Temporal Most Probable Explanation In Proceedings of the International Conference on Conceptual Structures, 2019
- [Gehrke et al. 19e]
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Taming Reasoning in Temporal Probabilistic Relational Models Technical report
- [Gehrke et al. 19f]
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Uncertain Evidence in Probabilistic Relational Models. In *Proceedings of the 32nd Canadian Conference on Artificial Intelligence, Canadian AI 2019*, 2019

Bibliography – Query Answering

- Ahmadi et al. (2013)
  Babak Ahmadi, Kristian Kersting, Martin Mladenov, and Sriraam Natarajan. Exploiting Symmetries for Scaling Loopy Belief Propagation and Relational Training. In *Machine Learning*. 92(1):91-132, 2013.
- B (2020)
  Tanya Braun. Rescued from a Sea of Queries: Exact Inference in Probabilistic Relational Models. PhD Thesis, 2020.
- B & Möller (2018)
  Tanya Braun and Ralf Möller. Parameterised Queries and Lifted Query Answering. In *IJCAI-18 Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 2018.
- B & Möller (2019)
  Tanya Braun and Ralf Möller. Exploring Unknown Universes in Probabilistic Relational Models. In *Proceedings of AI 2019: Advances in Artificial Intelligence*, 2019.

Bibliography

- Jaeger & Schulte (2018)
  Manfred Jaeger and Oliver Schulte. Inference, Learning, and Population Size: Projectivity for SRL Models. In *StaRAI-18 Workshop on Statistical Relational Artificial Intelligence*, 2018.
- Kersting et al. (2009)
  Kristian Kersting, Babak Ahmadi, and Sriraam Natarajan. Counting Belief Propagation. In *UAI-09 Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 2009.
- Lauritzen & Spiegelhalter (1988)
  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*, 50:157–224, 1988.
- Mittal et al. (2019)
  Happy Mittal, Ayush Bhardwaj, Vibhav Gogate, and Parag Singla. Domain-size Aware Markov Logic Networks. In *AISTATS-19 Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, 2019.

Bibliography

- Niepert & Van den Broeck (2014)
  Mathias Niepert and Guy Van den Broeck. Tractability through Exchangeability: A New Perspective on Efficient Probabilistic Inference. In *AAAI-14 Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 2014.
- Pearl (1982)
  Judea Pearl. Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach. In *AAAI-82 Proceedings of the 2nd National Conference on Artificial Intelligence*, 1982.
- Poole (2003)
  David Poole. First-order Probabilistic Inference. In *IJCAI-03 Proceedings of the 18th International Joint Conference on Artificial Intelligence*, 2003.
- Poole et al. (2014)
  David Poole, David Buchman, Seyed Mehran Kazemi, Kristian Kersting, and Sriraam Natarajan. Population Size Extrapolation in Relational Probabilistic Modeling. In *SUM-14 Proceedings of the 8th International Conference on Scalable Uncertainty Management*, 2014.

Bibliography

- Singla & Domingos (2008)
  Parag Singla and Pedro Domingos. Lifted First-order Belief Propagation. In *AAAI-08 Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, 2008.
- Taghipour et al. (2013)
  Nima Taghipour, Daan Fierens, Guy Van den Broeck, Jesse Davis, and Hendrik Blockeel. Completeness Results for Lifted Variable Elimination. In *AISTATS-13 Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, 2013.
- Taghipour et al. (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, 2013.
- Van den Broeck (2011)
  Guy Van den Broeck. On the Completeness of First-order Knowledge Compilation for Lifted Probabilistic Inference. In *NIPS-11 Advances in Neural Information Processing Systems 24*, 2011.

Bibliography

- Van den Broeck & Darwiche (2013)
  Guy Van den Broeck and Adnan Darwiche. On the Complexity and Approximation of Binary Evidence in Lifted Inference. In *NIPS-13 Advances in Neural Information Processing Systems 26*, 2013.
- Van den Broeck & Davis (2012)
  Guy Van den Broeck and Jesse Davis. Conditioning in First-Order Knowledge Compilation and Lifted Probabilistic Inference. In *AAAI-12 Proceedings of the 26th AAAI Conference on Artificial Intelligence*, 2012.
- Van den Broeck & Niepert (2015)
  Guy Van den Broeck and Mathias Niepert: Lifted Probabilistic Inference for Asymmetric Graphical Models. In *AAAI-15 Proceedings of 29th AAAI Conference on Artificial Intelligence*, 2015.

Bibliography – Temporal Models

- Ahmadi et al. (2013)
  Babak Ahmadi, Kristian Kersting, Martin Mladenov, and Sriraam Natarajan. Exploiting Symmetries for Scaling Loopy Belief Propagation and Relational Training. In *Machine Learning*. 92(1):91-132, 2013.
- Gehrke et al. (2018)
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Dynamic Junction Tree Algorithm. In *ICCS-18 Proceedings of the International Conference on Conceptual Structures*, 2018.
- Gehrke et al. (2019)
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Relational Forward Backward Algorithm for Multiple Queries. In *FLAIRS-32 Proceedings of the 32nd International Florida Artificial Intelligence Research Society Conference*, 2019.
- Gehrke et al. (2019a)
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Temporal Most Probable Explanation. In *ICCS-19 Proceedings of the International Conference on Conceptual Structures*, 2019.

Bibliography

- Gehrke et al. (2020)
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Taming Reasoning in Temporal Probabilistic Relational Models Explanation. In *Proceedings of the ECAI 2020*, 2020.
- Mladenov et al. (2017)
  Martin Mladenov, Leonard Kleinhans, Kristian Kersting: Lifted Inference for Convex Quadratic Programs. In *AAAI-17 Proceedings of 31st AAAI Conference on Artificial Intelligence*, 2017.
- Murphy (2002)
  Kevin P. Murphy. Dynamic Bayesian Networks: Representation, Inference and Learning. *PhD Thesis University of California, Berkeley*, 2002.
- Venugopal & Gogate (2014)
  Deepak Venugopal and Vibhav Gogate: Evidence-Based Clustering for Scalable Inference in Markov Logic. In *ECML PKDD 2014: Machine Learning and Knowledge Discovery in Databases*, 2014.

WWU MÜNSTER

Bibliography – Decision Making

- B et al. (2022)
  Tanya Braun, Marcel Gehrke, Florian Lau, and Ralf Möller: Lifting in Multi-agent Systems under Uncertainty. In *UAI-22 International Conference on Uncertainty in Artificial Intelligence*, 2022.
- Gehrke et al. (2019b)
  Marcel Gehrke, Tanya Braun, Ralf Möller, Alexander Waschkau, Christoph Strumann, and Jost Steinhäuser. Lifted Maximum Expected Utility. In *Artificial Intelligence in Health*, 2019.
- Gehrke et al. (2019c)
  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Temporal Maximum Expected Utility. In *CanadianAI-19 Proceedings of the 32nd Canadian Conference on Artificial Intelligence*, 2019.
- Nath & Domingos (2009)
  Aniruddh Nath and Pedro Domingos. A Language for Relational Decision Theory. In *Proceedings of the 6th International Workshop on Statistical Relational Learning*, 2009.

Bibliography – Decision Making

- Oliehoek & Amato (2016)
  Frans A. Oliehoek and Christopher Amato. A Concise Introduction to Decentralised POMDPs, 2019.
- Sanner & Kersting (2010)
  Scott Sanner and Kristian Kersting. Symbolic Dynamic Programming for First-order POMDPs. In *AAAI-10 Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 2010.