Statistical Relational AI & Robustness

Workshop on Robust AI for High-Stakes Applications

@KI2022

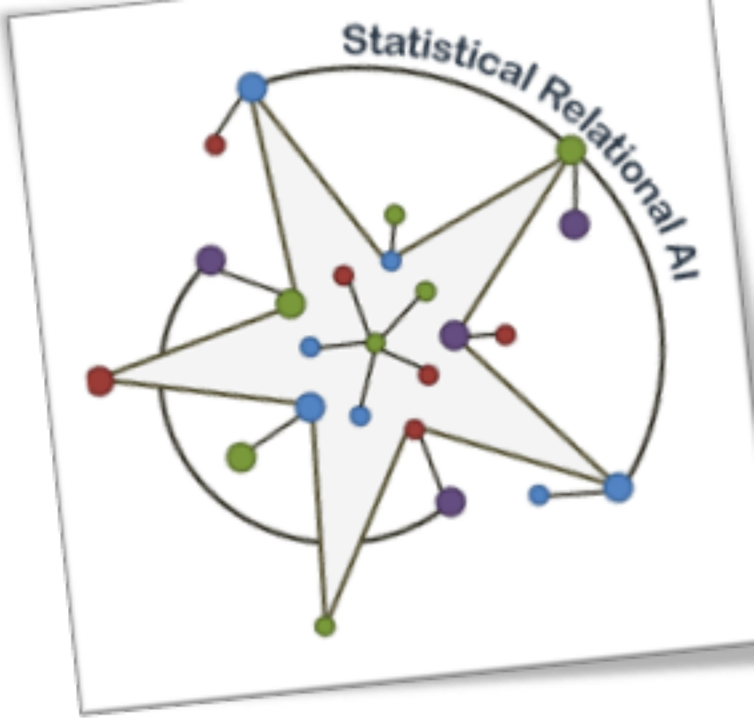

Tanya Braun

University of Münster

WWU MÜNSTER

Agenda

- Probabilistic relational models (PRMs)
  - Application example
  - Semantics
  - Query answering / basic inference
- Lifting Algorithms for More Robust Inference
  - Cluster trees for efficient multi-query inference
  - Adaptive inference in cluster trees
  - Changing domains
  - Keeping inference going over time
- Summary

Application: Epidemics

- Atoms: Parameterised random variables = PRVs
  - With logical variables
    - E.g., X, M
    - Possible values (domain):
    $\mathcal{D}(X) = \{alice, eve, bob\}$
    $\mathcal{D}(M) = \{injection, tablet\}$
  - With range
    - E.g., Boolean
    - $range(Travel(X))$
    $r(Travel(X))$

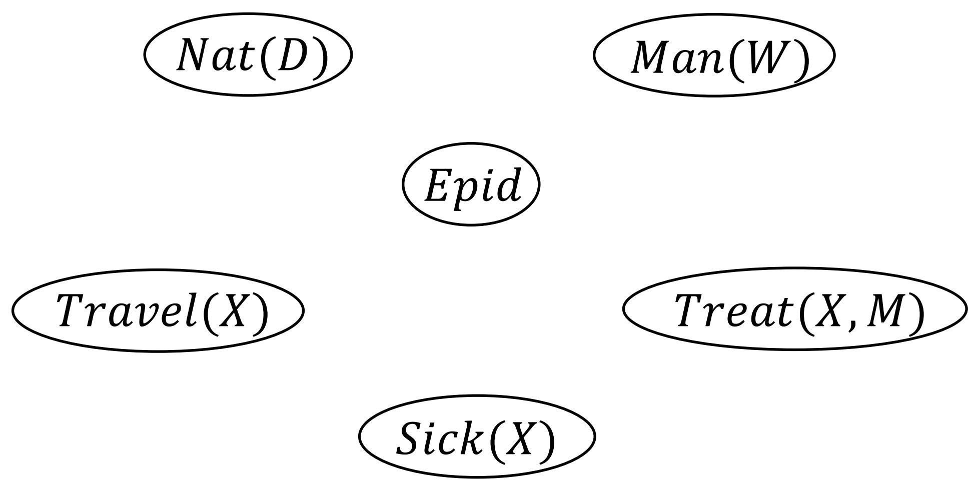

$Nat(D) = natural\ disaster(D)$
$Man(W) = man - made\ disaster(W)$

Encoding the Joint Distribution

- Factors with PRVs = parfactors
  - (Graphical) Model G
  - 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

Sparse encoding of joint distribution

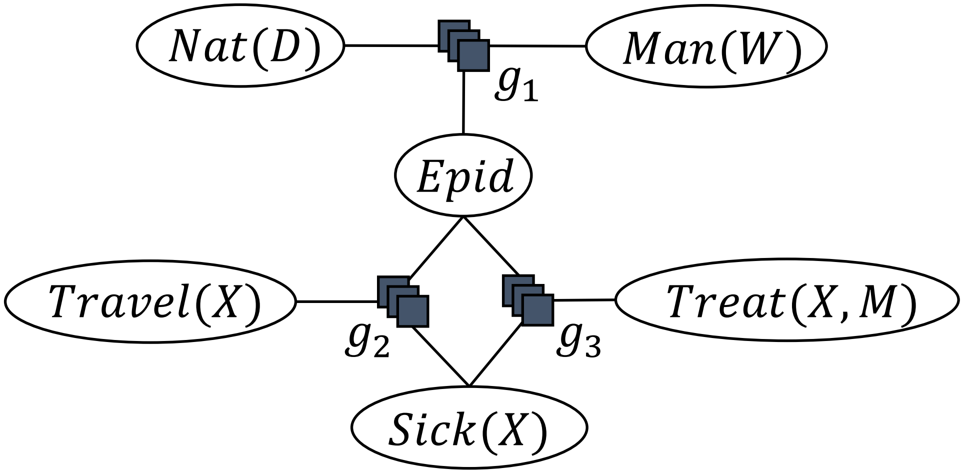

$3 \cdot 2^3 = 24$ entries in 3 parfactors, 6 PRVs

WWU MÜNSTER

Factors

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

$Travel(X)$	$Epid$

$Travel(alice)$	Epi		
false	*fals*		
false	*fals*		
false	*true*		
false	*true*		
true	*fals*		
true	*fals*		
true	*true*	*false*	2
true	*true*	*true*	9

$Travel(eve)$	$Epid$	$Sick(eve)$	f_2^1
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

(bob)	$Epid$	$Sick(bob)$	f_2^3
se	*false*	*false*	5
se	*false*	*true*	0
se	*true*	*false*	4
se	*true*	*true*	6
e	*false*	*false*	4
e	*false*	*true*	6
e	*true*	*false*	2
true	*true*	*true*	9

Semantics of a PRM

- Joint probability distribution P_G by grounding, multiplying all grounded factors, and normalising the result

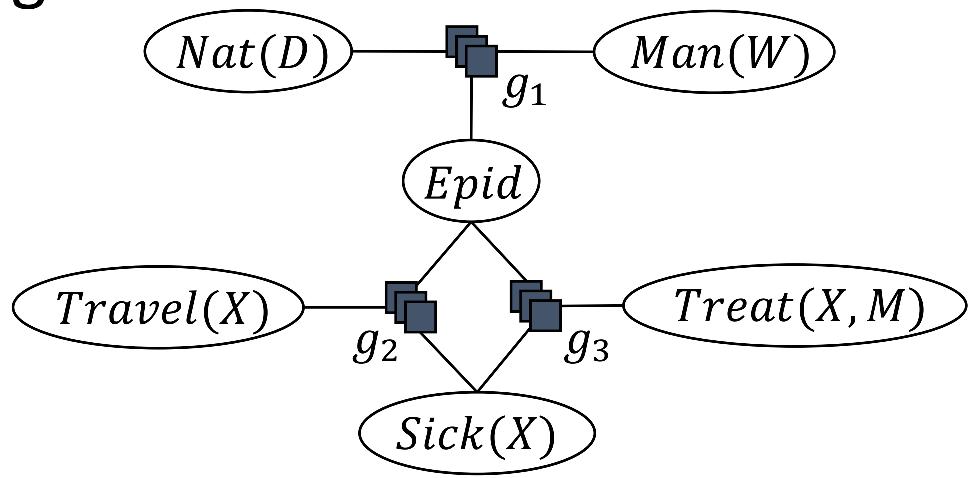

$$P_G = \frac{1}{Z} \prod_{f \in gr(G)} f$$

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

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

Grounded Model

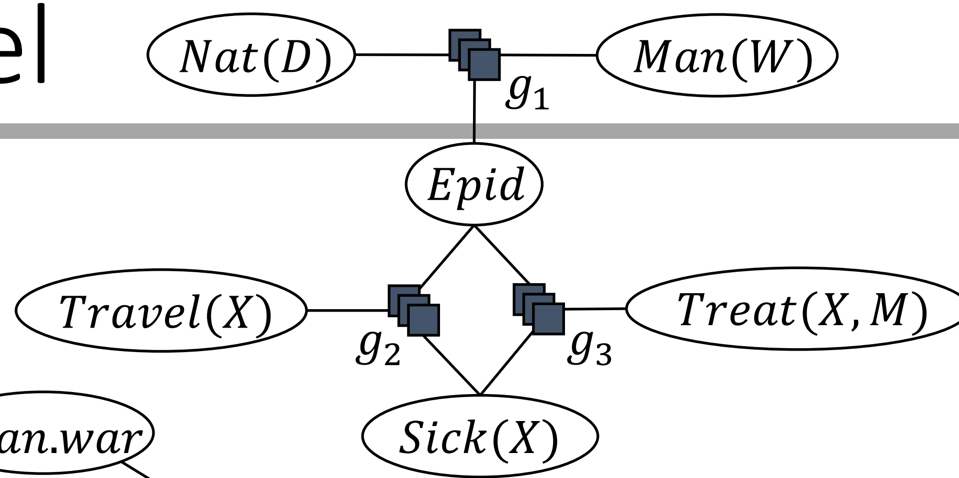

- Given domains
  - $\mathcal{D}(X) = \{alice, eve, bob\}$
  - $\mathcal{D}(M) = \{m_1, m_2\}$
  - $\mathcal{D}(D) = \{flood, fire\}$
  - $\mathcal{D}(W) = \{virus, war\}$

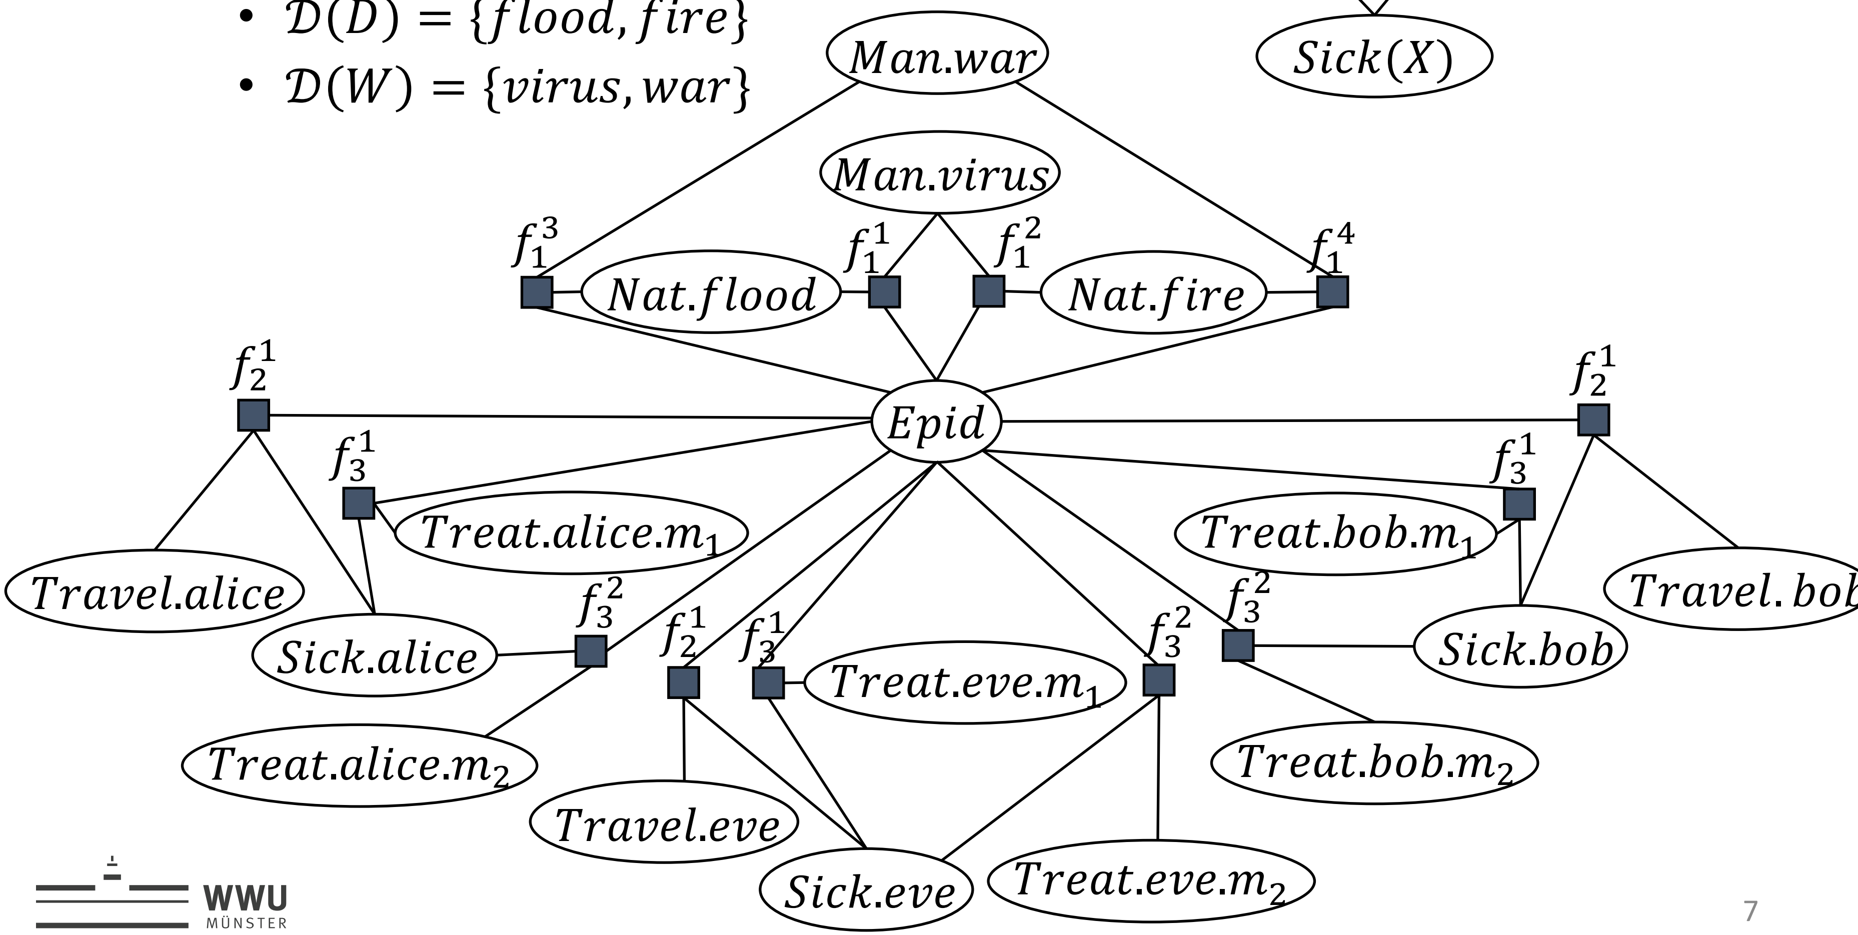

PRMs and Variants

- Probabilistic Relational Models (raw PRMs)
  [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]

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
  - MPE
  - MAP

Avoid groundings!

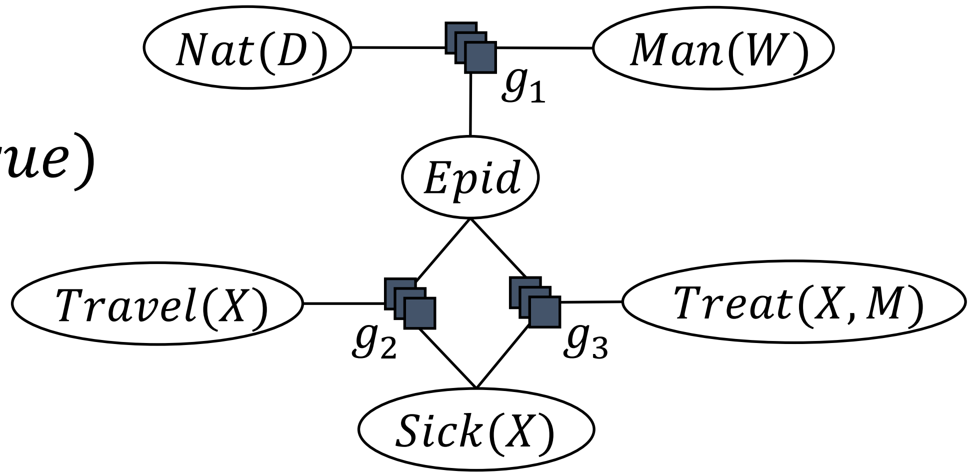

QA: Lifted Variable Elimination (LVE)

[Poole 03, de Salvo Braz et al. 05, 06, Milch et al. 08, Taghipour et al. 13, 13a]

- Eliminate all variables not appearing in query
- Lifted summing out
  1. Sum out representative instance as in propositional variable elimination
  2. Exponentiate result for isomorphic instances

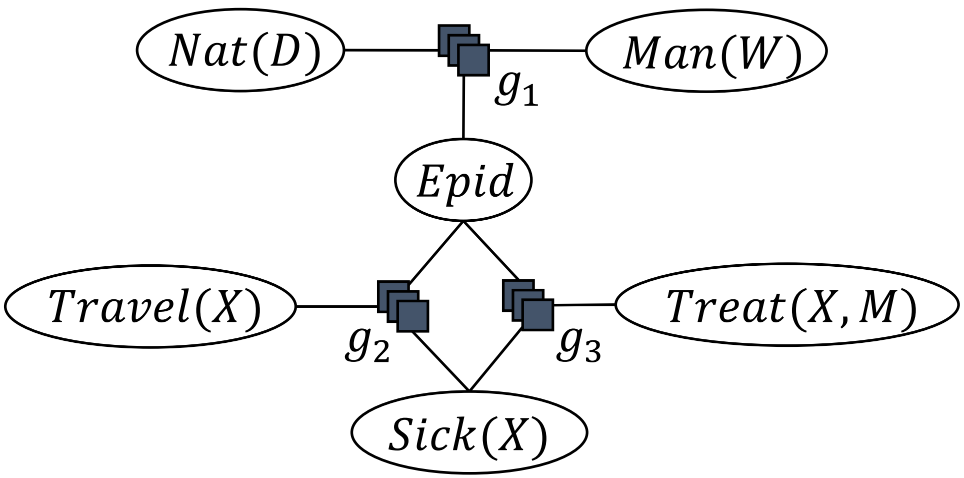

Avoid groundings!

QA: LVE in Detail

- E.g., marginal
  - $P(Travel(eve))$
  - Split atoms $P(\dots, X, \dots)$ w.r.t. eve if eve in $\mathcal{D}(X)$

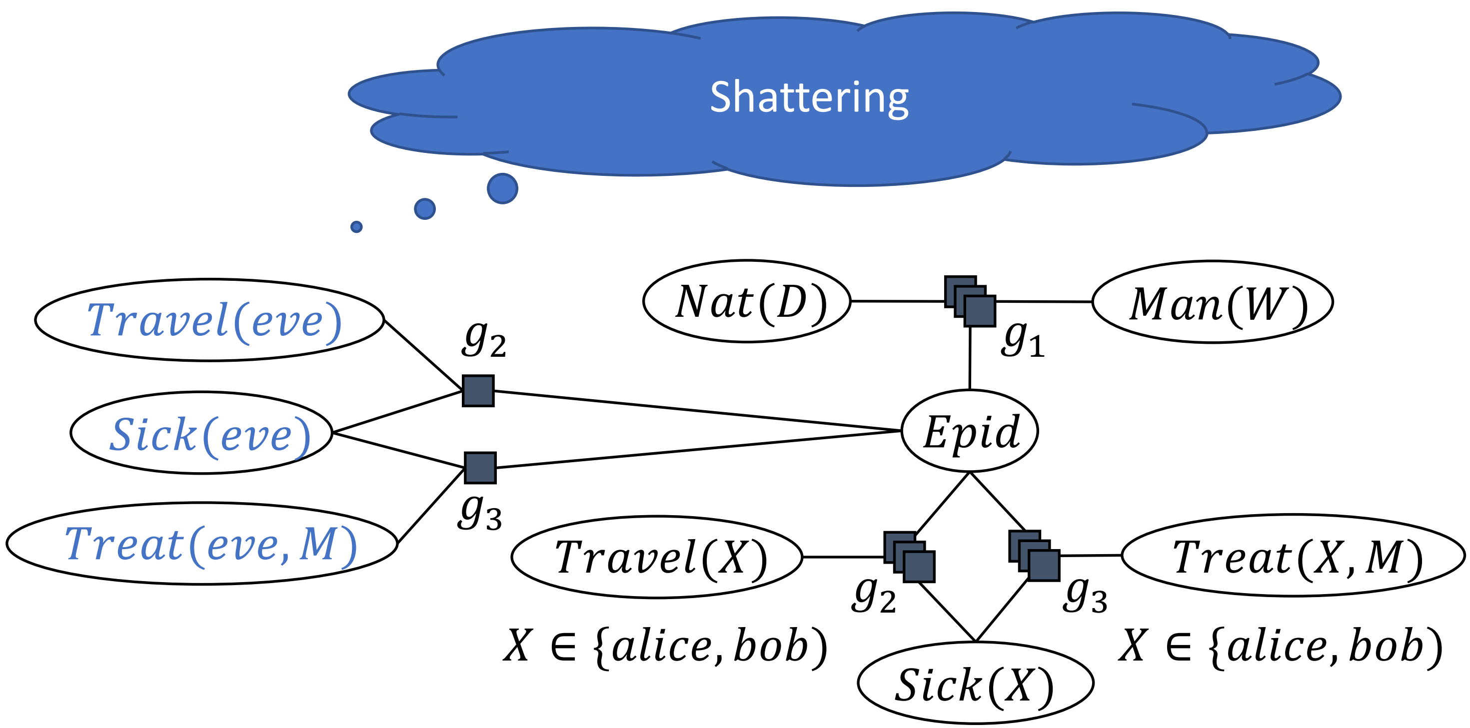

QA: LVE in Detail

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

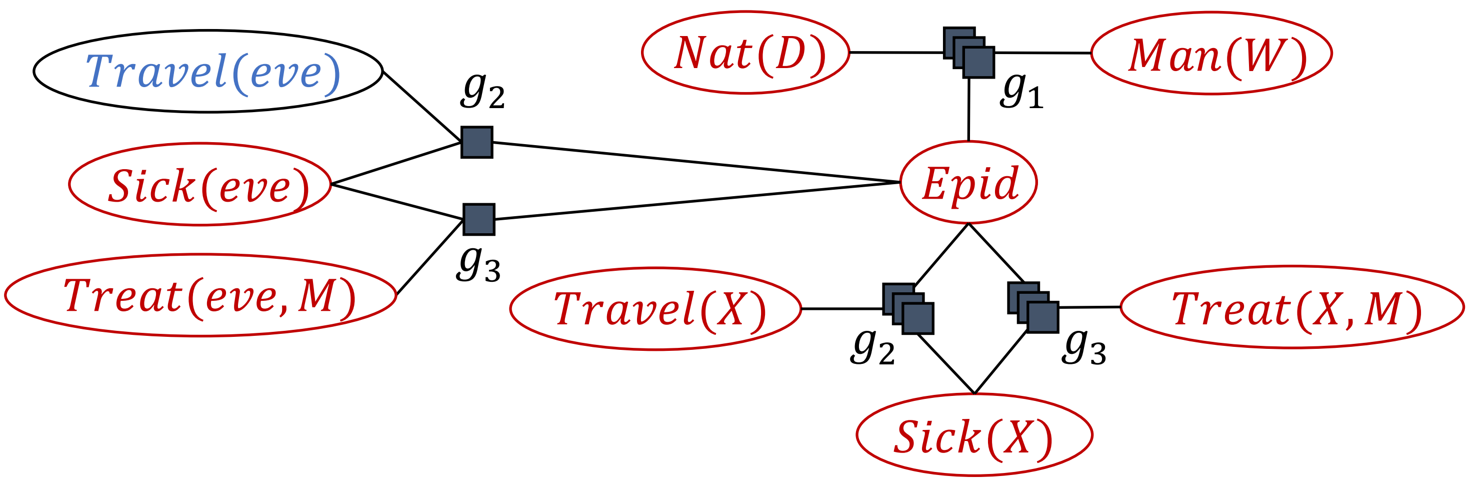

QA: LVE in Detail

- 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

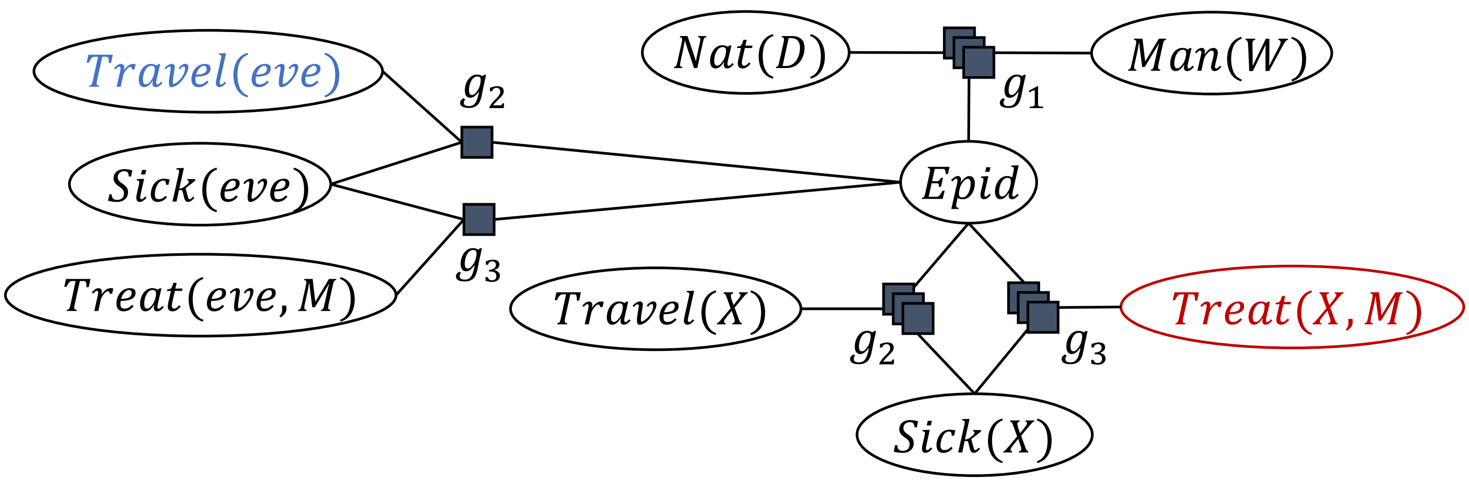

LVE in Detail: Lifted Summing Out

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

Only here, domain size comes into play → no change in structure / equation if domain size changes

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

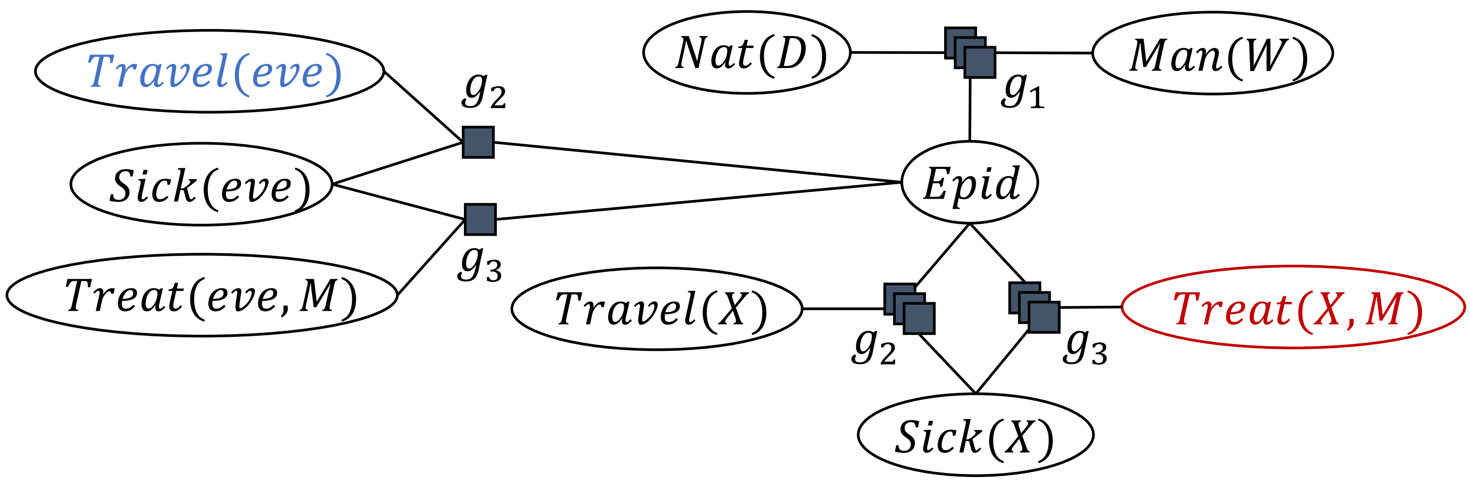

Tractability

- 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 and Van den Broeck (2014)

Agenda

- Probabilistic relational models (PRMs)
  - Application example
  - Semantics, static vs. dynamic behavior
  - Query answering / basic inference
- Algorithms for More Robust Inference
  - Cluster trees for efficient multi-query inference
  - Adaptive inference in cluster trees
  - Changing domains
  - Keeping inference going over time
- Summary

Cluster Trees for Efficient Multi-query Inference

Algorithms for More Robust Inference

WWU
MÜNSTER

Many Queries: LJT

- Set of queries
  - $P(Travel(eve))$
  - $P(Sick(bob))$
  - $P(Treat(eve, m_1))$
  - $P(Epid)$
  - $P(Nat(flood))$
  - $P(Man(virus))$
  - Combinations of variables
- Under evidence
  - $Sick(X') = true$
  - $X' \in \{alice, eve\}$

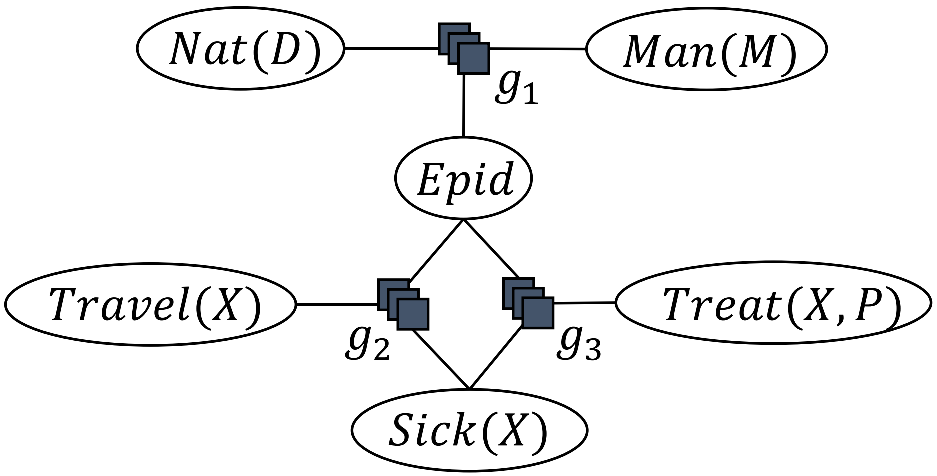

- **Challenges:**
  - **Do not start from scratch for every query**
  - **Support QA on subset of atoms**
  - **Avoid groundings**
- **Cluster tree based on (conditional) independences**

Solution: Submodels

Lauritzen and Spiegelhalter (1988), Shafer and Shenoy (1989), Jensen et al. (1990), Braun and Möller (2016)

- Identify submodel sufficient for query
  - Find PRVs that make submodel independent from remaining model
    - Separator
  - "Query" over separator collects all influences of remaining model on PRVs in submodel
    - PRVs of submodel = parcluster

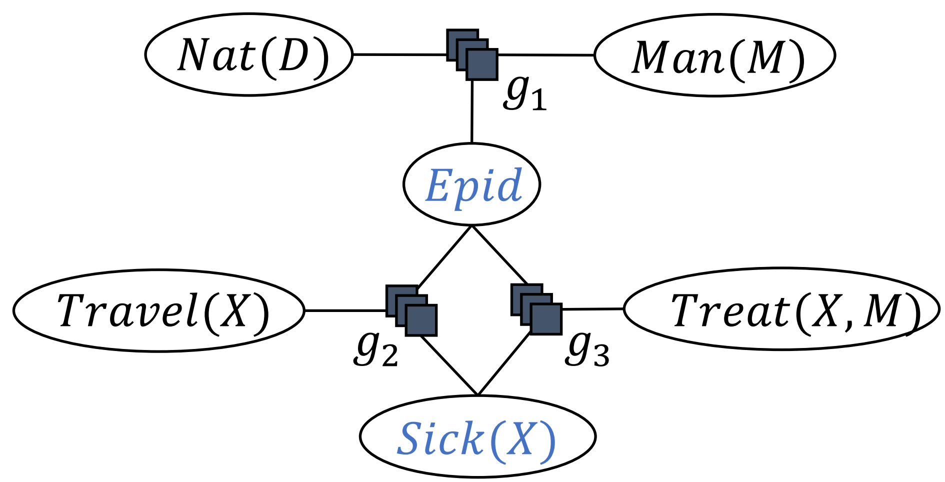

WWU MÜNSTER

Solution: Submodels

Lauritzen and Spiegelhalter (1988), Shafer and Shenoy (1989), Jensen et al. (1990), Braun and Möller (2016)

- Network of submodels with separators
  - Recursive “queries” to make submodels independent from each other
  - (First-order) Junction tree
    - Acyclic, running intersection property
- Recursive queries from each node
  - Arrange queries using dynamic programming
    - Also known as message passing

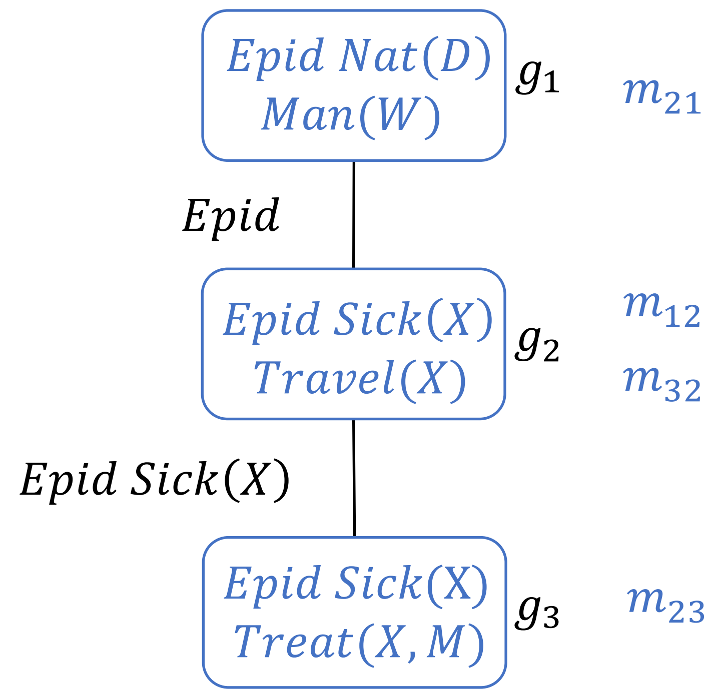

Answer queries on subtree over the query terms
- Use middle cluster for $P\big(Sick(eve)\big)$

Adaptive Inference in Cluster Trees

Algorithms for More Robust Inference

Adaptive Inference

- After changes in queries, evidence, model:
  Avoid starting from scratch to fast reach the point of answering queries again
  → adaptive inference
- Small, local model changes may preserve much of tree
  - If only local changes, up to half of messages still valid
    - Only resend messages
      if local model or incoming information changed

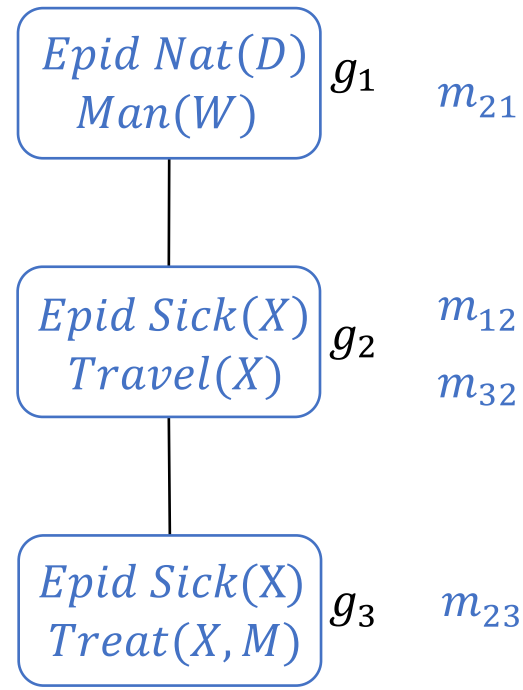

Tanya B and Ralf Möller: Adaptive Inference on Probabilistic Relational Models. In *Proceedings of the 31st Australasian Joint Conference on Artificial Intelligence*, 2018

Adaptive Inference: Changes

- Queries: no change
- Evidence: changes local models
  - New observations incoming
- Model
  - Potentials: changes local models
  - Domain sizes: changes local models
    - Nice property of relational models: No effect on model structure!
      - E.g., more people in $\mathrm{dom}(X)$
      - Propositional models: number of variables changes, which changes the tree structure
  - Parfactors (addition, deletion): changes tree structure

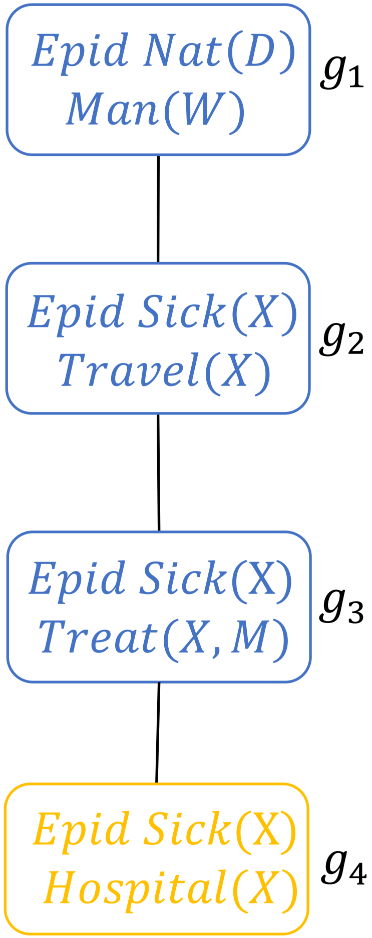

Changing Domains

Algorithms for More Robust Inference

Changing Domains

- Keep semantics as before
  - Assume that parfactors accurately describe world
- Posterior probabilities change depending on domain sizes
  - Example by Poole (2003)

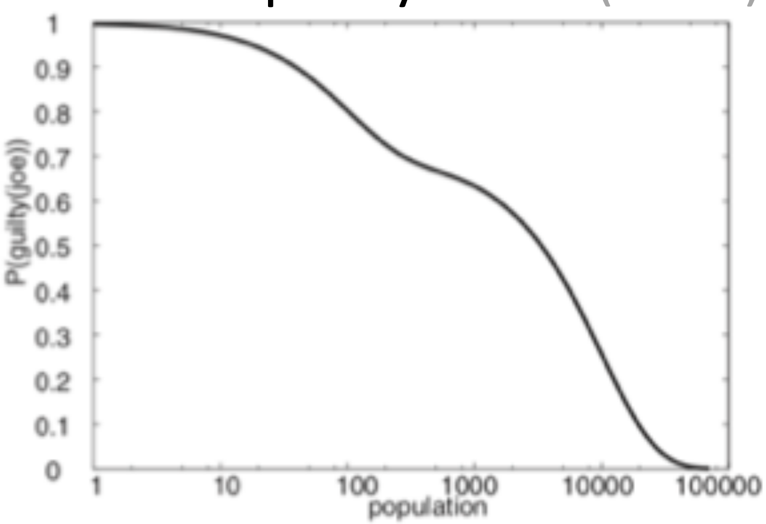

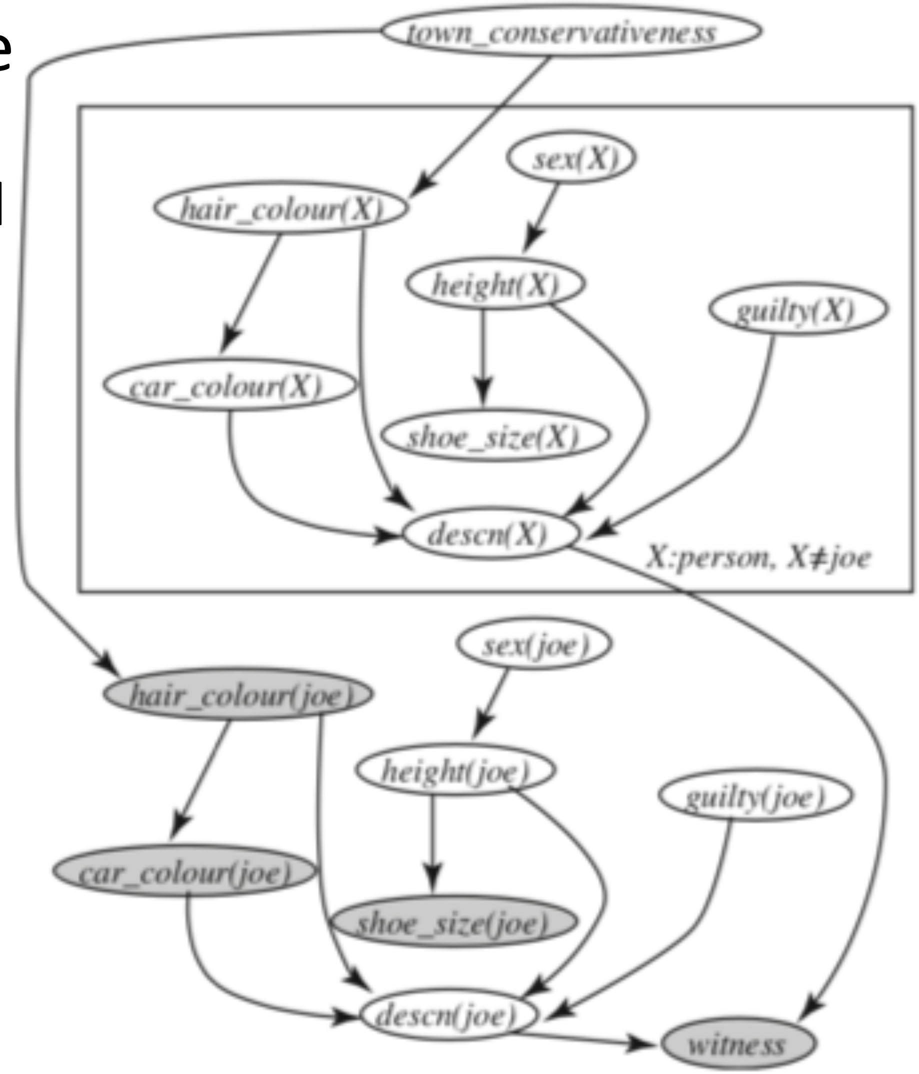

… Without Effects

- (Conditional) Independence
  PRVs, containing logical variables $\boldsymbol{X}$, that are (conditionally) independent from query terms → domains of $\boldsymbol{X}$ have no influence on query results
- E.g., given $Epid = e$,
  - $\mathcal{D}(D)$ and $\mathcal{D}(W)$ do not matter for queries regarding $Travel$, $Sick$, and $Treat$
  - $\mathcal{D}(X)$ and $\mathcal{D}(M)$ do not matter for queries regarding Nat and Man

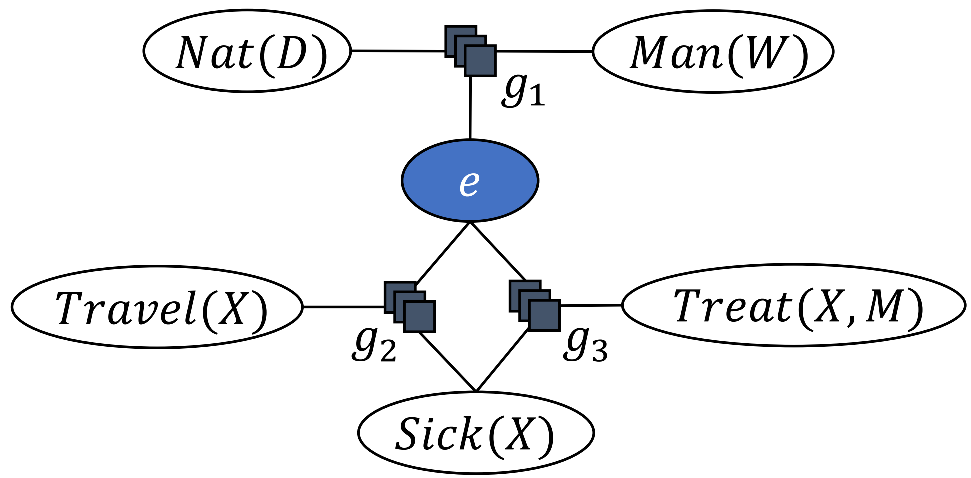

→ Partly invariant under increasing domain sizes

… Without Effects

- A simple case of so-called projectivity
  After shattering, query terms are independent of model parts containing logical variables $\boldsymbol{X}$ → domains of $\boldsymbol{X}$ have no influence on query results
  - Depends on model structure
  - More by Jaeger and Schulte (2018)
- E.g., $P(Sick(x_1))$
  - $\mathcal{D}(X) = \{x_1, \dots, x_n\}$
  - After shattering:
    - $\mathcal{D}(X) = \{x_2, \dots, x_n\}$
    - Upper part independent from lower part; $\mathcal{D}(X)$ irrelevant

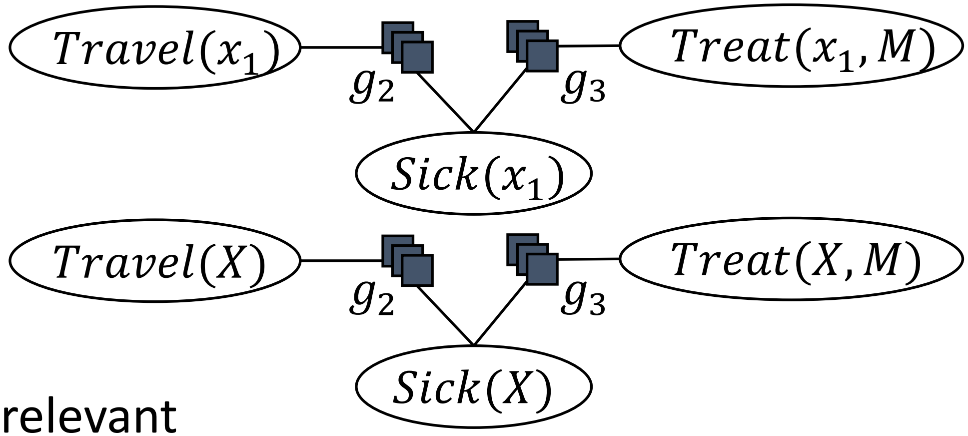

→ Partly invariant under increasing domain sizes

Growing Domain Sizes

Poole et al. (2014)

- Let domain size n grow
  - With grounding semantics, posteriors change
    - Can lead to extreme behaviour in the posteriors
- Example: $Epid$ gets more and more neighbours with n rising

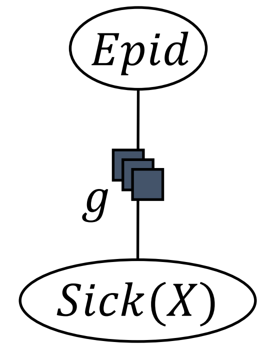

$$P(Epid) \propto \left(\sum_{s \in r(Sick(X))} g(Epid, Sick(x) = s) \right)^n$$

$$= \left(g'(Epid)\right)^n = g''(Epid) = g^\alpha(Epid)$$

$Epid$	g'
$false$	a
$true$	b

$Epid$	g''
$false$	a^n
$true$	b^n

$Epid$	g^α
$false$	$\frac{a^n}{a^n + b^n}$
$true$	$\frac{b^n}{a^n + b^n}$

$$\rightarrow \underbrace{\frac{1}{1 + \left(\frac{b}{a}\right)^n}}_{\text{Sigmoid function}} \xrightarrow{n \to \infty} \begin{cases} 0 & b > a \\ 1 & a > b \end{cases}$$

Growing Domain Sizes

Mittal et al. (2019)

- How to avoid extreme behaviour?

→ Adapt values in model dependent on domain size

- Approach for MLNs: Domain-size aware MLNs
  - Assume predicates $P_1, \dots, P_m$ occur in a first-order formula F
    - Count number of connections c_j for each predicate P_j given *new* domains
    - Build a connection vector $[c_1, \dots, c_m]$
    - Choose $\max_{c_i}[c_1, \dots, c_m]$ as scaling-down factor
      - Instead of $\max$, other functions possible
      - Works best if the values in $[c_1, \dots, c_m]$ do not vary that much
  - Given an MLN with a set of formulas F_i with weights w_i
    - Rescale each w_i with scaling-down factor s_i computed for F_i as $\frac{w_i}{s_i}$
- Analogous approach possible for parfactors

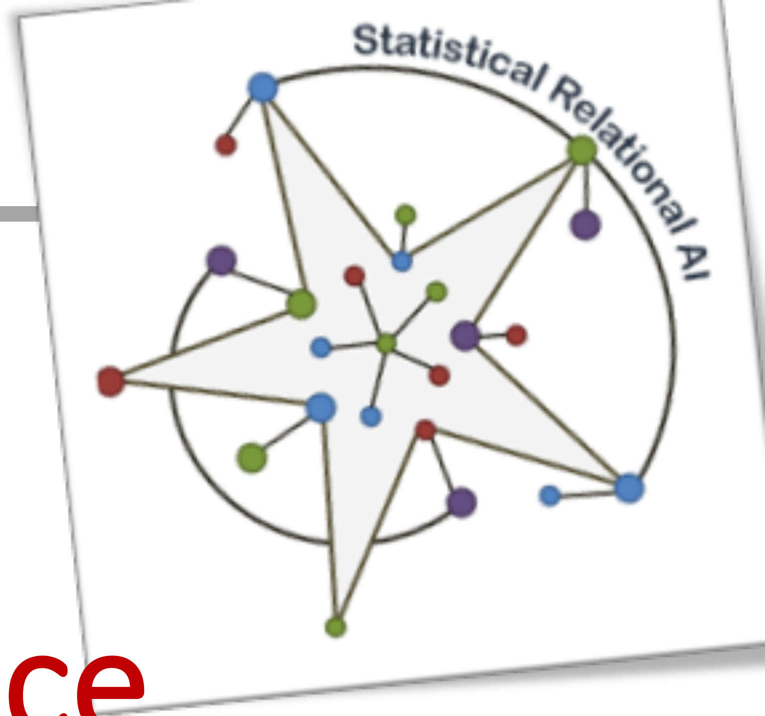

Keeping Inference Going over Time

Algorithms for More Robust Inference

Dynamic PRMs

- Marginal distribution query: $P(A_\pi^i \mid E_{0:t})$ w.r.t. the model:
  - 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?)
- MPE, MAP on temporal sequence

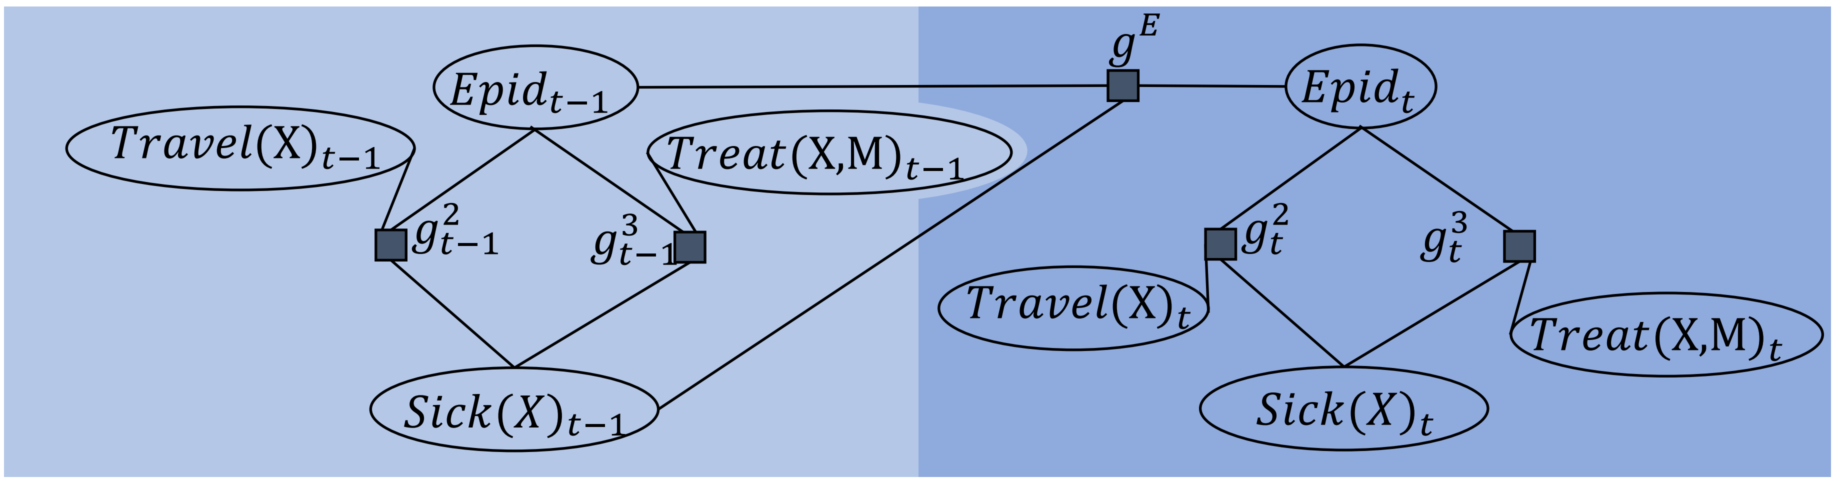

- **Define the interface for relational case (avoid groundings)**
- **Taming reasoning w.r.t. evidence over time (avoid creeping groundings)**

Reasoning over Time: Interfaces

- Main idea: Use temporal conditional independences to perform inference on smaller model
  - Normally only a subset of random variables influence next time step → interface variables
  - State description of interface variables from time slice $t-1$ suffice to perform inference on time slice t

→ Makes present independent from past / future

- Procedure
  - Build a helper structure of clusters (junction tree)
  - Proceed forward one time step at a time (forward message), using the same structure (vanilla junction tree)
  - Algorithms:
    - Propositional: Interface Algorithm (Murphy, 2002)
    - Lifted: Lifted Dynamic Junction Tree Algorithm (Gehrke et al, 2018)

Taming Reasoning

Gehrke et al. (2020)

- 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
- Vanilla junction trees for each time step
- Forward message carries over splits, leading to slowly grounding a model over time

Undoing Splits

Gehrke et al. (2020)

- Need to undo splits to

  keep reasoning polynomial w.r.t. domain sizes
- Where can splits be undone efficiently?
- How to undo splits?
- Is it reasonable to undo splits?
  - Effect of slight differences in evidence?
  - Impact of evidence vs. temporal behaviour of model?

Where Can Splits Be Undone Efficiently?

Gehrke et al. (2020)

- Evidence causes splits in a logical variable in the same way in all factors in a model
- LDJT always instantiates a vanilla junction tree
- Forward message carries over splits

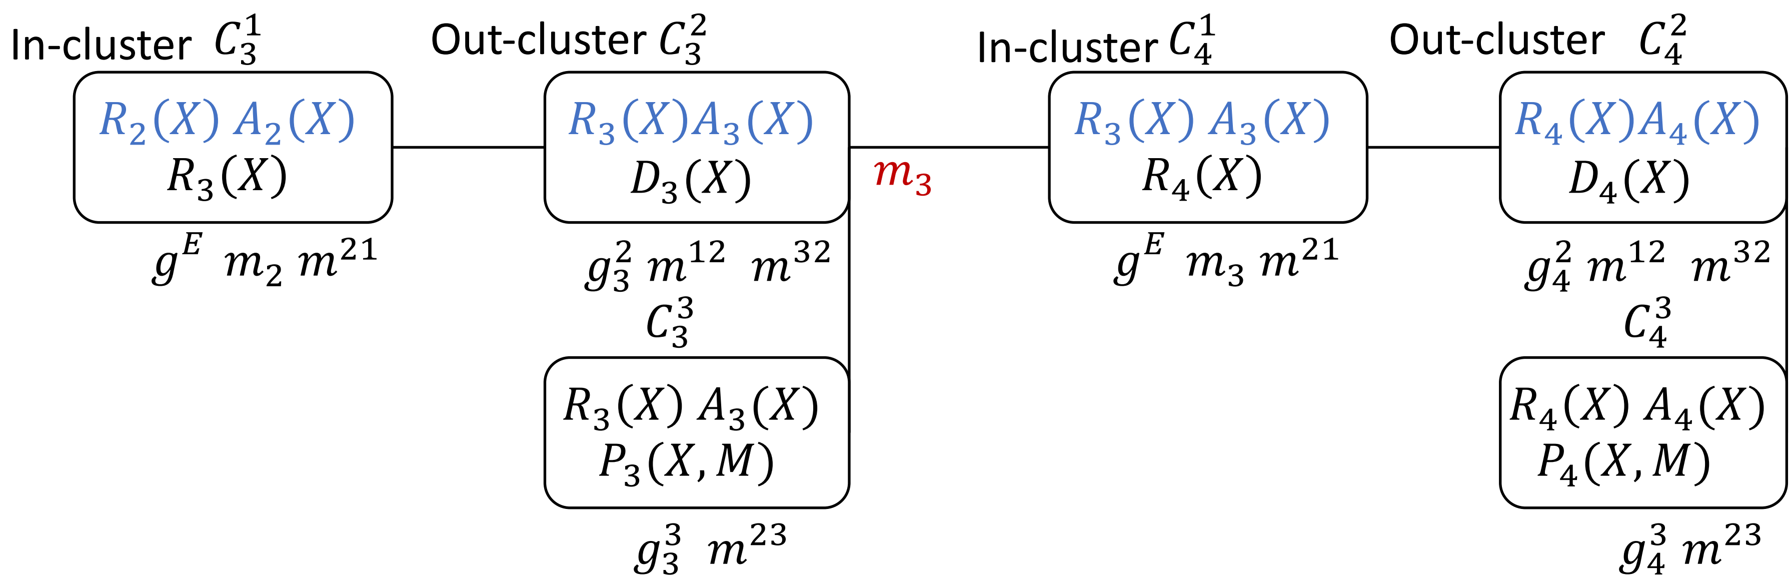

Identifying Similar Groups

Gehrke et al. (2020)

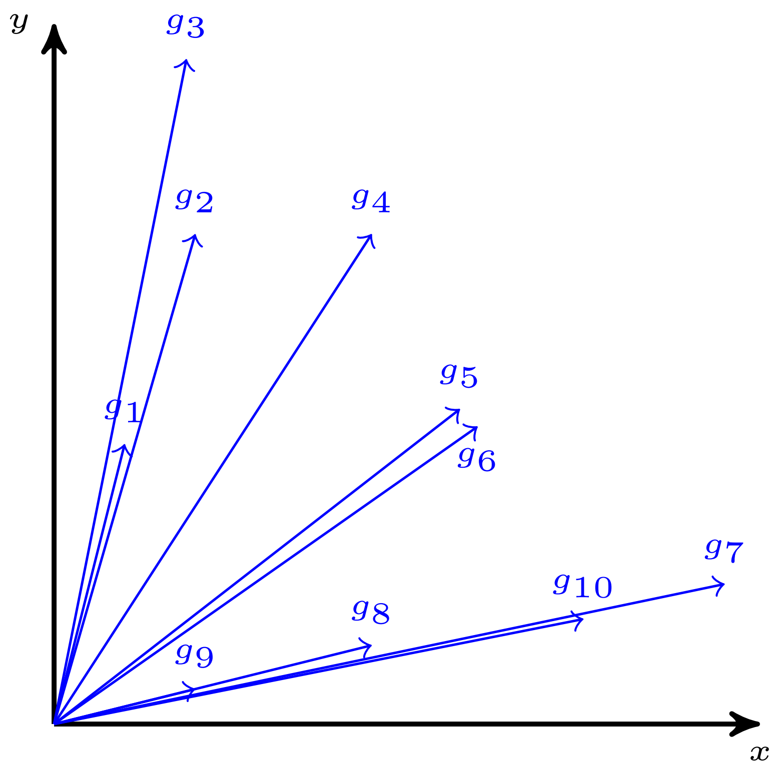

- Groups are equal if they have the same full joint distribution
- Full joint distribution computationally hard to get

→ Use parfactors as vector

→ If vectors of two groups point in same direction, they have the same full joint distribution

Cluster Groups

Gehrke et al. (2020)

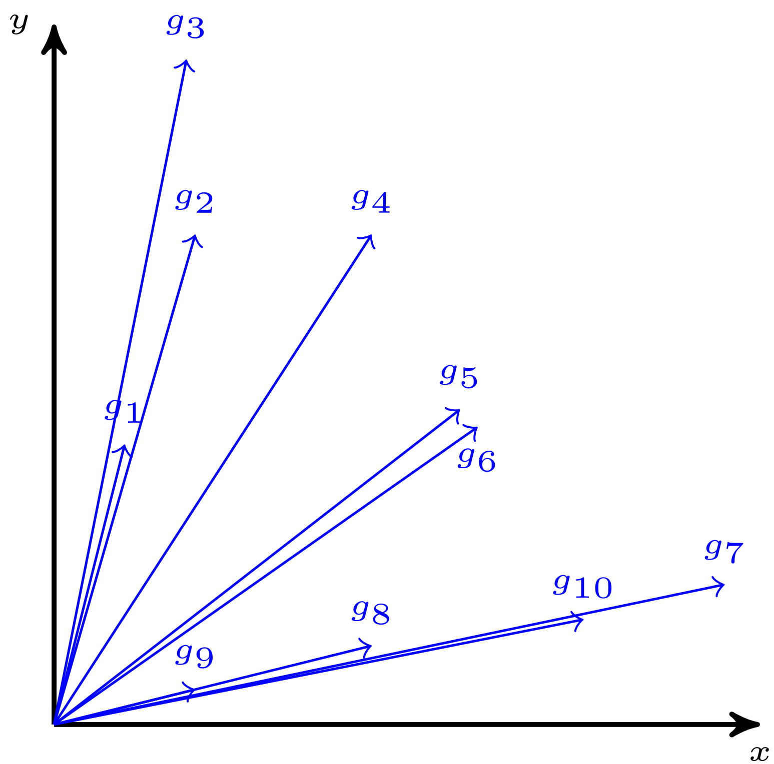

- Density-based clustering as unknown number of clusters
- Cosine similarity as distance function

Cluster Groups

Gehrke et al. (2020)

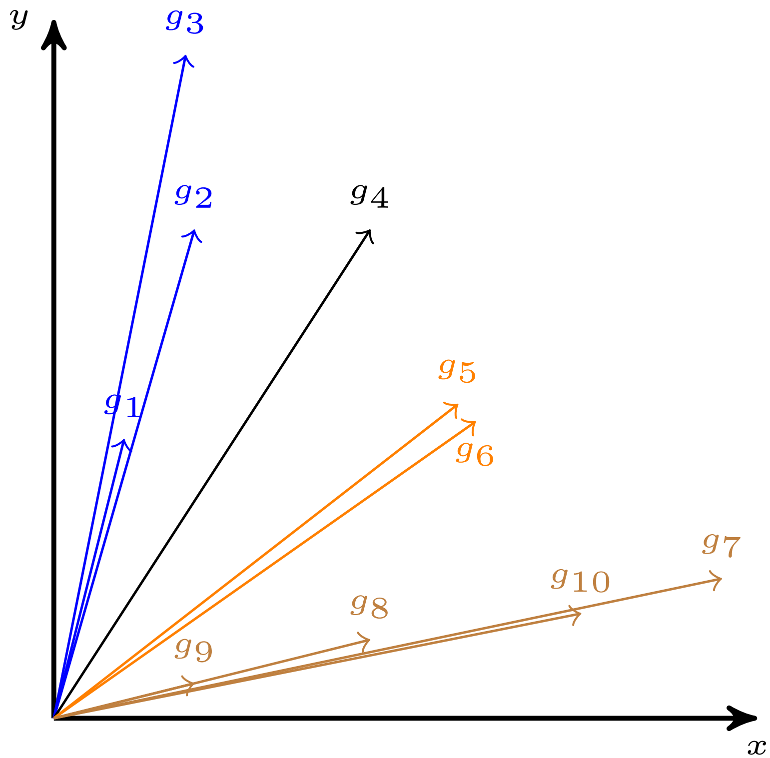

- Density-based clustering as unknown number of clusters
- Cosine similarity as distance function

Merge Clusters

Gehrke et al. (2020)

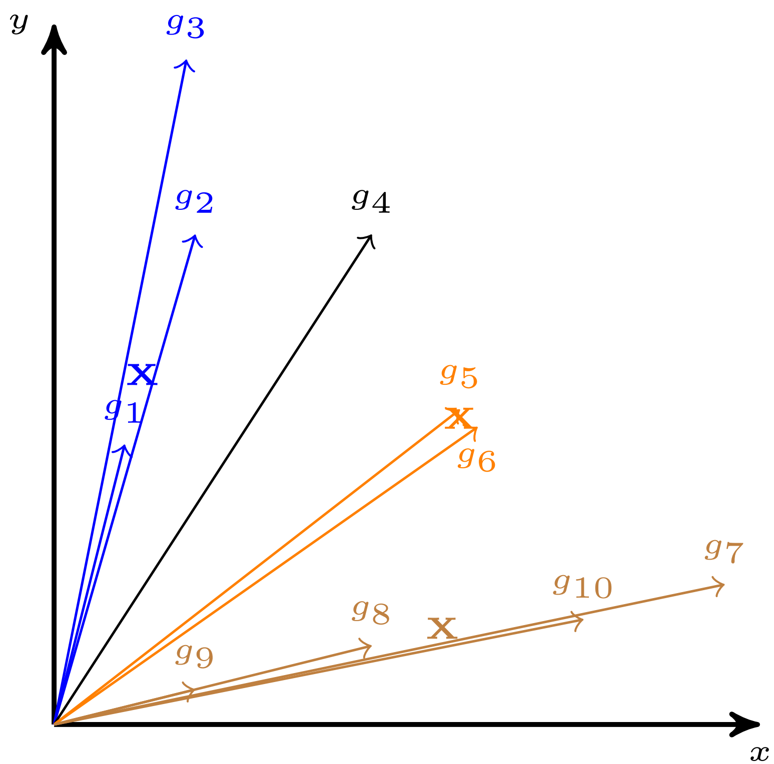

- Merge groups of cluster by calculating mean of cluster while accounting for groundings
- Replace old groups with merged group in temporal message

Is It Reasonable to Undo Splits?

Gehrke et al. (2020)

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

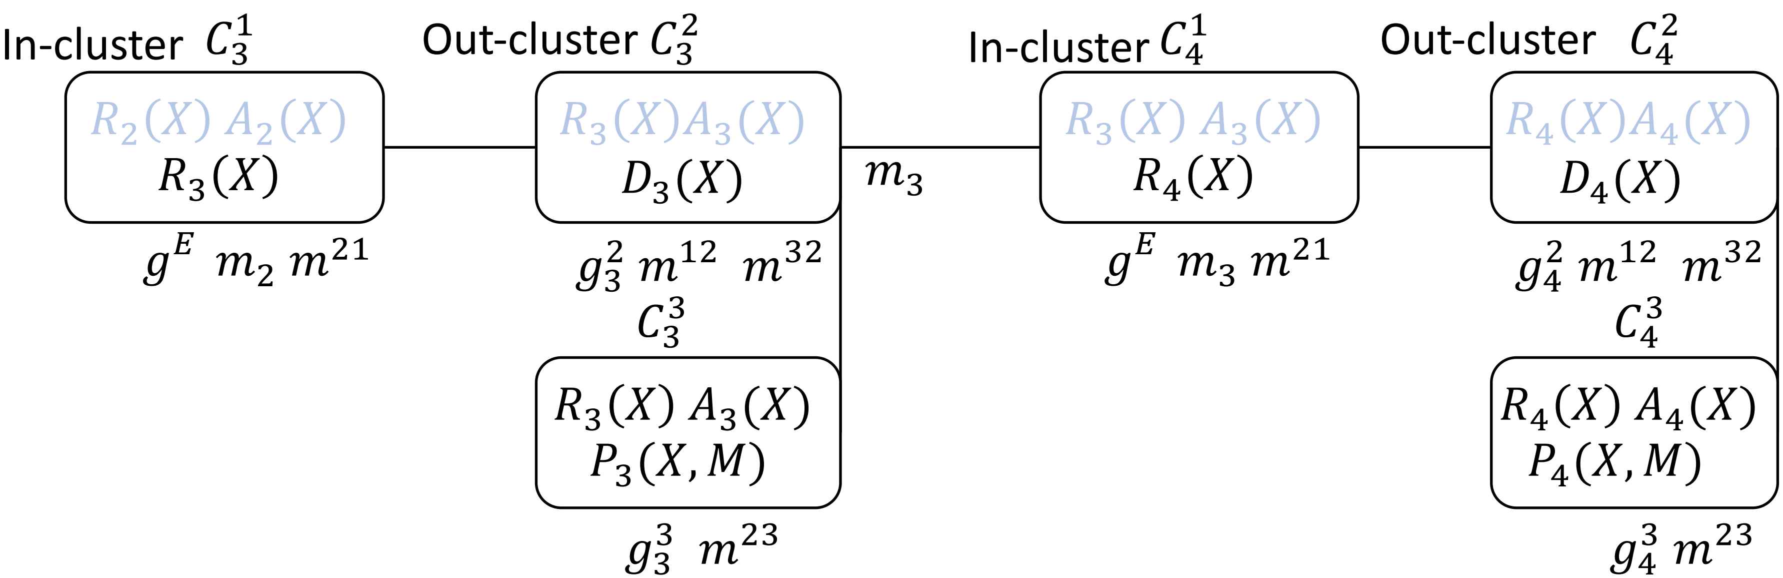

Taming Reasoning

Gehrke et al. (2020)

- Need to undo splits to

  keep reasoning polynomial w.r.t. domain sizes

- Where can splits be undone efficiently?
  - Undo splits in a forward message
- How to undo splits?
  - Find approximate symmetries
  - Merge based on groundings
- Is it reasonable to undo splits
  - Yes, due to the temporal model behaviour (indefinitely bounded error)

Results

Gehrke et al. (2020)

- DBSCAN for Clustering
- ANOVA for checking fitness of clusters

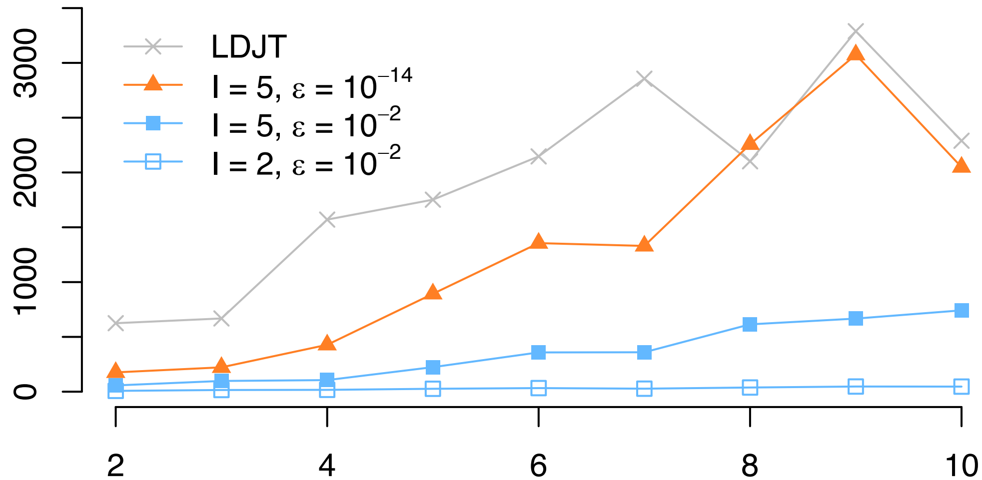

π	Max	Min	Average
0	0.0001537746121	0.0000000001720	0.0000191206488
2	0.0000000851654	0.0000000000001	0.0000000111949
4	0.0000000000478	0	0.0000000000068

WWU MÜNSTER

Agenda

- Probabilistic relational models (PRMs)
  - Application example
  - Semantics
  - Query answering / basic inference
- Algorithms for More Robust Inference
  - Cluster trees for efficient multi-query inference
  - Adaptive inference in cluster trees
  - Changing domains
  - Keeping inference going over time
- Summary

The Finish Line

- PRMs as a compact encoding of a full joint
  - Exploit symmetries
- Lifted inference
  - Use information about regular structures to speed up inference
- More robust infernce
  - Multi-query answering using junction tree as helper structure
  - Adaptive inference to get to the point of answering queries again fast
  - Changing domains with minimal effect
  - Keeping inference going over time

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

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 – Static 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.
- B (2020)
  Tanya Braun. Rescued from a Sea of Queries: Exact Inference in Probabilistic Relational Models. PhD Thesis, 2020.
- B and 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 and 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 and 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 and 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 and 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 and 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 and 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 and 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.

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 and 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.

Bibliography

- Van den Broeck and 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 and 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.