Dynamic StarAI

Dynamic Models and Statistical Relational AI

Tutorial at KI 2019

Tanya Braun, Marcel Gehrke, Ralf Möller
Universität zu Lübeck

UNIVERSITÄT ZU LÜBECK

https://www.ifis.uni-luebeck.de/index.php?id=612

Agenda: Dynamic Models and Statistical Relational AI

- Probabilistic relational models (PRMs) (Ralf)
  - Application example
  - Basic semantics, static vs. dynamic behavior
  - Exact multi-query answering
  - Variants: raw PRMs, MLNs, PSL, …
- Answering static queries (Tanya)
  - Lifted Junction Tree Algorithm (LJT)
- Answering continuous queries (Marcel)
  - Lifted Dynamic Junction Tree Algorithm (LDJT)
  - Relational interfaces
  - Taming reasoning w.r.t. lots of evidence over time
- Take home messages (Ralf)
  - LJT and LDJT research
    relevant for all variants of PRMs

Goal:
Overview of
central ideas

Application: Epidemics

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

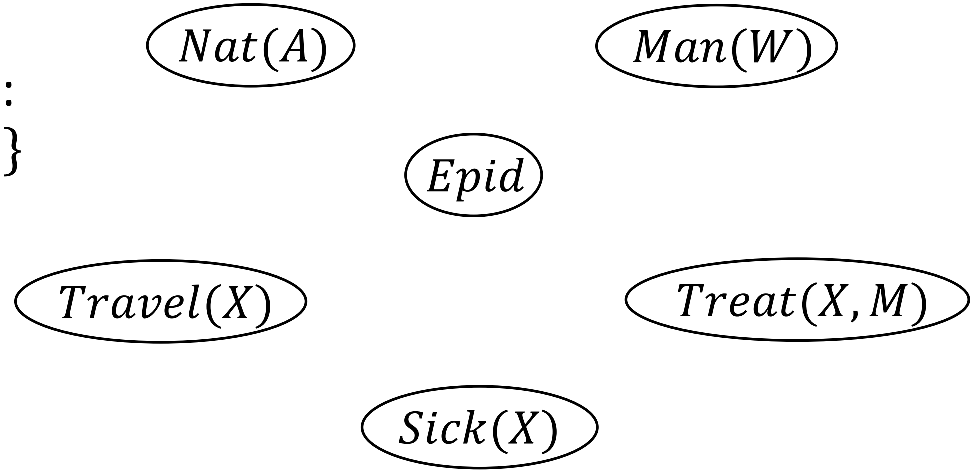

Encoding the Joint Distribution

- 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

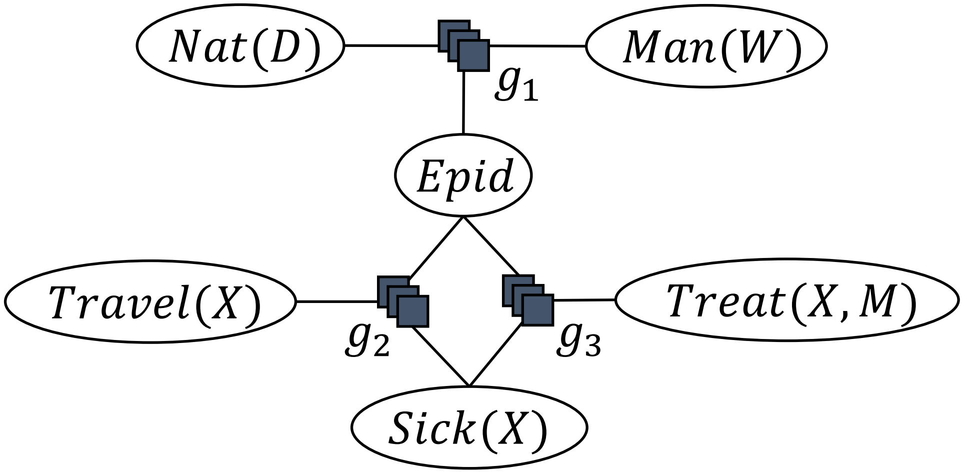

Nat(D) = natural disaster (D)
Man(W) = man-made disaster (W)

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

Factors

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

$Travel(eve)$	$Epid$	$Sick(eve)$	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

$Travel(X)$ $Epid$

$Travel(alice)$	Epi		
false	*fals*		
false	*fals*		
false	*tru*		
false	*tru*		
true	*fals*		
true	*fals*		
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

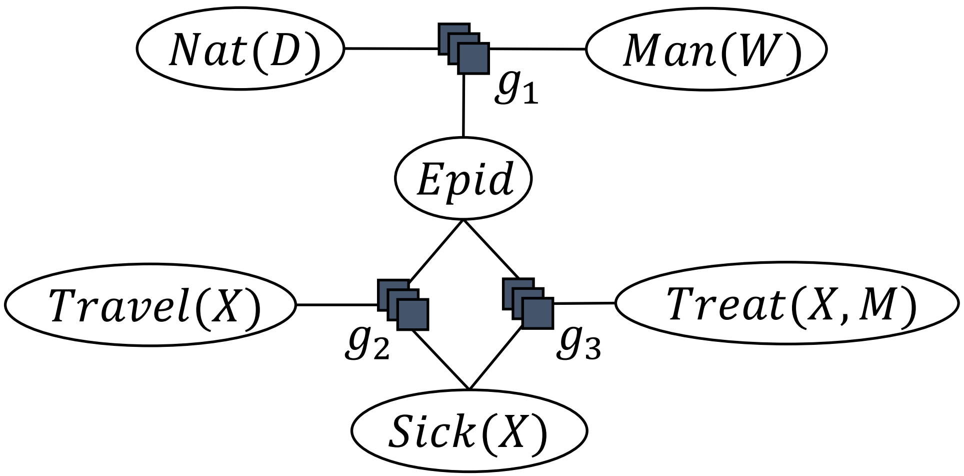

$$P_G = \frac{1}{Z} \prod_{f \in gr(G)} f$$

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

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

QA: Queries

- Marginal distribution
  - $P(Sick(eve))$
  - $P(Travel(eve,)\ Treat(eve, m_1))$
- Conditional distribution
  - $P(Sick(eve)|Epid)$
  - $P(Epid|Sick(eve) = true)$
- Most probable assignment
  - MPE
  - MAP

Avoid groundings!

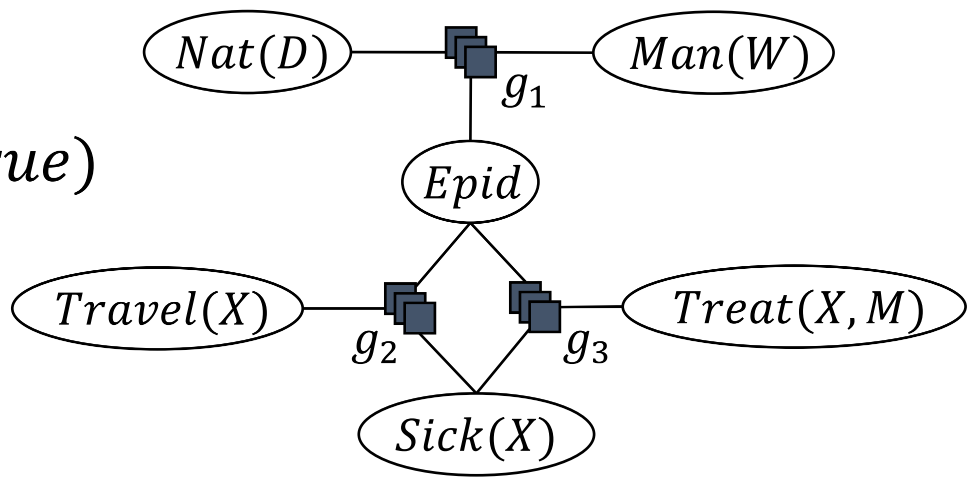

Problem: Many Queries

- 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
  - $Sick(X') = true$
  - $X' \in \{alice, eve\}$

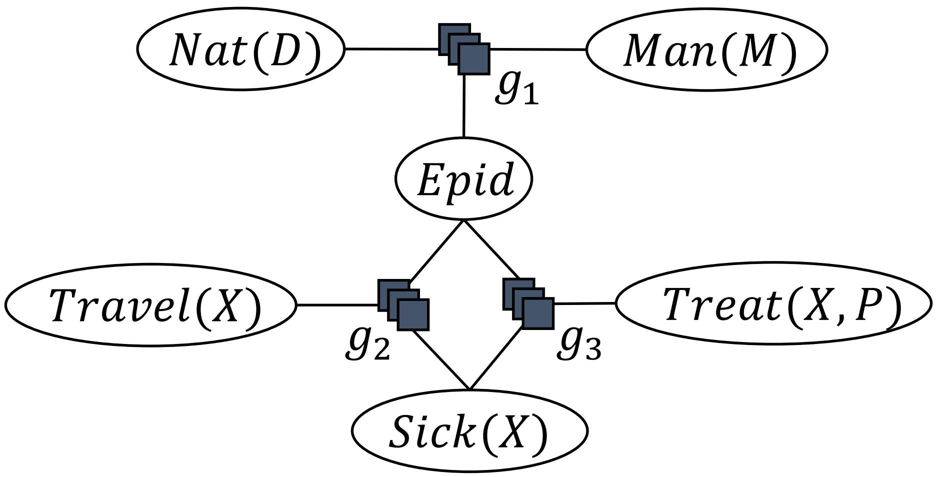

- **Challenges:**
  - **Do not start from scratch for every query**
  - **Support QA on subset of atoms**
  - **Avoid groundings**

Dynamic PRMs

- Marginal distribution query: $P\left(A_\pi^i \mid E_{0:t}\right)$ w.r.t. the model:
  - Hindsight: $\pi < t$ (Was there an epidemic $t - \pi$ days ago?)
  - Filtering: $\pi = t$ (Is there an currently an epidemic?)
  - Prediction: $\pi > t$ (Is there an epidemic in $\pi - t$ days?)
- MPE, MAP on temporal sequence

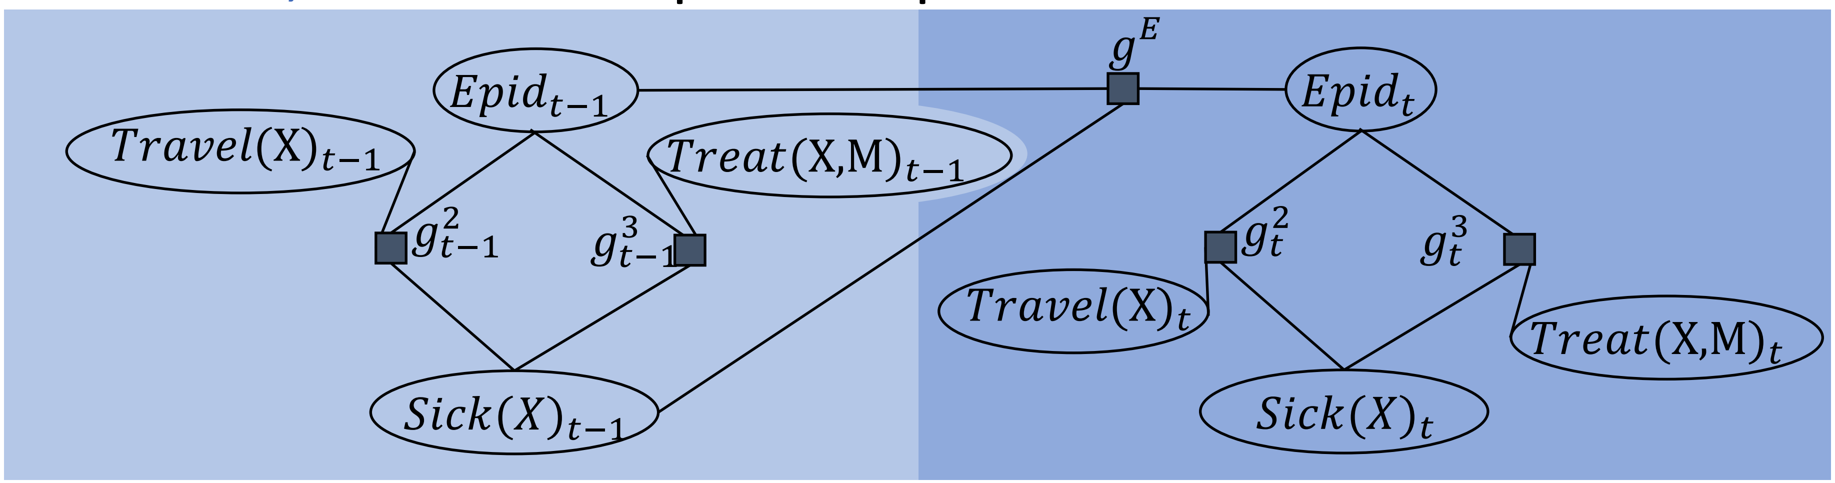

- **Define the interface for relational case (avoid groundings)**
- **Taming reasoning w.r.t. evidence over time (avoid creeping groundings)**

PRMs and variants

- Dynamic Probabilistic Relational Models (PRMs)
- Markov Logic Networks (MLNs)
- Probabilistic Soft Logic (PSL)
- Maximum Entropy Semantics (not covered here)

  [Thimm et al. 2010]

  Partial specification of discrete joint with "uniform completion"

Next: Answering static queries for PRMs

Bibliography *

- [Thimm et al. 10]
  Matthias Thimm, Marc Finthammer, Sebastian Loh, Gebriele Kern-Isberner, and Christoph Beierle. A System for Relational Probabilistic Reasoning on Maximum Entropy. In: Proc. FLAIRS-10: 116–121, 2010.

*PRMs are a true backbone of AI, and this tutorial emphasized only some central topics. We definitely have not cited all publications relevant to the whole field of PRMs here. We would like to thank all our colleagues for making their slides available (see some of the references to papers for respective credits). Slides or parts of it are almost always modified.