




Dynamic Probabilistic Relational Models

Lifted Exact Inference: Lifted Junction Tree Algorithm



Contents

1. Introduction

- StaRAI: Agent, context, motivation

2. Foundations

- Logic
- Probability theory
- Probabilistic graphical models (PGMs)

3. Probabilistic Relational Models (PRMs)

- Parfactor models, Markov logic networks
- Semantics, inference tasks

4. Exact Lifted Inference

- Lifted Variable Elimination
- Lifted Junction Tree Algorithm
- First-Order Knowledge Compilation

5. Lifted Sequential Models and Inference

- Parameterised models
- Semantics, inference tasks, algorithm

6. Lifted Decision Making

- Preferences, utility
- Decision-theoretic models, tasks, algorithm

7. Approximate Lifted Inference

8. Lifted Learning

- Parameter learning
- Relation learning
- Approximating symmetries

Outline: 4. Lifted Inference


Exact Inference

- i. Lifted Variable Elimination for Parfactor Models
 - Idea, operators, algorithm, complexity
- ii. Lifted Junction Tree Algorithm
 - Idea, helper structure: junction tree, algorithm
- iii. First-order Knowledge Compilation for MLNs
 - Idea, helper structure: circuit, algorithm

Problem: Many Queries

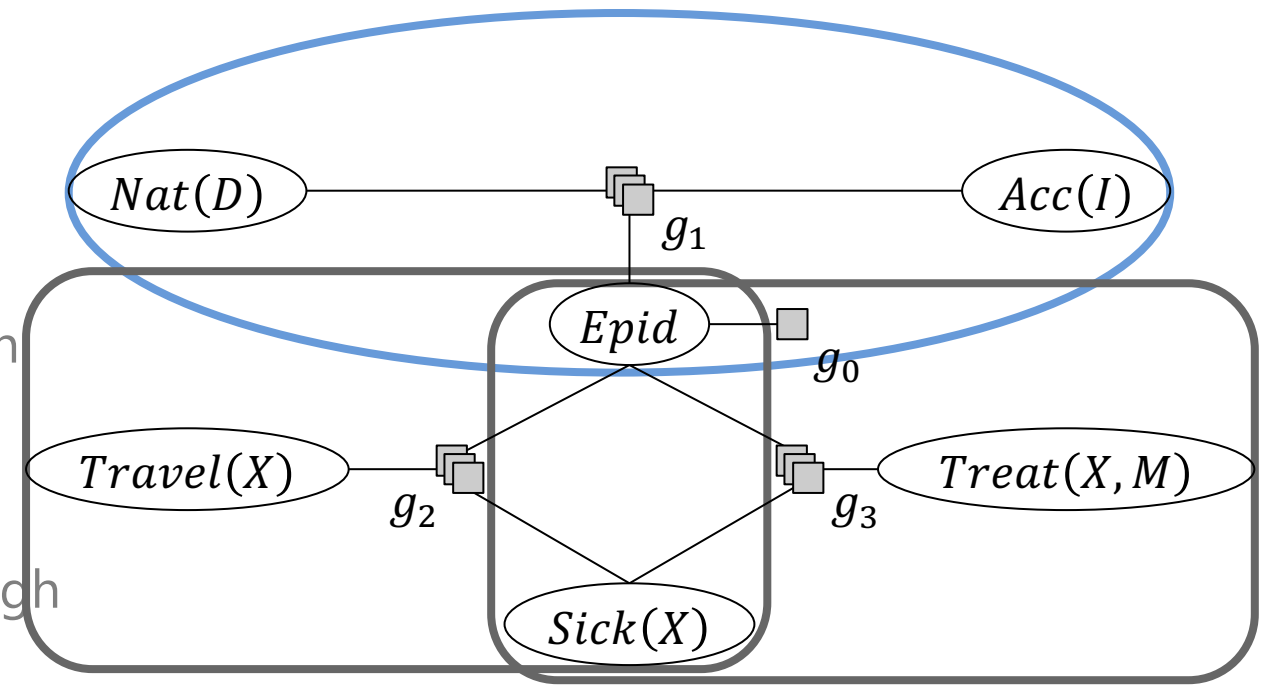
- Set of queries
 - $P(\text{Travel}(\text{eve}))$
 - $P(\text{Sick}(\text{bob}))$
 - $P(\text{Treat}(\text{eve}, m_1))$
 - $P(\text{Epid})$
 - $P(\text{Nat}(\text{flood}))$
 - $P(\text{Acc}(\text{chem}))$
 - Combinations of variables
- Under evidence
 - $\text{Sick}(X') = \text{true}$
 - $X' \in \{\text{alice}, \text{eve}\}$
- LVE restarts with initial model for each query

Build a helper structure to precompute parts



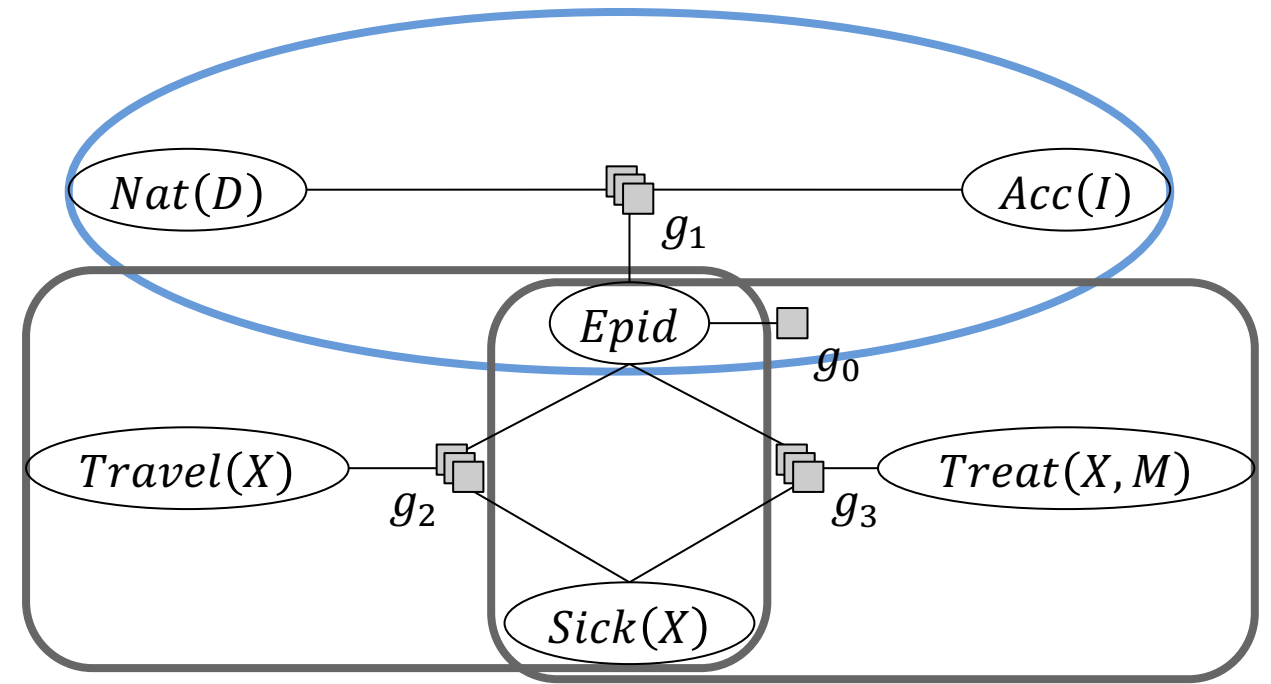
Clustering of Models

- Idea: Find subsets (clusters) of PRVs that are “enough” for certain queries
 - E.g.,
 - For queries about instances of *Nat(D)*, *Acc(I)*, *Epid*
 - *Nat(D)*, *Acc(I)*, *Epid* enough
 - For queries about instances of *Travel(X)*, *Sick(X)*, *Epid*
 - *Travel(X)*, *Sick(X)*, *Epid* enough
 - For queries about instances of *Treat(X, M)*, *Sick(X)*, *Epid*
 - *Treat(X, M)*, *Sick(X)*, *Epid* enough



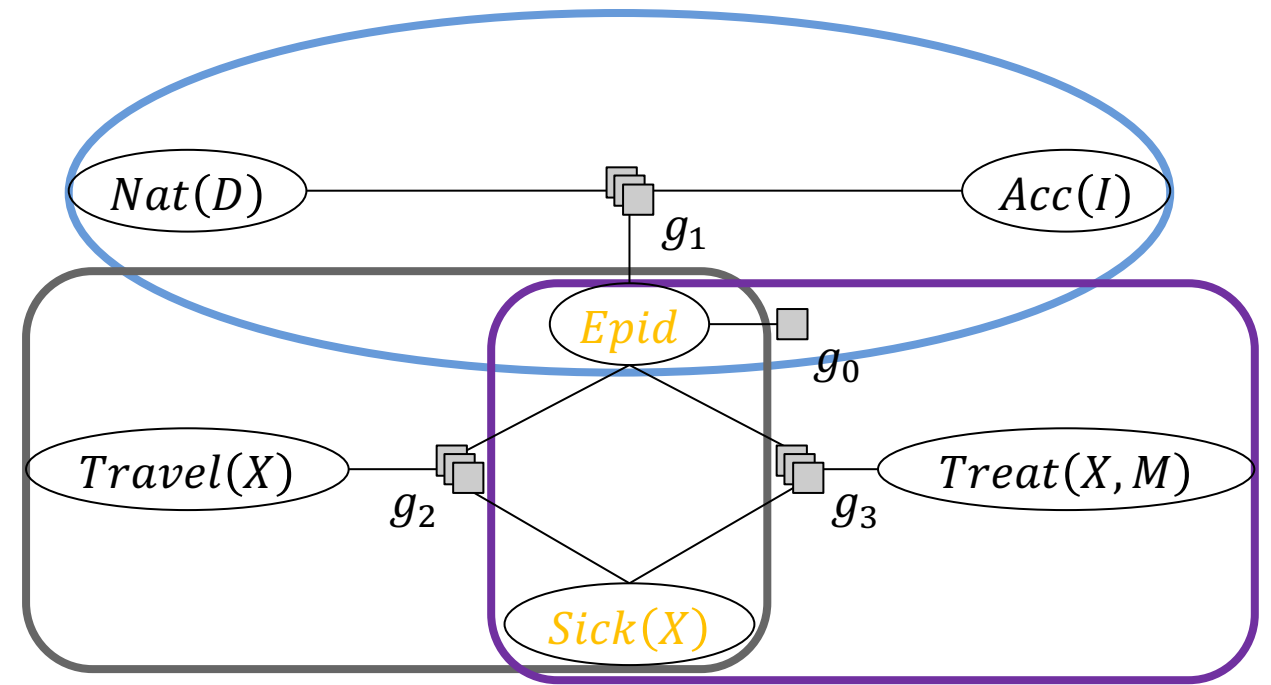
Clustering of Models

- But: If only parfactors used that contain the PRVs of a cluster, information stored in all other parfactors ignored
 - E.g.,
 - $Nat(D), Acc(I), Epid: g_1$
→ misses g_2, g_3
 - $Travel(X), Sick(X), Epid: g_2$
→ misses g_1, g_3
 - $Treat(X, M), Sick(X), Epid: g_3$
→ misses g_1, g_2
 - Whatever we do with $g_0...$
- Only correct if clusters are *independent* from each other
 - How can we achieve independence?



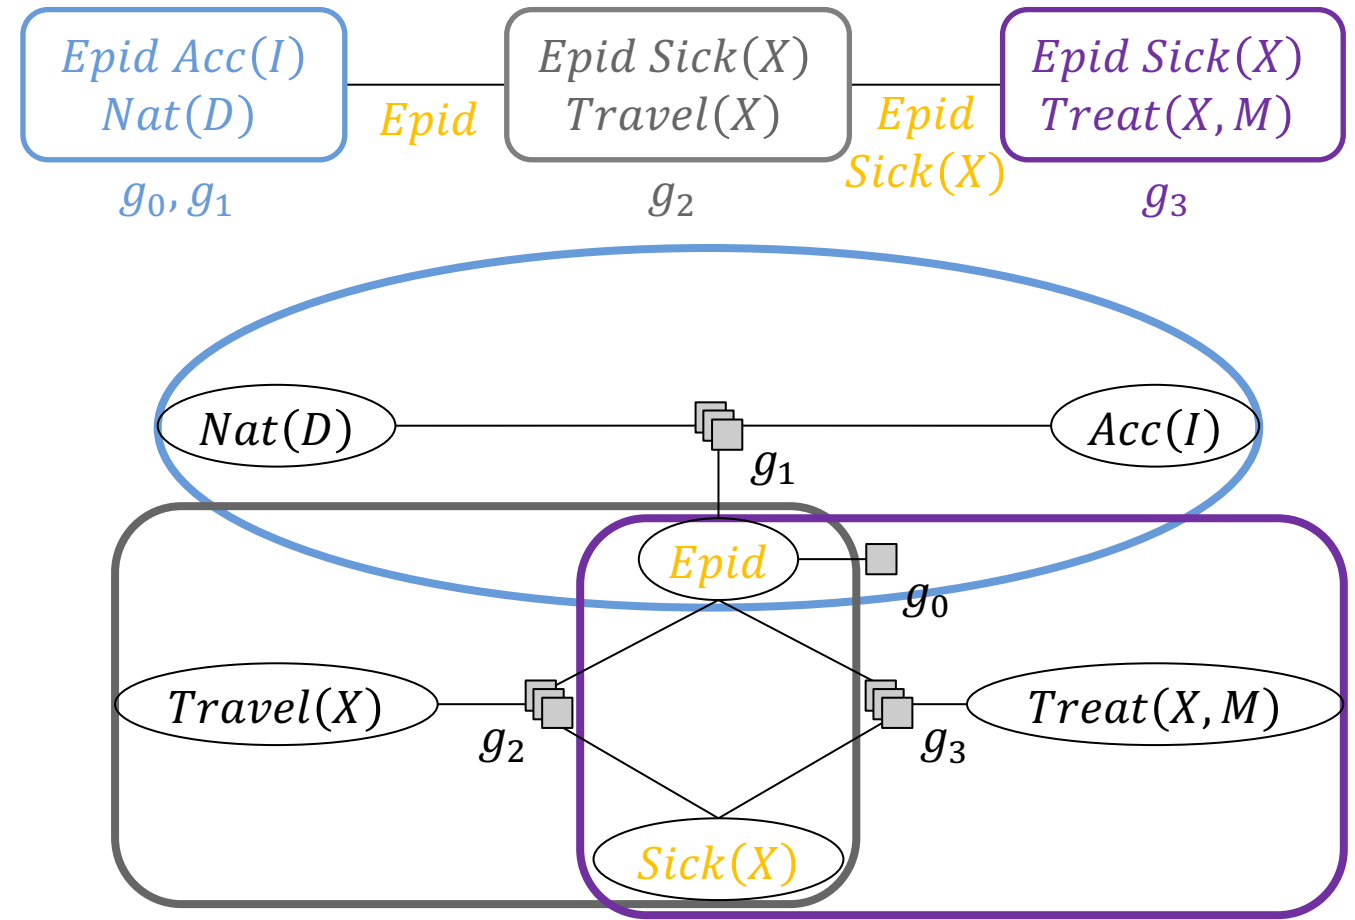
Clustering of Models

- Factorised models encode independences:
 - Any two subsets of variables are conditionally independent given a *separating* subset S
 - *Separating* subset S : All paths from one subset to the other run through S
 - Also known as *global Markov property*
 - E.g.,
 - $Nat(D), Acc(I), Epid: g_1$
→ independent of the rest given $Epid$
 - $Travel(X), Sick(X), Epid: g_2$
→ independent of the rest given $Epid, Sick(X)$
 - $Treat(X, M), Sick(X), Epid: g_3$
→ independent of the rest given $Epid, Sick(X)$




Clustering of Models

- Put clusters and their separators into a graph structure where
 - Nodes are clusters with parafactors assigned containing the cluster PRVs (*local model*)
 - Edges are labelled with the *separator* (separating subset) between neighbouring nodes
 - If two nodes contain the same PRV, every node on the path between the two nodes contain the PRV (*running intersection property*)



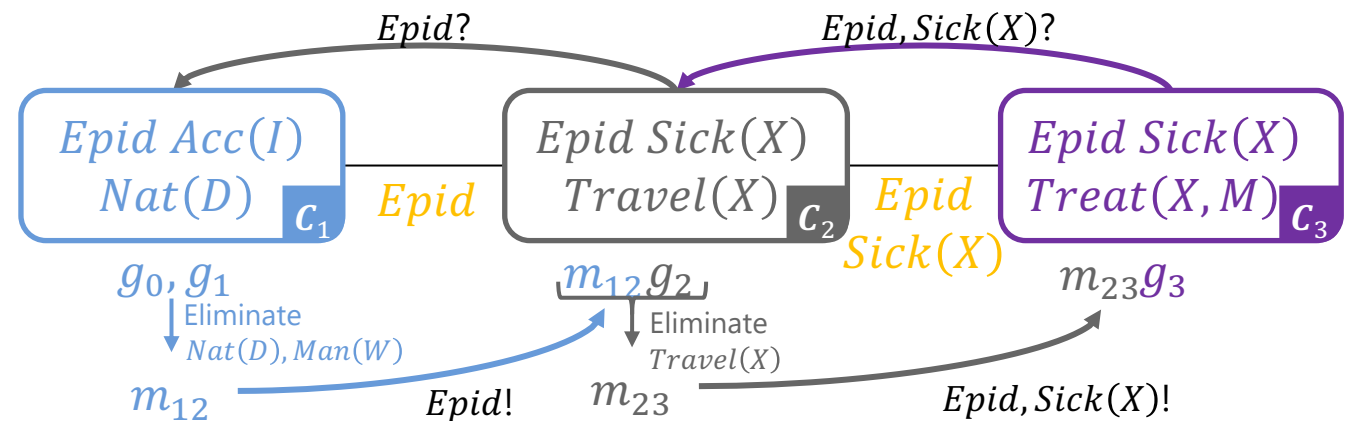
Clustering of Models

- Next: Make clusters actually independent of each other
 - Each cluster i asks its neighbours $j \in nbs(i)$ for information about the separator \mathcal{S}_{ij} between them
 - Other clusters have to collect all the information from the model that lies behind the separator on its part, eliminate the non-separator PRVs from that information using LVE, and send the result in a message m_{ji} , i.e., a set of parfactors, back
 - Having the information on the separators to all neighbours makes a cluster independent from its neighbours and therefore all other parts of the model
 - Ensures that each cluster of PRVs has all model information needed available for query answering on instances of its cluster PRVs



Clustering of Models

- Next: Make clusters actually independent of each other
 - E.g., $C_3: g_3 \rightarrow$ independent of the rest given $Epid, Sick(X)$
 - Asks neighbour C_2 for information on $Epid, Sick(X)$
 - C_2 asks neighbour C_1 for information on $Epid$
 - » C_1 sends information on $Epid$ in a message m_{12}
 - » Eliminates $Nat(D), Acc(I)$ from g_0, g_1 for m_{12}
 - C_2 sends information on $Epid, Sick(X)$ to C_3 in a message m_{23}
 - » Eliminates $Travel(X)$ from g_2 and m_{12} for m_{23}
 - With m_{23} , C_3 is independent from its neighbour C_2 and therefore also from C_1
 - As C_2 is independent given m_{12} from C_1




The same has to be done for C_2 and C_1

Clustering of Models

- With each cluster i independent of the rest, each i can answer queries about instances of its PRVs based on its local model and the messages received
 - Query terms: grounded instances or parameterised versions of its PRVs
 - Conjunctive queries if terms only concern the cluster PRVs
 - E.g., $C_3: g_3 \rightarrow$ independent of the rest given $Epid, Sick(X)$
 - Based on g_3 and m_{23} , C_3 can answer queries about $Epid, Sick(X), Treat(X, M)$ such as

$P(Sick(X)),$
 $P(Treat(eve, m_2)),$
 $P(Epid, Sick(alice))$

- Cannot answer any queries about $Nat(D), Acc(I), Travel(X)$ but C_1 and C_2 , respectively, can

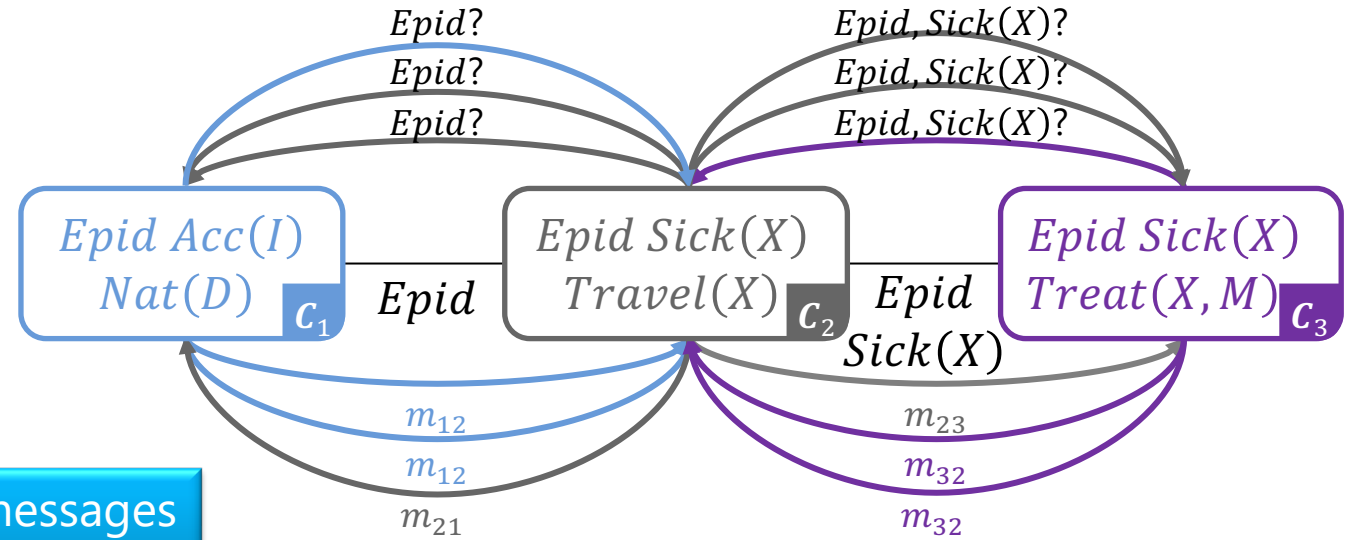


Clustering of Models

- Problem left: If each cluster asks for information on separators, some messages are sent multiple times

– E.g.,


- C_3 asks C_2 , which asks C_1
 - Messages calculated and sent: m_{12}, m_{23}
- C_2 asks C_1 and C_3
 - Messages calculated and sent: m_{12}, m_{32}
- C_1 asks C_2 , which asks C_3
 - Messages calculated and sent: m_{32}, m_{21}



Organise in way that messages are calculated only once

Clustering of Models

- Use dynamic programming to organise the order of asking or rather sending information in messages:
 - If a node i has received all information from neighbours but one, j , node i sends a message with its information on the separator \mathcal{S}_{ij} to j
 - If a node i has received all messages, then it sends messages to all neighbours j that have not received a message yet
- When computing the message, i takes into consideration
 - its local model G_i as well as
 - the messages m_{ki} received from all other neighbours k but the receiving neighbour j




Clustering of Models

Graph structured called (first-order) junction tree and algorithm called (lifted) junction tree algorithm

- Observations:

- If a node i has received all information from neighbours but one, j , node i sends a message with its information on the separator S_{ij} to j
 - Trivially true at leaf nodes (*periphery*), can start sending immediately to its only neighbour (in *parallel*!)
 - From periphery inbound, new nodes trigger this first condition
- If a node i has received all messages, then it sends messages to all neighbours j that have not yet received a message

- As messages are sent further inwards, a first node at the centre will have received all messages and will start sending messages outbound, leading to new nodes triggering this second condition




Two passes from periphery to centre and back suffice to distribute all information and make the clusters independent from each other*

* Shown by Steffen L. Lauritzen and David J. Spiegelhalter: Local Computations with Probabilities on Graphical Structures and Their Application to Expert Systems. In: *Journal of the Royal Statistical Society. Series B: Methodological*, 1988.


Lifted Junction Tree Algorithm (LJT)

- Inputs
 - Model G
 - Evidence e as evidence parfactors
 - Set of query terms $\{Q_i\}_{i=1}^m$
 - Queries on G : $P(Q_i | e)$, $i \in \{1, \dots, m\}$
- LJT consists of four steps
 - Build FO jtree J for model G
 - Enter evidence e in J
 - Pass messages in J
 - Answer queries $\{Q_i\}_{i=1}^m$




Evidence:
 $sick(eve)$

Queries:
 $\{\{Epid\}, \{Travel(eve), Treat(eve, m_1)\}\}$



First-order Jtree (FO Jtree)

- As seen on the earlier slides
 - Acyclic graph
 - Nodes contain PRVs, which form clusters
 - Edges are based on the separators between the clusters
 - Nodes have parfactors assigned
- Next slides:
 - Formal definition
 - Construction
 - Get an *acyclic* structure with valid *separators* and each parfactor of a model assigned to a *local model*



Parameterised Clusters

- Node of an FO jtree: Set of PRVs called parameterised cluster (**parcluster**)
- Let \mathbf{X} be a set of logical variables, \mathbf{A} a set of PRVs with $lv(\mathbf{A}) \subseteq \mathbf{X}$, and (\mathcal{X}, C_x) a constraint on \mathbf{X}
- Then, a parcluster \mathbf{C} is given by

$$\forall x \in C_x : \mathbf{A}|_{(\mathcal{X}, C_x)}$$


- $\mathbf{A}|_{(\mathcal{X}, C_x)}$ for short
 - Again, (\mathcal{X}, C_x) can be omitted if T constraint encoded
- Depicted as a round shape containing \mathbf{A} or just \mathbf{A}
 - Again, constraint usually not depicted
- E.g., parcluster \mathbf{C}_2

$$\forall x \in dom(X) : \{Epid, Sick(x), Travel(x)\}|_{(\mathcal{X}, dom(X))}$$

$$= \{Epid, Sick(X), Travel(X)\}|_{(\mathcal{X}, D(X))}$$



$$= \{Epid, Sick(X), Travel(X)\}$$

Epid Sick(X)
Travel(X)



FO Jtree

- An FO jtree J for a model G is a cycle-free graph (V, E)
 - Set of nodes $V \subseteq 2^{rv(G)}$
 - I.e., nodes are sets of PRVs (parclusters)
 - $2^{rv(G)}$ denotes the power set of $rv(G)$
 - Set of edges $E \subseteq \{\{i, j\} \mid i, j \in V, i \neq j\}$,
 - Has to be cycle free, which includes no self-loops
 - E.g., as depicted on the left
 - But at this point in the definition, could be any subsets of PRVs



FO Jtree

- An FO jtree J for a model G is a cycle-free graph (V, E)

- Has to satisfy three properties:

- $\forall \mathcal{C} \in V : \mathcal{C} \subseteq rv(G)$

- Every parcluster consists of PRVs from G

- $\forall g \in G : \exists \mathcal{C} \in V : rv(g) \subseteq \mathcal{C}$


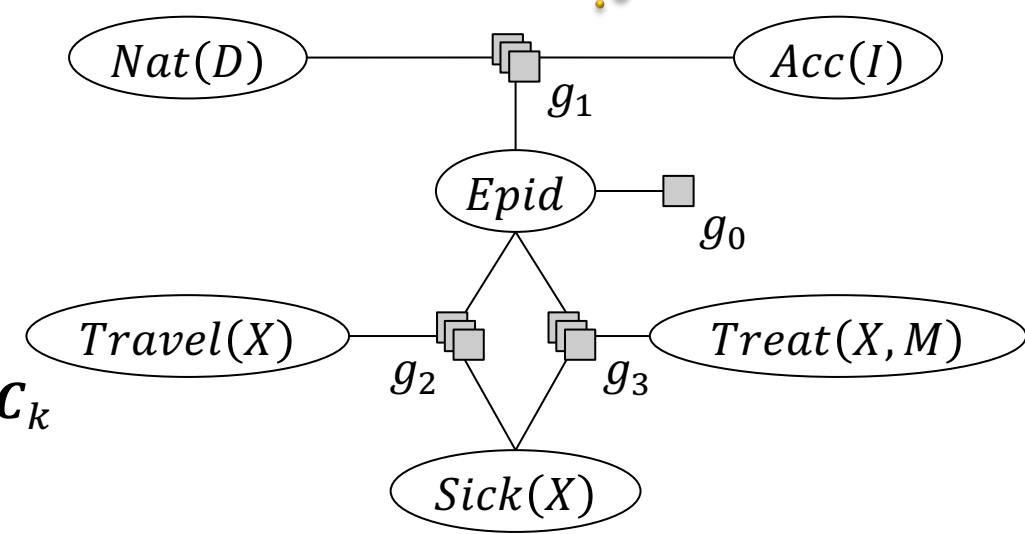
- Arguments of every parfactor in G occur in a parcluster

- If $\exists A \in rv(G) : A \in \mathcal{C}_i \wedge A \in \mathcal{C}_j$ with $\mathcal{C}_i, \mathcal{C}_j \in V$, then $\forall \mathcal{C}_k \in V$ on the path between $\mathcal{C}_i, \mathcal{C}_j : A \in \mathcal{C}_k$ (running intersection property)

- If a PRV occurs in two parclusters, it also occurs in every parcluster on the path between them



- E.g., as depicted on the left

Other valid FO jtrees to build?



FO Jtree

- An FO jtree J for a model G is a cycle-free graph (V, E)
 - Is **minimal** if by removing a PRV from a parcluster, the FO jtree ceases to be an FO jtree
 - I.e., no longer fulfils at least one property
 - E.g., depicted on the left
 - Cannot remove any PRV from any parcluster
 - Otherwise, a parfactor would no longer have its arguments in one parcluster



FO Jtree

- An FO jtree J for a model G is a cycle-free graph (V, E)
 - Set S_{ij} called **separator** of edge $\{i, j\} \in E$, defined by

$$S_{ij} = C_i \cap C_j$$

- Term $nbs(i)$ refers to the neighbours of C_i , defined by



$$nbs(i) = \{j \mid \{i, j\} \in E\}$$

- Each C_i has a **local model** G_i and

$$\forall g \in G_i : rv(g) \subseteq C_i$$

- Local models G_i partition G , i.e.,

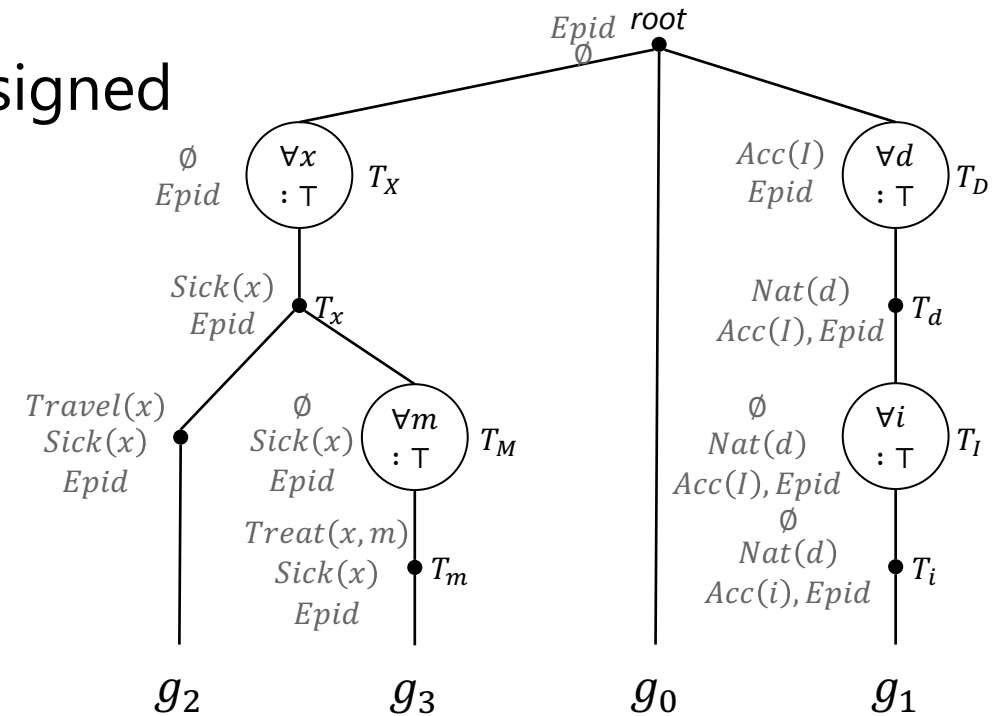
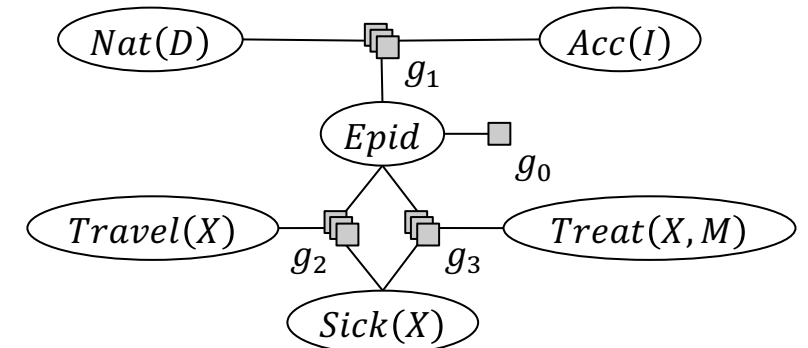

$$G = \bigcup_{i \in V} G_i, \forall i, j \in V, i \neq j : G_i \cap G_j = \emptyset, G_i \neq \emptyset$$



Construction

- Where do we get the FO jtree from s.t. the jtree
 - is acyclic
 - fulfils the three FO jtree properties
 - has the model parfactors automatically assigned to fitting parclusters?

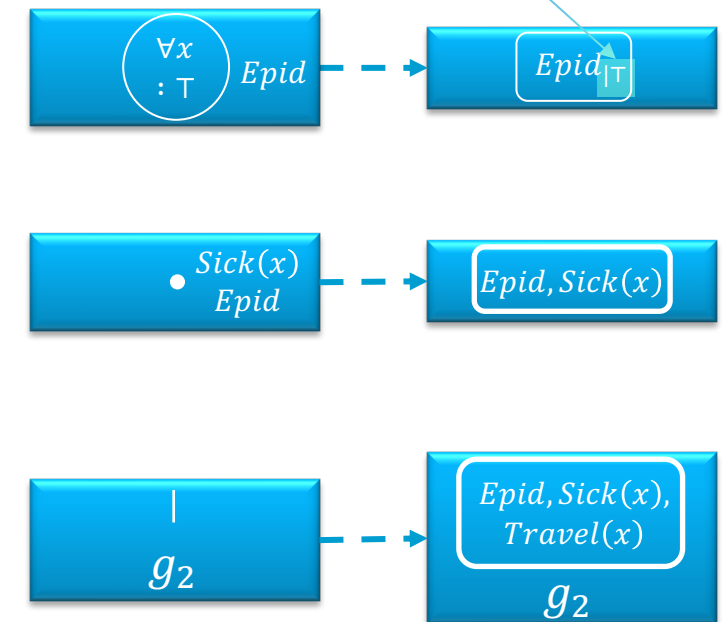
→ Clusters of an FO dtree
 + undirected dtree edges
 + minimisation
 = FO jtree



Clusters → Parclusters

- Given an FO dtree T for a model G with clusters for each node
- Given a cluster $\{A_1, \dots, A_n\}$ of a DPG node (X, x, C)
 - Resulting parcluster $\mathcal{C}_j = \{A_1, \dots, A_n\}_{|C}$
 - Local model $G_j = \emptyset$
- Given a cluster $\{A_1, \dots, A_n\}$ of a VE node
 - Resulting parcluster $\mathcal{C}_j = \{A_1, \dots, A_n\}_{|T}$
 - Local model $G_j = \emptyset$
- Given a cluster $\{A_1, \dots, A_n\}$ from a leaf node with parfactor g_i
 - Resulting parcluster $\mathcal{C}_j = \{A_1, \dots, A_n\}_{|T}$
 - Local model $G_j = \{g_i\}$

Let's carry the constraint around for a bit to make it explicit

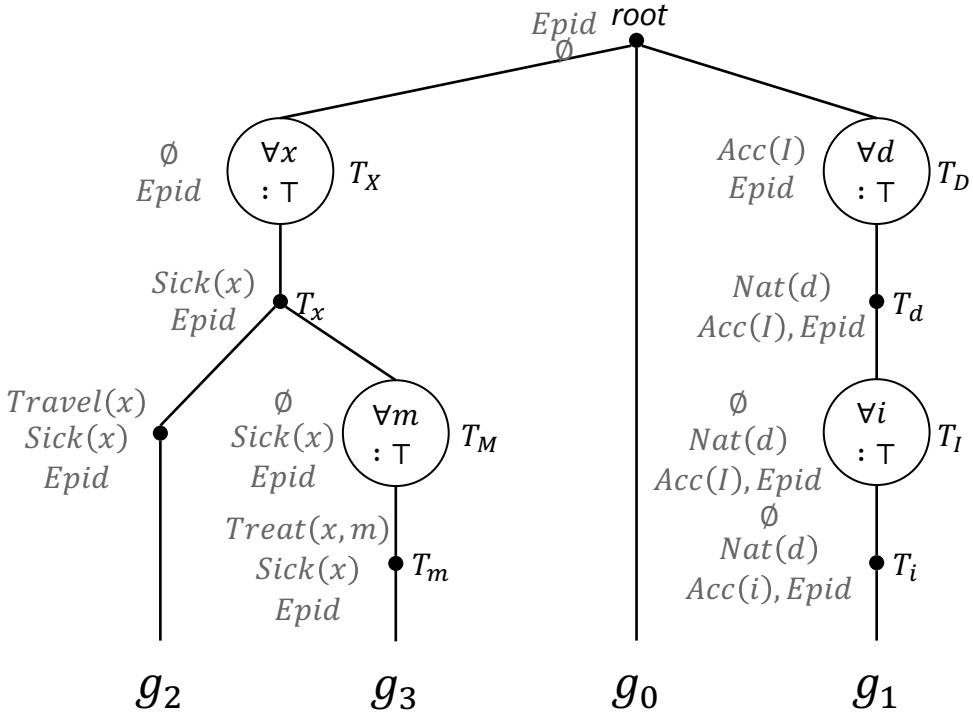
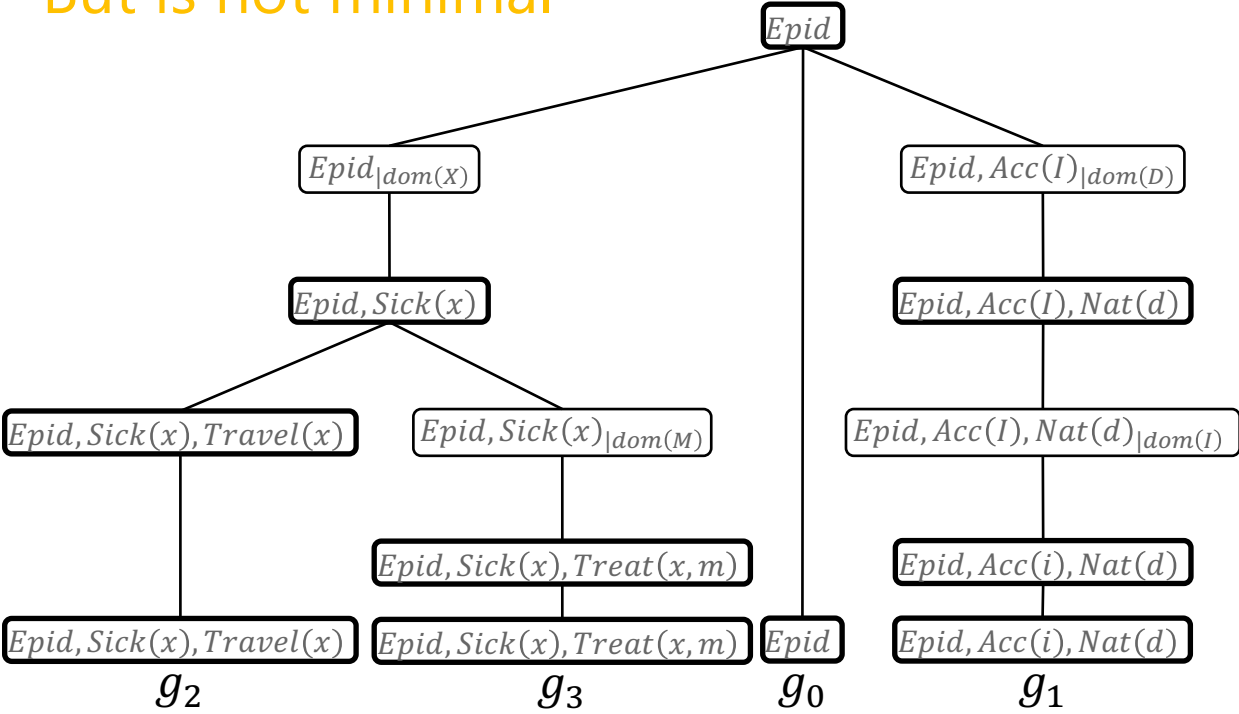



FO Dtree \rightarrow FO Jtree

- Forming an FO jtree J from an FO dtree T of a model G
- Nodes of J
 - Parclusters resulting from clusters of T as shown on previous slide
 - Each parcluster has a source node in T
- Edges of J
 - Add an edge between two parclusters whenever there is an edge between the source nodes of the two parclusters in T

FO Dtree \rightarrow FO Jtree

- Result after transformation
 - Fulfills the three jtree properties
 - But is not minimal



FO Dtree \rightarrow FO Jtree

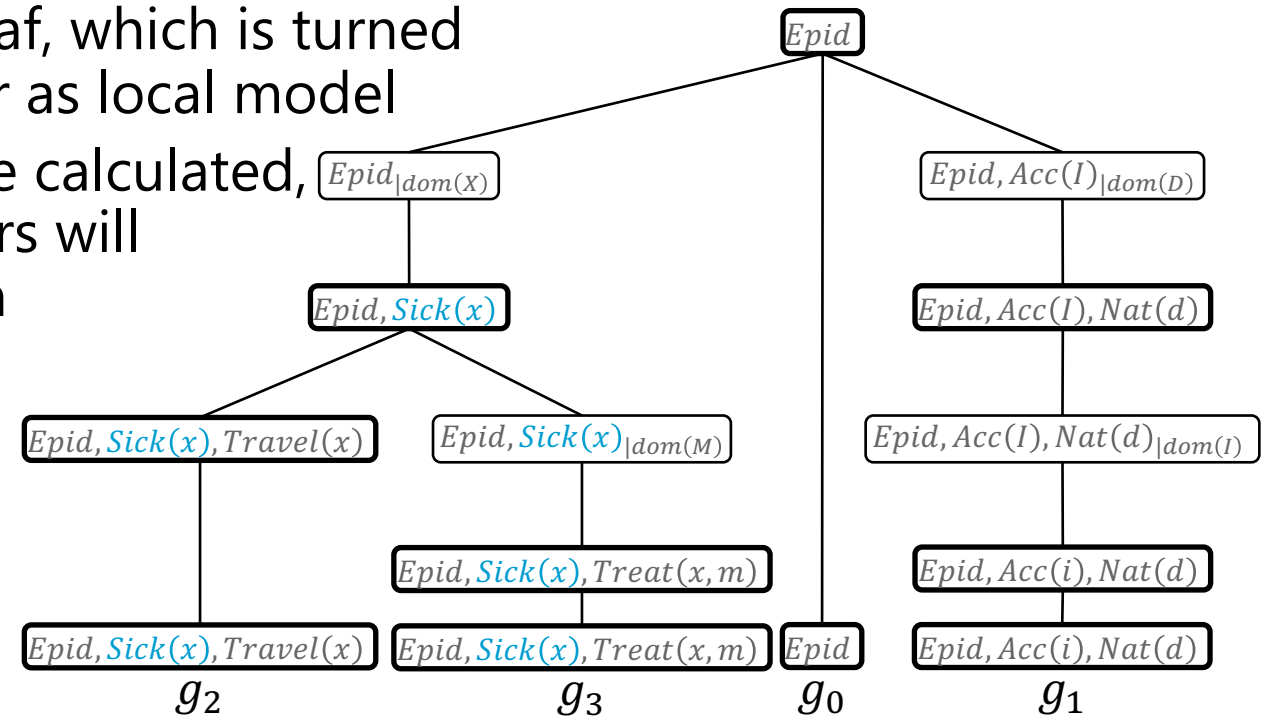
- Transformation result fulfils the three jtree properties

- Properties hold by construction of the FO dtree

1. Parclusters can only contain model PRVs
2. Each parfactor occurs at a dtree leaf, which is turned into a parcluster with the parfactor as local model
3. Based on how cutset & context are calculated, a PRV that occurs in two parclusters will occur in all parclusters on the path between them*

- E.g., *Sick(X)*

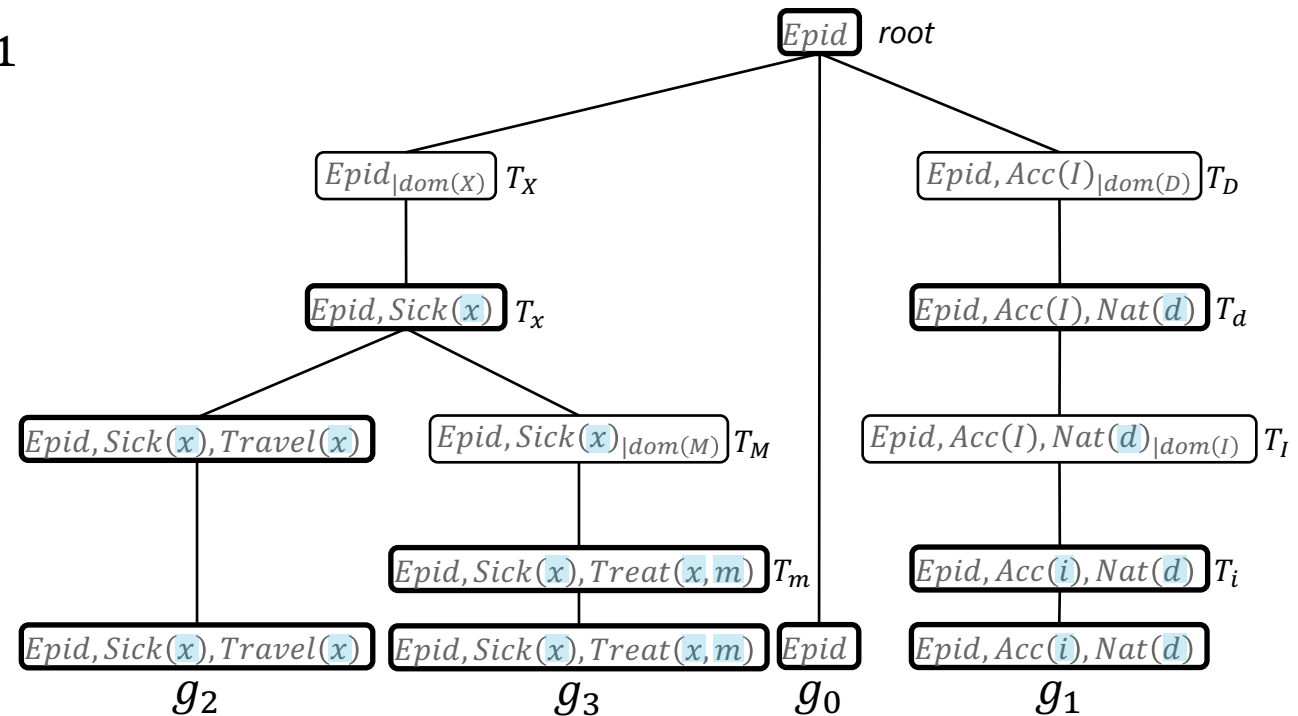
1. $\forall C \in V : C \subseteq rv(G)$
2. $\forall g \in G : \exists C \in V : rv(g) \subseteq C$
3. If $\exists A \in rv(G) : A \in C_i \wedge A \in C_j$ with $C_i, C_j \in V$, then $\forall C_k \in V$ on the path between $C_i, C_j : A \in C_k$



FO Dtree \rightarrow FO Jtree

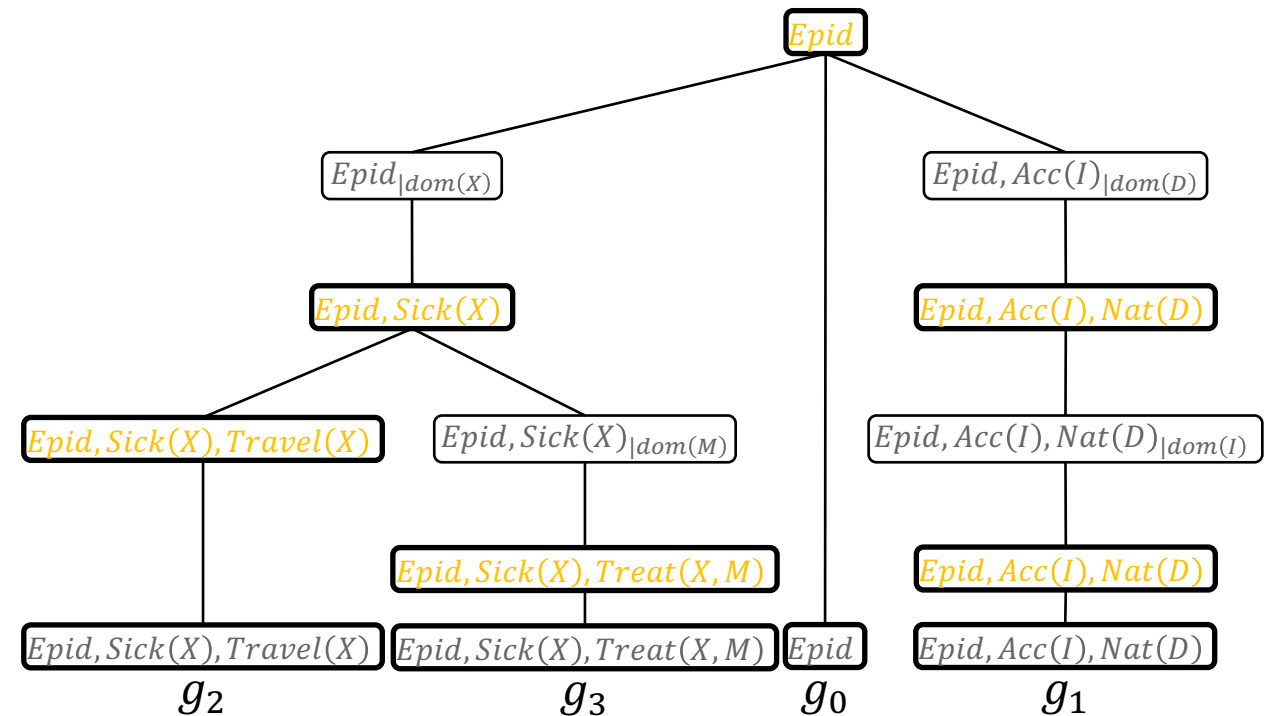
- But: Parclusters may contain a logical variable X or its representative x
- For each source DPG node T_X
 - Apply the inverse substitution θ^{-1} to the one applied during FO dtree construction to all parclusters that come from descendants of T_X :

$$\begin{aligned} & \theta^{-1} \\ &= \{X \rightarrow x\}^{-1} \\ &= \{x \rightarrow X\} \end{aligned}$$



FO Dtree \rightarrow FO Jtree

- Result after transformation **not minimal**
 - Can remove complete parclusters and still have an FO jtree
 - Even if we keep parclusters that carry constraint information that we would otherwise lose
 - E.g.,
 - Parclusters **marked**
- Observation
 - Parclusters are subsets of other parclusters
 - Use for minimisation



Minimisation

- Merge parclusters C_i, C_j with local models G_i, G_j iff $gr(C_i) \subseteq gr(C_j) \vee gr(C_j) \subseteq gr(C_i)$
 - Assuming T constraints and same names for logical variables that reference the same domain (from normal form of FO dtree), then the following suffices:

$$C_i \subseteq C_j \vee C_j \subseteq C_i$$

- Checking on a PRV and logical variable level instead of a grounded level
- Merging parclusters C_i, C_j into parcluster C_k

- $C_k = C_i \cup C_j$

- $G_k = G_i \cup G_j$

- Changes in FO jtree (V, E)

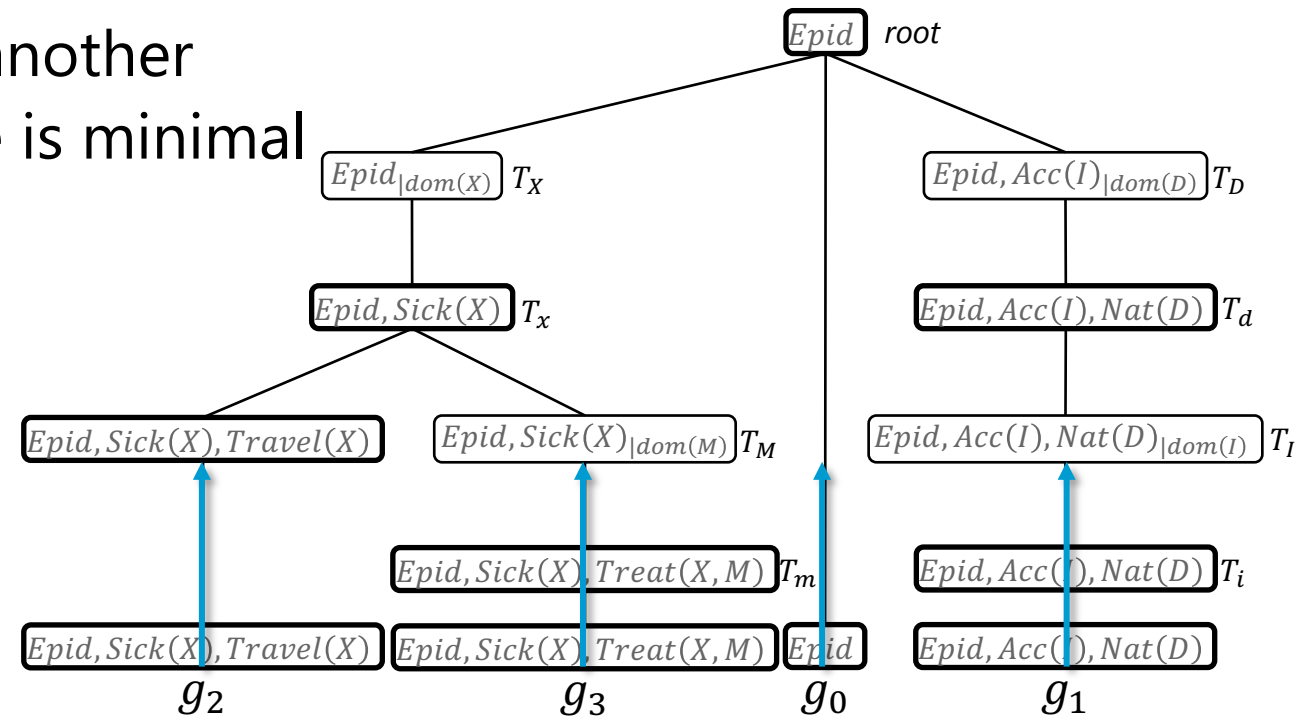
- $V = V \setminus \{C_i, C_j\} \cup C_k$

- $E = E \setminus \{\{i, l\} \mid l \in nbs(i)\} \setminus \{\{j, l\} \mid l \in nbs(j)\} \cup \{\{k, l\} \mid l \in nbs(i) \vee l \in nbs(j), l \neq i, l \neq j\}$

Reassigns all neighbours
of C_i, C_j to C_k

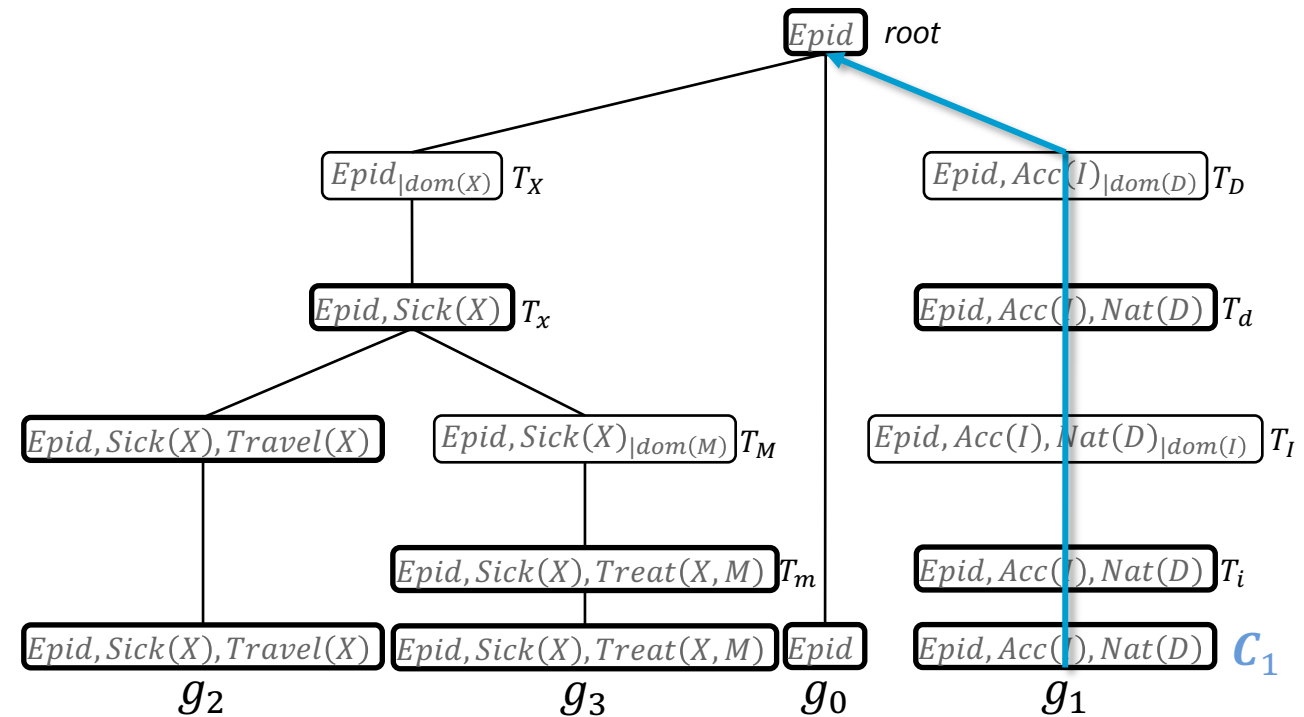
Minimisation

- Possible merging strategy
 - Start at leaves and merge **inbound**
 - Until no further merging possible
 - I.e., no parcluster a subset of another
- After merging, the resulting FO jtree is minimal
- E.g.,
 - Start at leaves with
 - local model $\{g_0\}$
 - local model $\{g_1\}$
 - local model $\{g_2\}$
 - local model $\{g_3\}$



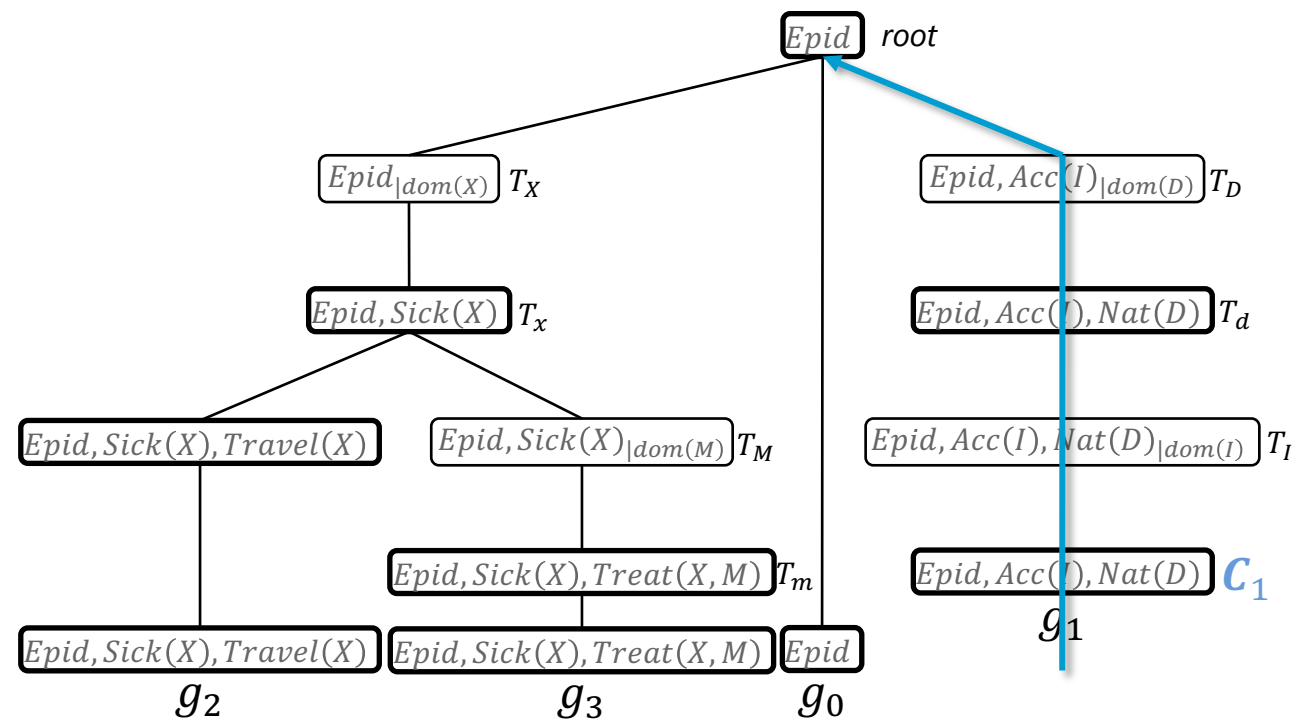
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_1\}$
 - Let us call it C_1
 - Merge **inbound**
 - C_1 and T_i parcluster identical
 - merge (call result C_1 again)



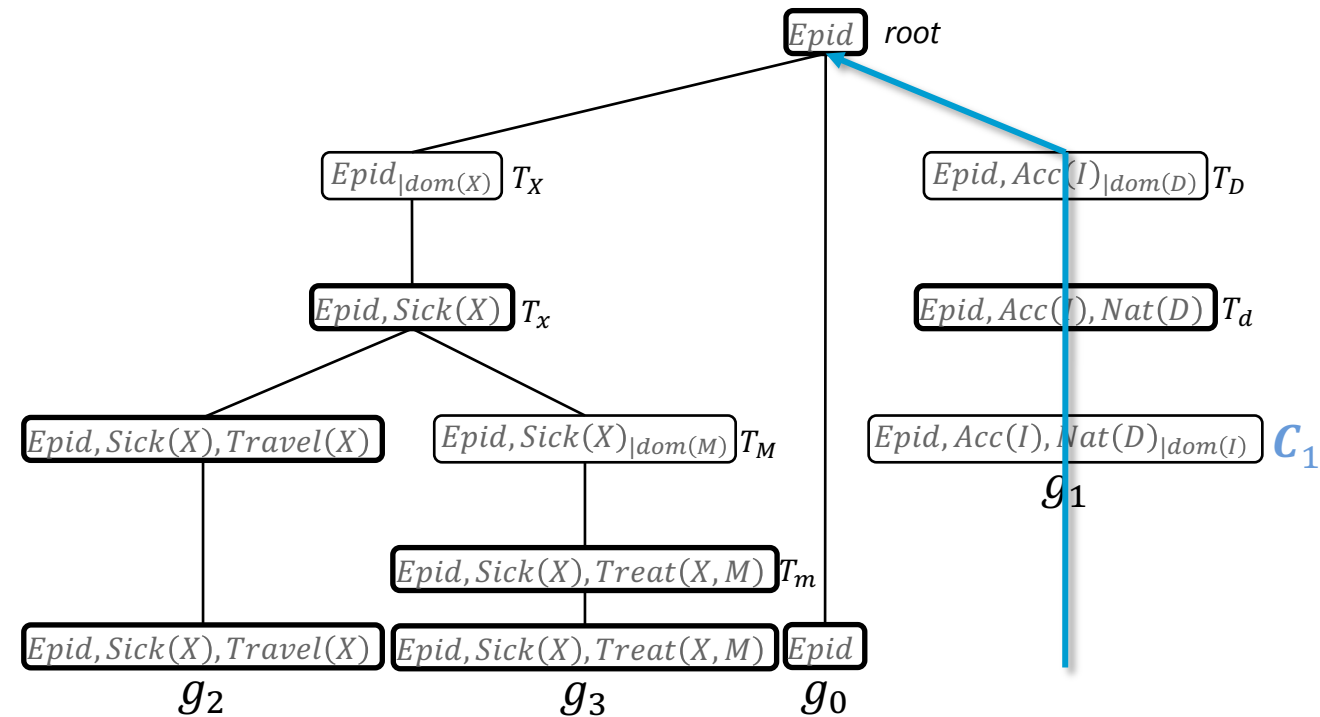
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_1\}$
 - Let us call it C_1
 - Merge **inbound**
 - C_1 and T_i parcluster identical
→ merge (call result C_1 again)
 - C_1 and T_I parcluster identical
→ merge (call result C_1 again)



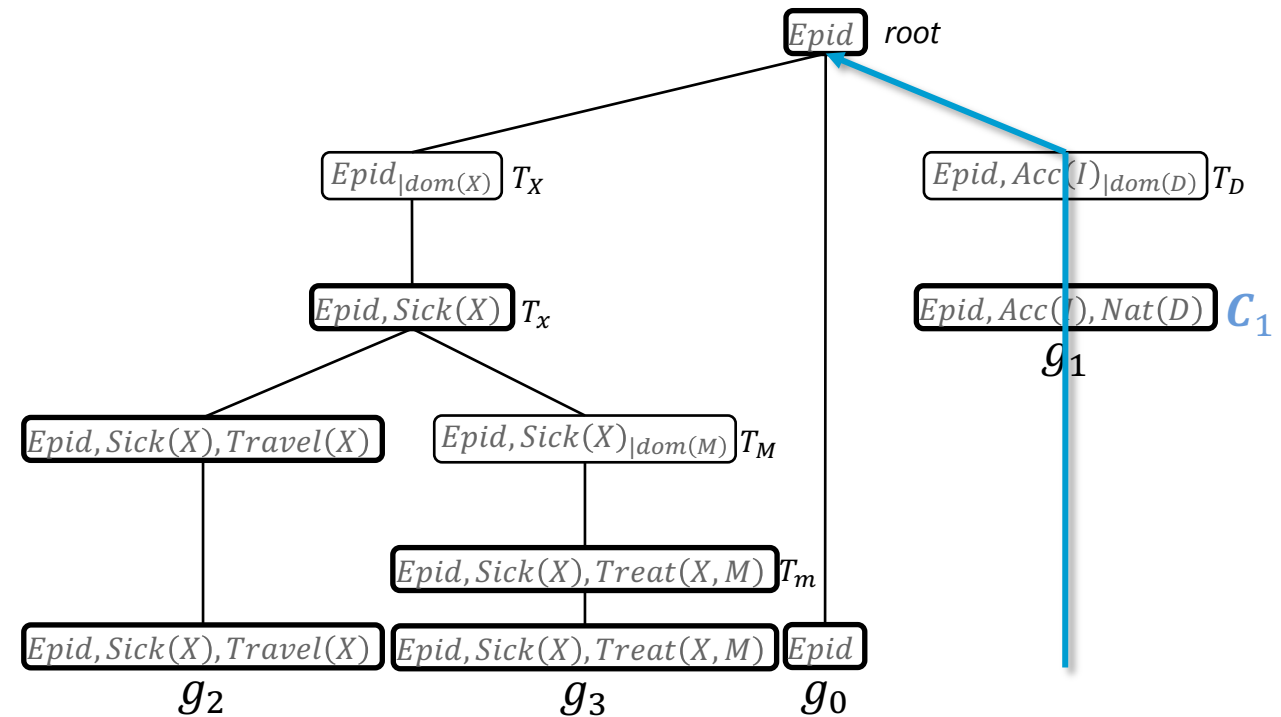
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_1\}$
 - Let us call it \mathcal{C}_1
 - Merge **inbound**
 - \mathcal{C}_1 and T_i parcluster identical \rightarrow merge (call result \mathcal{C}_1 again)
 - \mathcal{C}_1 and T_l parcluster identical \rightarrow merge (call result \mathcal{C}_1 again)
 - \mathcal{C}_1 and T_d parcluster identical \rightarrow merge (call result \mathcal{C}_1 again)



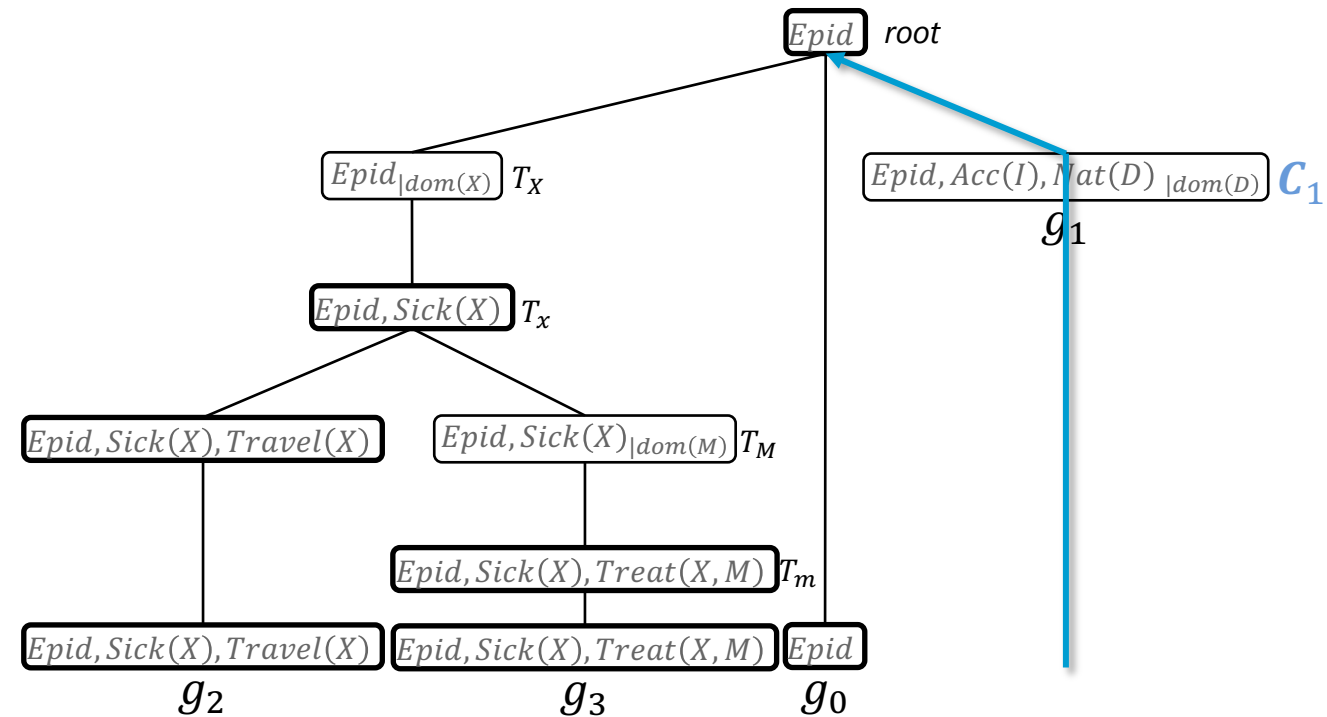
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_1\}$
 - Let us call it \mathcal{C}_1
 - Merge **inbound**
 - \mathcal{C}_1 and T_i parcluster identical \rightarrow merge (call result \mathcal{C}_1 again)
 - \mathcal{C}_1 and T_I parcluster identical \rightarrow merge (call result \mathcal{C}_1 again)
 - \mathcal{C}_1 and T_d parcluster identical \rightarrow merge (call result \mathcal{C}_1 again)
 - T_D parcluster subset of \mathcal{C}_1 \rightarrow merge (call result \mathcal{C}_1 again)



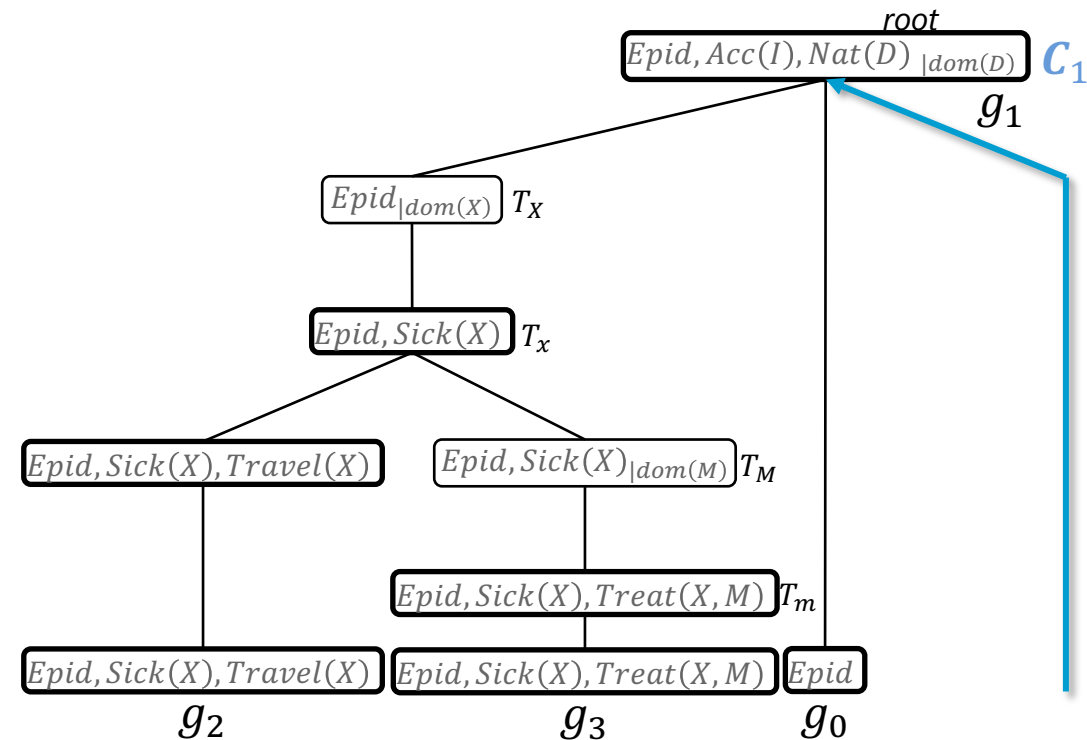
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_1\}$
 - Let us call it \mathcal{C}_1
 - Merge **inbound**
 - \mathcal{C}_1 and T_i parcluster identical
→ merge (call result \mathcal{C}_1 again)
 - \mathcal{C}_1 and T_I parcluster identical
→ merge (call result \mathcal{C}_1 again)
 - \mathcal{C}_1 and T_d parcluster identical
→ merge (call result \mathcal{C}_1 again)
 - T_D parcluster subset of \mathcal{C}_1
→ merge (call result \mathcal{C}_1 again)
 - Root parcluster subset of \mathcal{C}_1
→ merge (call result \mathcal{C}_1 again)



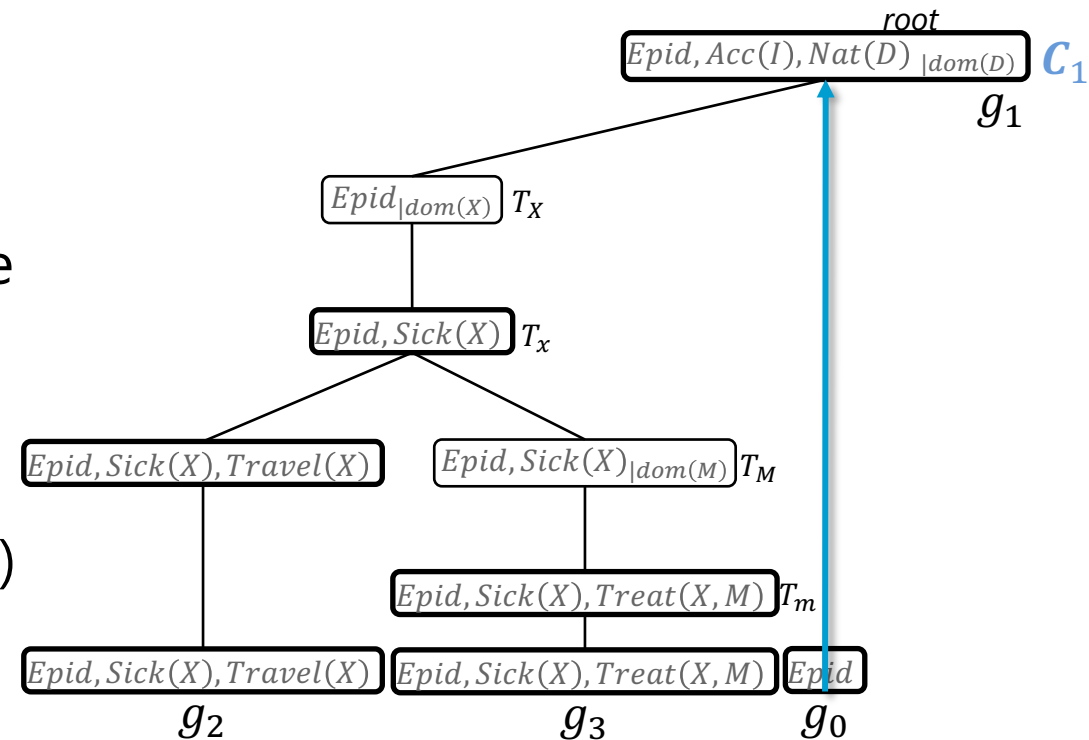
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_1\}$
 - Let us call it C_1
 - At this point, we have reached the former root and cannot merge further inbound
 - Also: the T_x parcluster contains logical variable X , which is not a subset or superset of the logical variables of C_1 (D, I)
 - Merging *stops*



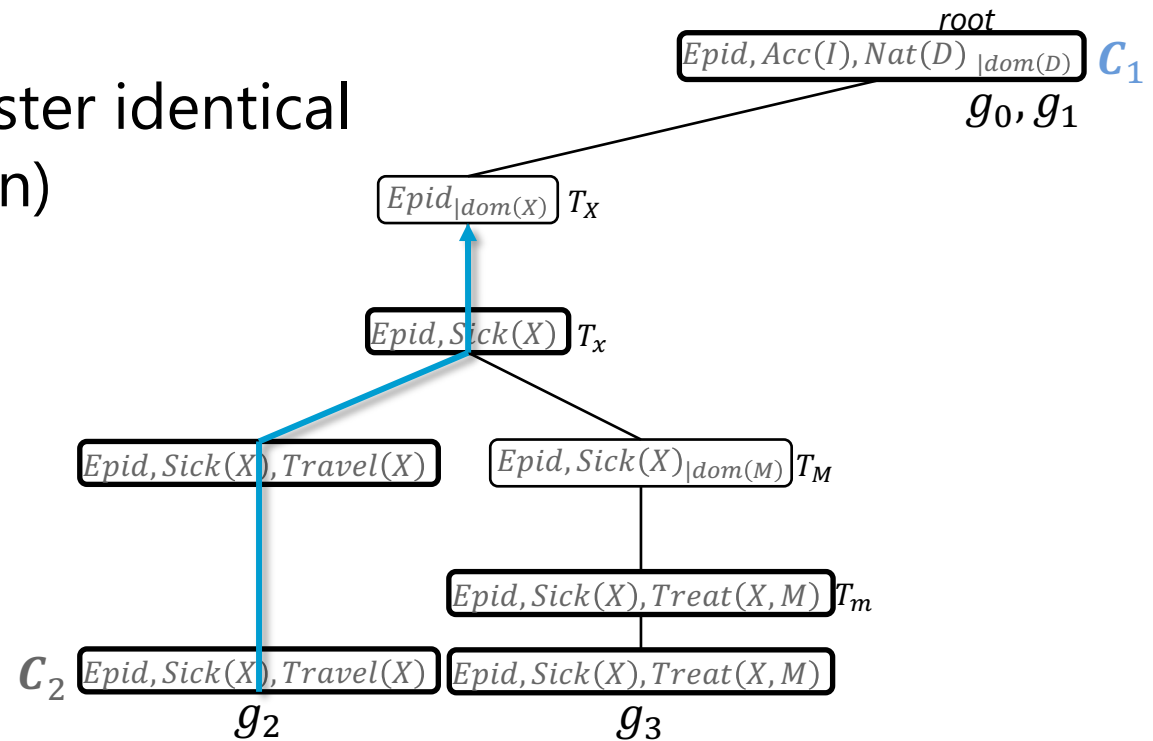
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_0\}$
 - Let us call it \mathcal{C}_0
 - Merge **inbound**
 - \mathcal{C}_0 subset of \mathcal{C}_1
 - merge (call result \mathcal{C}_1 again)
 - At this point, we have reached the former root again and cannot merge further inbound
 - Also again: the T_X parcluster contains logical variable X , which is not a subset or superset of the logical variables of \mathcal{C}_1 (D, I)
 - Merging *stops*



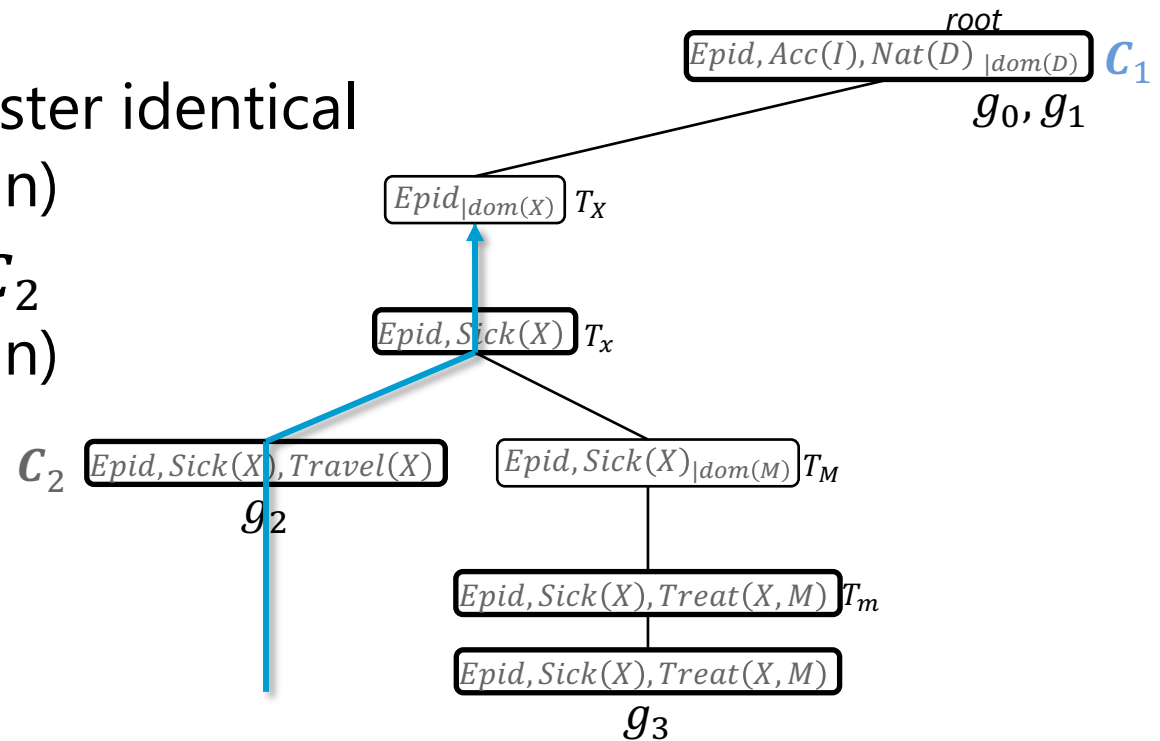
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_2\}$
 - Let us call it \mathcal{C}_2
 - Merge **inbound**
 - \mathcal{C}_2 and neighbouring parcluster identical
→ merge (call result \mathcal{C}_2 again)



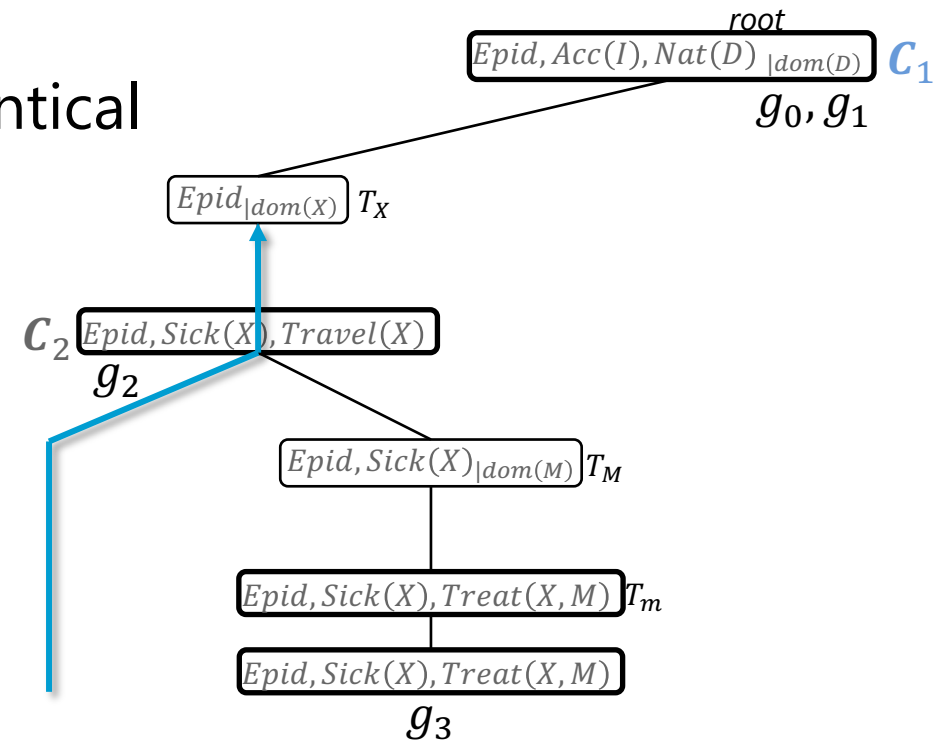
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_2\}$
 - Let us call it \mathcal{C}_2
 - Merge **inbound**
 - \mathcal{C}_2 and neighbouring parcluster identical
→ merge (call result \mathcal{C}_2 again)
 - T_x parcluster is a subset of \mathcal{C}_2
→ merge (call result \mathcal{C}_2 again)



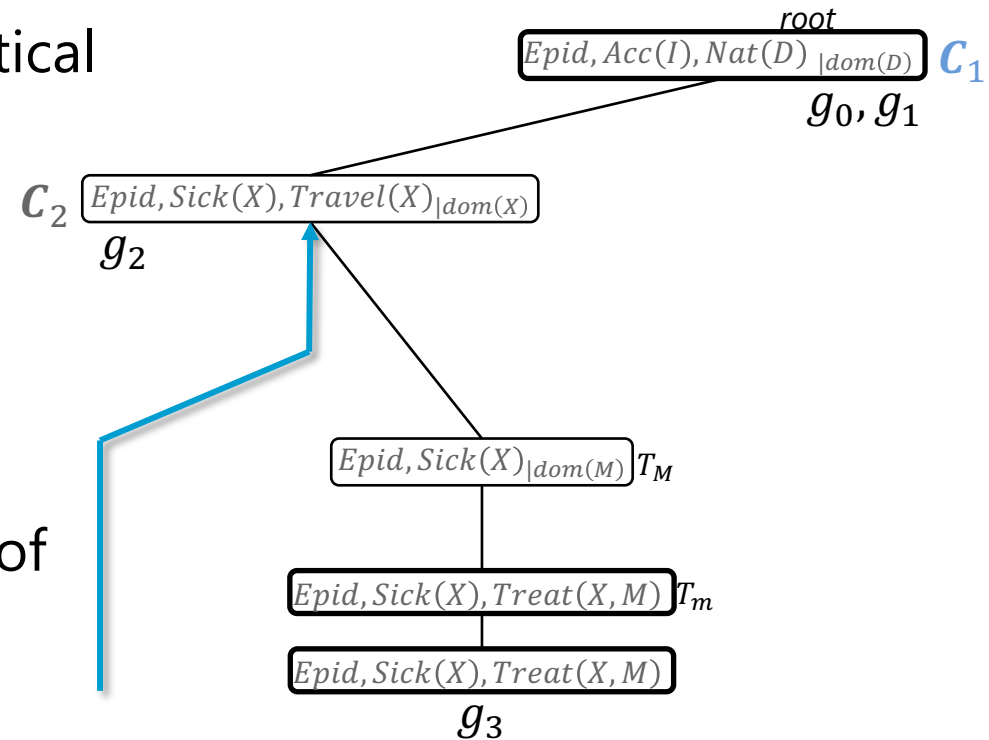
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_2\}$
 - Let us call it \mathcal{C}_2
 - Merge **inbound**
 - \mathcal{C}_2 and neighbouring parcluster identical
→ merge (call result \mathcal{C}_2 again)
 - T_x parcluster is a subset of \mathcal{C}_2
→ merge (call result \mathcal{C}_2 again)
 - T_X parcluster is a subset of \mathcal{C}_2
→ merge (call result \mathcal{C}_2 again)



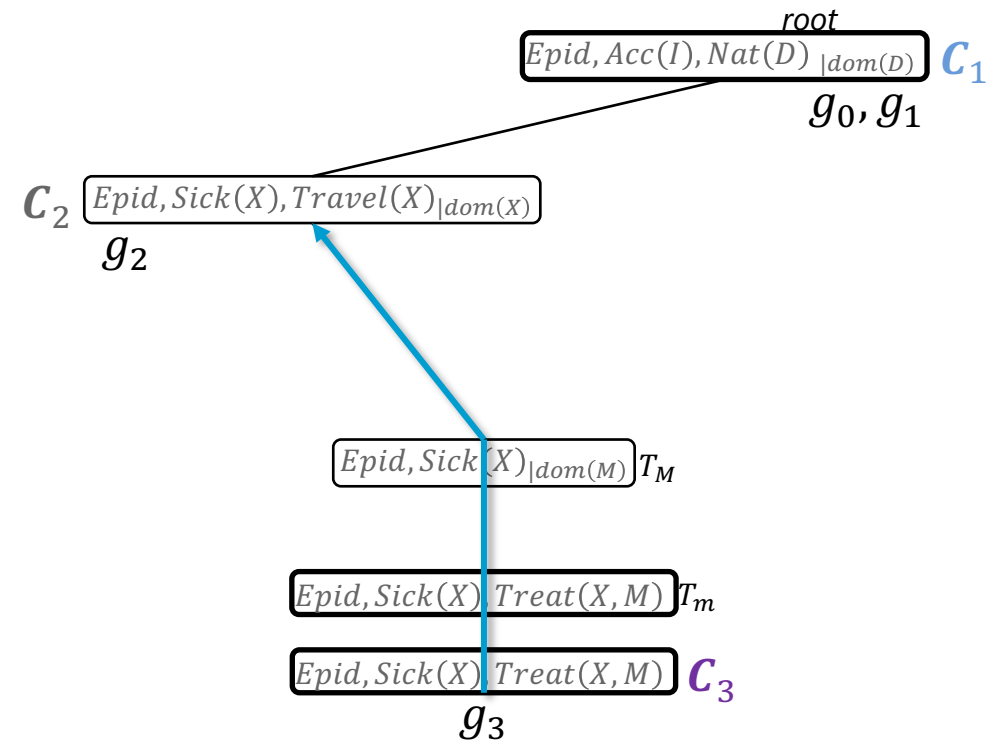
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_2\}$
 - Let us call it \mathcal{C}_2
 - Merge **inbound**
 - \mathcal{C}_2 and neighbouring parcluster identical
→ merge (call result \mathcal{C}_2 again)
 - T_x parcluster is a subset of \mathcal{C}_2
→ merge (call result \mathcal{C}_2 again)
 - T_X parcluster is a subset of \mathcal{C}_2
→ merge (call result \mathcal{C}_2 again)
 - Merging cannot move further inbound
 - \mathcal{C}_1 is neither a subset nor a superset of \mathcal{C}_2
 - Merging *stops*




Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_3\}$
 - Let us call it C_3
 - Merge **inbound**
 - C_3 and T_m parcluster identical
→ merge (call result C_3 again)



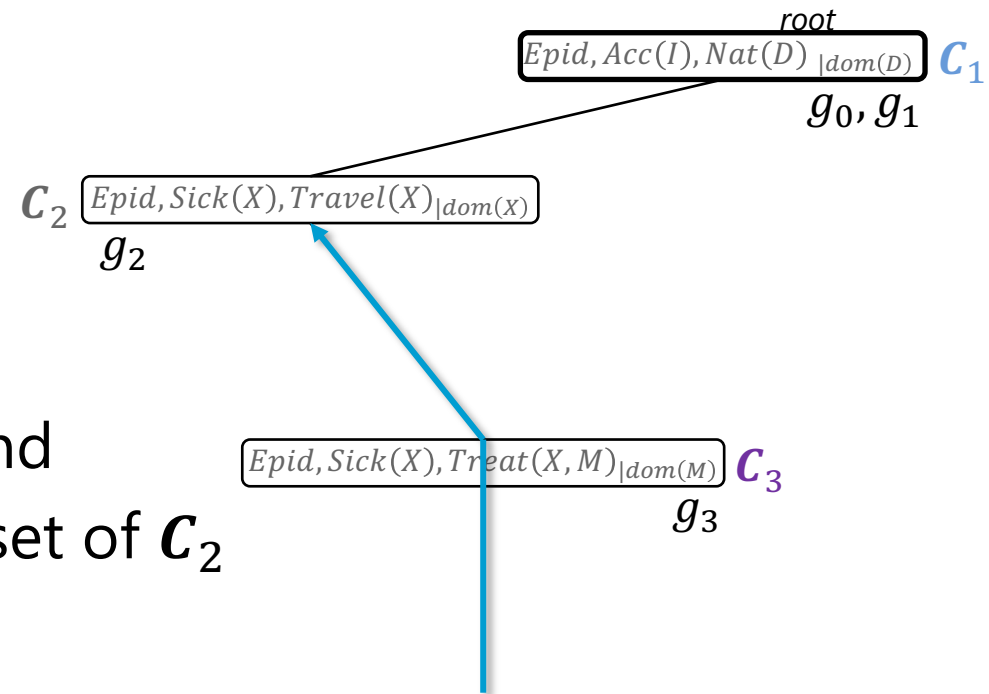
Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_3\}$
 - Let us call it \mathcal{C}_3
 - Merge **inbound**
 - \mathcal{C}_3 and T_m parcluster identical
→ merge (call result \mathcal{C}_3 again)
 - T_M parcluster is a subset of \mathcal{C}_3
→ merge



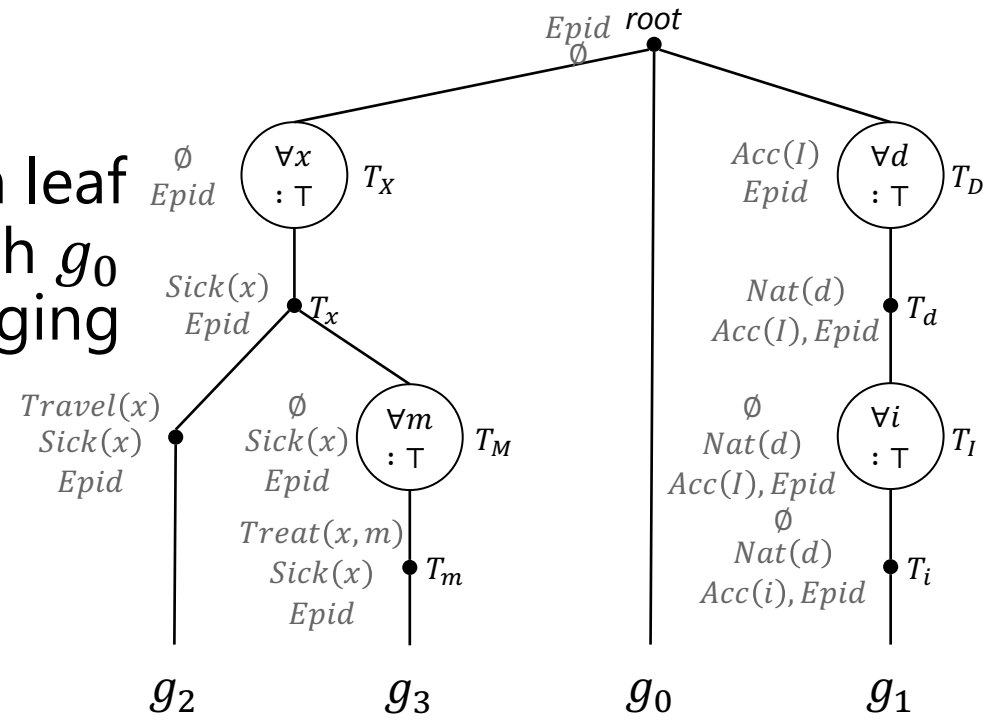


Minimisation: Example Continued

- Consider leaf parcluster with local model $\{g_3\}$
 - Let us call it \mathcal{C}_3
 - Merge **inbound**
 - \mathcal{C}_3 and T_m parcluster identical
→ merge (call result \mathcal{C}_3 again)
 - T_M parcluster is a subset of \mathcal{C}_3
→ merge (call result \mathcal{C}_3 again)
 - Merging cannot move further inbound
 - \mathcal{C}_3 is neither a subset nor a superset of \mathcal{C}_2
 - Merging *stops*



Minimisation: Example Continued

- Resulting FO jtree J from FO dtree T given model G
 - If we had started merging from leaf with g_3 inbound before merging from leaf with g_2 , C_2 and C_3 would be switched
 - g_0 could have made one of the other parclusters if we had started merging from leaf with g_2 or g_3 before merging from leaf with g_0 or by starting at leaf with g_0 and then merging from leaf with g_2 or g_3



FO Jtree Construction


- Given a model G , the following steps are necessary
 1. Bring G into the required normal form for FO dtree construction
 2. Construct an FO dtree T for G
 3. Translate T into an FO jtree J
 4. Apply inverse substitutions to parclusters of descendants of DPG nodes in J
 5. Minimise J

Construction

- Next?
- FO jtrees for query answering
 - Messages need to be passed to ensure independence
 - What about evidence?

Message Passing in FO Jtrees

- Ensure independence between parclusters
- Send messages based on two conditions
 - If a node i has received all messages from neighbours but one, j , node i calculates and sends a message to j
 - If a node i has received all messages, then it calculates and sends messages to all neighbours j that have not received a message yet




Message Passing in FO Jtrees

- Message m_{ij} from sender C_i to receiver C_j
 - Set of parfactors $\{g_l\}_{l=1}^n$ with $rv(g_l) \subseteq S_{ij}$
 - To calculate
 - Collect necessary information from local model and received messages:

$$G_{ij} = G_i \cup \bigcup_{k \in nbs(i), k \neq j} m_{ki}$$

- Ignore the message that came from C_j (if it already exists)
- Call slightly modified LVE with G_{ij} as input model, S_{ij} as query, and no evidence: $LVE\text{-}MSG(G_{ij}, S_{ij})$
 - Specification of LVE–MSG: next slide



LVE for Message Passing

LVE-MSG(G, S)

$G \leftarrow$ Shatter G on itself

No shattering on separator (due to construction) or evidence, no absorption (will have been handled)

- Model might need to be shattered on itself because of splits introduced by messages

while G contains non-query terms **do**

if a PRV A fulfils the preconditions of sum-out **then**

$G \leftarrow$ Apply sum-out to A in G

else

$G \leftarrow$ Apply an enabling operator on some parfactors in G


return G

No normalisation (and multiplication of the remaining factors to be able to normalise) at the end

- Interim result returned


Message Passing in FO Jtrees: Example

- Message m_{12} from C_1 to C_2
 - Collect $G_{12} = \{g_1\} \cup \emptyset$
 - No further neighbours except C_2
 - Call $LVE\text{--}MSG(\{g_1\}, \{Epid\}, \emptyset)$
 - $LVE\text{--}MSG$ eliminates $Nat(D), Acc(I)$ from $\{g_1\}$
 - Count-converting $Nat(D)$ into $\#_D[Nat(D)]$
 - Summing out $Acc(I)$
 - Summing out $\#_D[Nat(D)]$
 - Returning $\{g'_1\}$
 - Send $\{g'_1\}$ as m_{12} to C_2




Message Passing in FO Jtrees: Example

- Message m_{32} from C_3 to C_2
 - Collect $G_{32} = \{g_3\} \cup \emptyset$
 - No further neighbours except C_2
 - Call $LVE\text{--}MSG(\{g_3\}, \{Epid, Sick(X)\}, \emptyset)$
 - $LVE\text{--}MSG$ eliminates $Treat(X, M)$ from $\{g_3\}$
 - Summing out $Treat(X, M)$
 - Returning $\{g'_3\}$
 - Send $\{g'_3\}$ as m_{32} to C_2




Message Passing in FO Jtrees: Example

- Message m_{21} from C_2 to C_1
 - Collect $G_{21} = \{g_2\} \cup m_{32}$
 - Further neighbour: C_3 , sent message $m_{32} = \{g'_3\}$
 - Call $LVE\text{--}MSG(\{g_2, g'_3\}, \{Epid\}, \emptyset)$
 - $LVE\text{--}MSG$ eliminates $Travel(X), Sick(X)$ from $\{g_2, g'_3\}$
 - Summing out $Travel(X)$ from g_2 , yielding g'_2
 - Summing out $Sick(X)$ from product of g'_2 and g'_3 , yielding g'_{23}
 - Returning $\{g'_{23}\}$
 - Send $\{g'_{23}\}$ as m_{21} to C_1



Message Passing in FO Jtrees: Example

- Message m_{23} from C_2 to C_3
 - Collect $G_{23} = \{g_2\} \cup m_{12}$
 - Further neighbour: C_1 , sent message $m_{12} = \{g'_1\}$
 - Call $LVE^*(\{g_2, g'_1\}, \{Epid, Sick(X)\}, \emptyset)$
 - LVE^* eliminates $Travel(X)$ from $\{g_2, g'_1\}$
 - Summing out $Travel(X)$ from g_2 , yielding g'_2
 - Returning $\{g'_2, g'_1\}$
 - Send $\{g'_2, g'_1\}$ as m_{23} to C_3



Message Passing: Overview

- Given an FO jtree J , send messages if one of the two conditions is true
 - If a node i has received all messages from neighbours but one, j , node i calculates and sends a message to j
 - If a node i has received all messages, then it calculates and sends messages to all neighbours j that have not received a message yet
- To calculate a message:
 - Collect necessary information from local model and received messages:

$$G_{ij} = G_i \cup \bigcup_{k \in \text{nbs}(i), k \neq j} m_{ki}$$

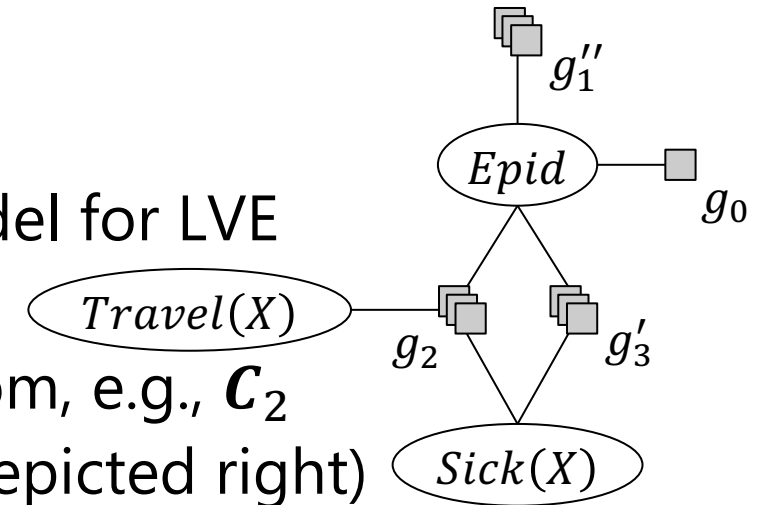
- Call $\text{LVE-MSG}(G_{ij}, \mathbf{S}_{ij})$

Message Passing

Query Answering in FO Jtrees

- Idea


- Pick parcluster in which query terms occur
- Use local model and outside messages as input model for LVE
- E.g., for $P(Epid)$
 - All parclusters contain *Epid*, choose one at random, e.g., C_2
 - Collect $G_{Epid} = \{g_2\} \cup m_{12} \cup m_{32} = \{g_2, g'_1, g'_3\}$ (depicted right)
 - Call $LVE(\{g_2, g'_1, g'_3\}, Epid, \emptyset)$, yielding a parfactor g containing the probability distribution over *Epid*



- What if query terms occur outside of one parcluster?

- E.g.,

$P(\text{Travel}(\text{eve}), \text{Treat}(\text{eve}, m_1))$



Query Answering in FO Jtrees

How can we find a subgraph?


- For query terms Q , possibly contained in more than one parcluster
 - Find a subgraph J' of the FO jtree J such that $Q \subseteq rv(J')$
 - Use local models in J' and messages from outside J' as basis for calling LVE
 - No duplicate information used
 - E.g., query on R_1, R_5 using C_i, C_k, C_j
 - Ignore inside messages $m_{ki}, m_{ik}, m_{jk}, m_{kj}$

- Subgraph should be minimal in the number of PRVs in it for optimal performance:

$$\operatorname{argmin}_{J'} |rv(J')|$$



s.t. $Q \subseteq rv(J')$

- Trade-off between finding a subgraph fast and finding a minimal one
- It is not about the number of parclusters!



Query Answering in FO Jtrees: Example

- E.g., $P(\text{Travel}(\text{eve}), \text{Treat}(\text{eve}, m_1))$
 - Subgraph: C_2, C_3
 - Submodel for query answering: $G_Q = (g_2, g_3, m_{12})$
 - Depicted right
 - Call LVE with G_Q and $Q = \{\text{Travel}(\text{eve}), \text{Treat}(\text{eve}, m_1)\}$
 - Split off query terms
 - Eliminate all non-query terms
 - Normalise the result



Query Answering in FO Jtrees


- After message passing, parclusters independent from each other given messages
 - Prepared for query answering

- For each query with query terms Q
 - Find subtree $J' = (V', E')$ s.t. $Q \subseteq rv(J')$
 - Collect information from local model and messages, i.e,

$$G_Q = \bigcup_{i \in V'} G_i \cup \bigcup_{\substack{j \in nbs(i) \\ j \notin V'}} m_{ji}$$

- Call $LVE(G_Q, Q, \emptyset)$ and return or store result of the call


- What about evidence?



Evidence in FO Jtrees

- Evidence applies to PRVs in some parclusters
 - Changes the distributions in local models
 - Information sent in messages might change
 - Even if summed out and therefore hidden from the other parclusters
- Therefore, handle evidence before sending messages
 - Only then send messages


- Given a set of evidence parfactors $\{\phi_e(A_e)|_{C_e}\}_{e=1}^m$
- For each $\phi_e(A_e)|_{C_e}$
 - For each parcluster C_i where $A_e \in C_i$
 - Shatter G_i on C_e
 - Absorb $\phi_e(A_e)|_{C_e}$ in G_i **Evidence Entering**



Evidence in FO Jtrees: Example


- Given $Sick(eve) = true$ as evidence g_e
 - In C_2
 - Shatter $G_2 = \{g_2\}$ on $Sick(eve)$, yielding $\{g_2^e, g_2'\}$
 - Absorb g_e in g_2^e , yielding $g_2^{e'}$
 - Result: $G_2 = \{g_2^{e'}, g_2'\}$
 - In C_3
 - Shatter $G_3 = \{g_3\}$ on $Sick(eve)$, yielding $\{g_3^e, g_3'\}$
 - Absorb g_e in g_3^e , yielding $g_3^{e'}$
 - Result: $G_3 = \{g_3^{e'}, g_3'\}$

- After evidence handling, send messages based on the local models that have absorbed the evidence
 - $G_0 = \{g_0, g_1\}$ (unchanged)
 - $G_2 = \{g_2^{e'}, g_2'\}$
 - $G_3 = \{g_3^{e'}, g_3'\}$




Evidence in FO Jtrees

- E.g., given $Sick(eve) = true$ as evidence in g_e
 - Message m_{12} does not change compared to previous example
 - Message m_{32} calculated based on $\{g_3^{e'}, g_3'\}$
 - Call $LVE-MSG(\{g_3^{e'}, g_3'\}, \{Epid, Sick(X)\}, \emptyset)$, yielding $\{g_3^{e''}, g_3''\}$
 - Message m_{23} calculated based on $\{g_2^{e'}, g_2'\} \cup m_{12}$
 - Call $LVE-MSG(\{g_2^{e'}, g_2', g_1'\}, \{Epid, Sick(X)\}, \emptyset)$, yielding $\{g_2^{e''}, g_2'', g_1'\}$
 - Message m_{21} calculated based on $\{g_2^{e'}, g_2'\} \cup m_{32}$
 - Call $LVE-MSG(\{g_2^{e'}, g_2', g_3^{e''}, g_3''\}, \{Epid\}, \emptyset)$, yielding $\{g_2^{e''}, g_2'', g_3^{e''}, g_3''\}$



Evidence and Queries in FO Jtrees

- After evidence handling
 - All queries are answered in an FO jtree with handled evidence $\{g_e\}_{e=1}^m$ yield results conditional on $\{g_e\}_{e=1}^m$
 - So, given evidence $\{g_e\}_{e=1}^m$ and query terms $\{Q_i\}_{i=1}^n$ for a model G
 - The posed queries are $P(Q_i | \{g_e\}_{e=1}^m)$, $1 \leq i \leq n$, w.r.t. P_G
- FO jtree constructed without specific evidence
 - Reuse for different evidence sets
 - As long as model stays the same
 - Reset the local models before entering new evidence



LJT: Algorithm

$\text{LJT}(G, \{Q_i\}_{i=1}^n, \{g_e\}_{e=1}^m)$

Construct an FO jtree J for G

Enter evidence $\{g_e\}_{e=1}^m$ into J

Pass message in J

Answer queries with query terms $\{Q_i\}_{i=1}^n$ in J


- Look for blue boxes on the previous slides to find the descriptions of each step

Step Name

- *Constant overhead* for FO jtree construction
- *Payoff* if given multiple queries

Foundations of Clustering

- History in propositional probabilistic inference:
 - Based on **probability propagation** introduced by Pearl (1988)
 - If a BN is a **polytree**, i.e., the underlying undirected graph has no trivial cycles, then
 - Treat each node in a BN as a cluster with the random variables of the accompanying conditional probability table as the cluster random variables
 - Send messages along the edges (to parents and children), eliminating random variables not occurring in the parent or child nodes

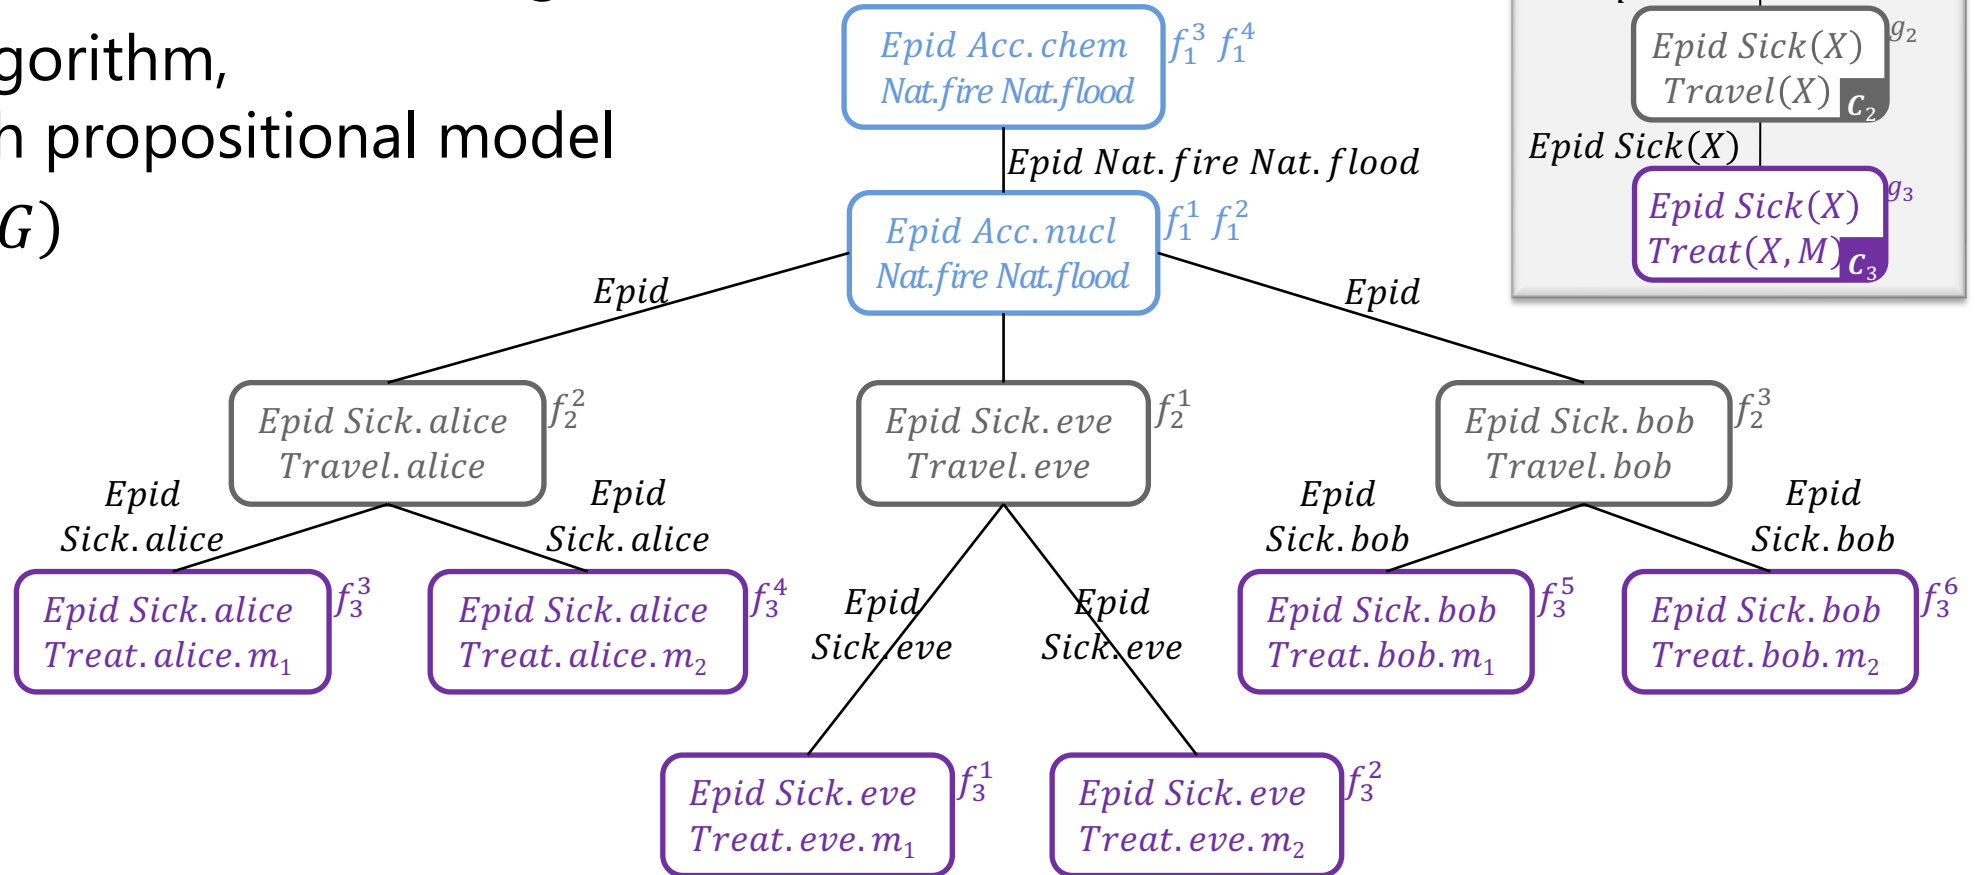


Foundations of Clustering

- History in propositional probabilistic inference:
 - If no polytree, the cycles mess up the message passing along the edges (information arrives multiple times)
 - Send messages nonetheless (exact if polytree, approximate otherwise): called **belief propagation** as an algorithm for *approximate* inference
 - Exact inference required → put the cycles into one cluster
 - Graph formed called a junction tree (**jtree**)
 - A first-order version of a jtree was induced on the previous slides
 - Also known as *clique tree* (since the cycles often form cliques in the model graph) or join tree
 - Propositional version introduced by Lauritzen and Spiegelhalter (1988)
 - Shenoy and Shafer (1989) introduce three axioms of *local computations* to show correctness of doing computations locally

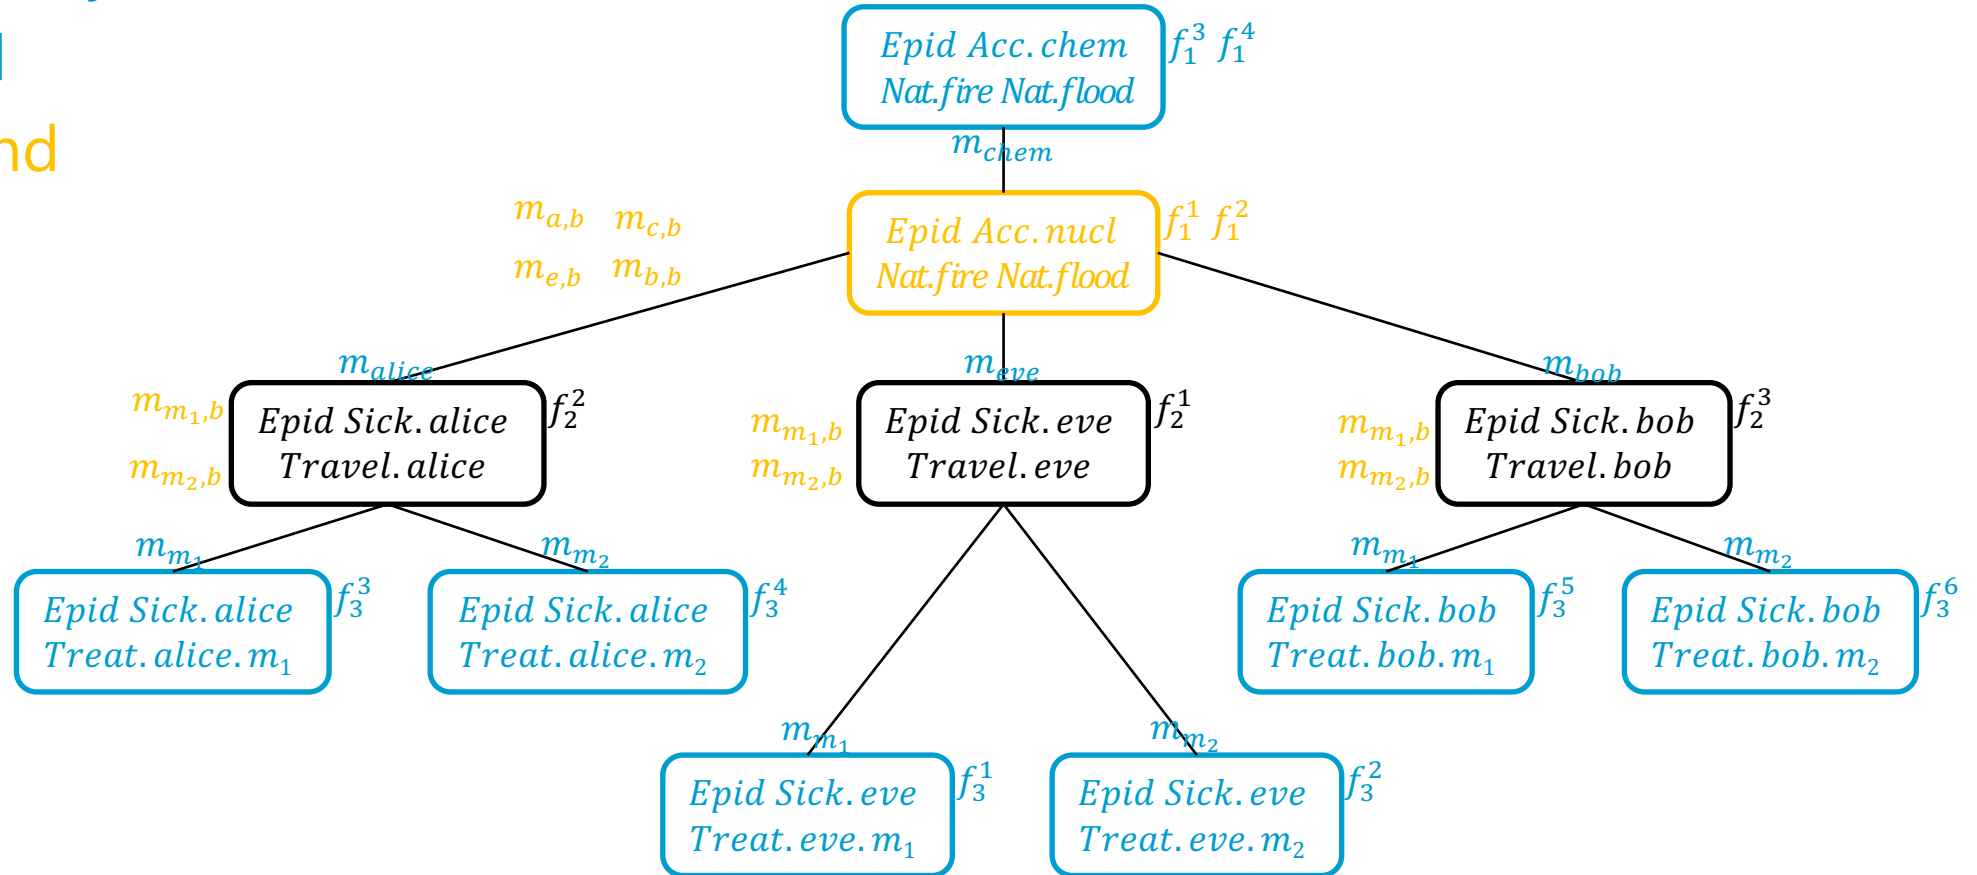
Comparison to Ground Inference

- Propositional Junction Tree Algorithm (JT)
 - Same algorithm, only with propositional model
 - E.g., $gr(G)$



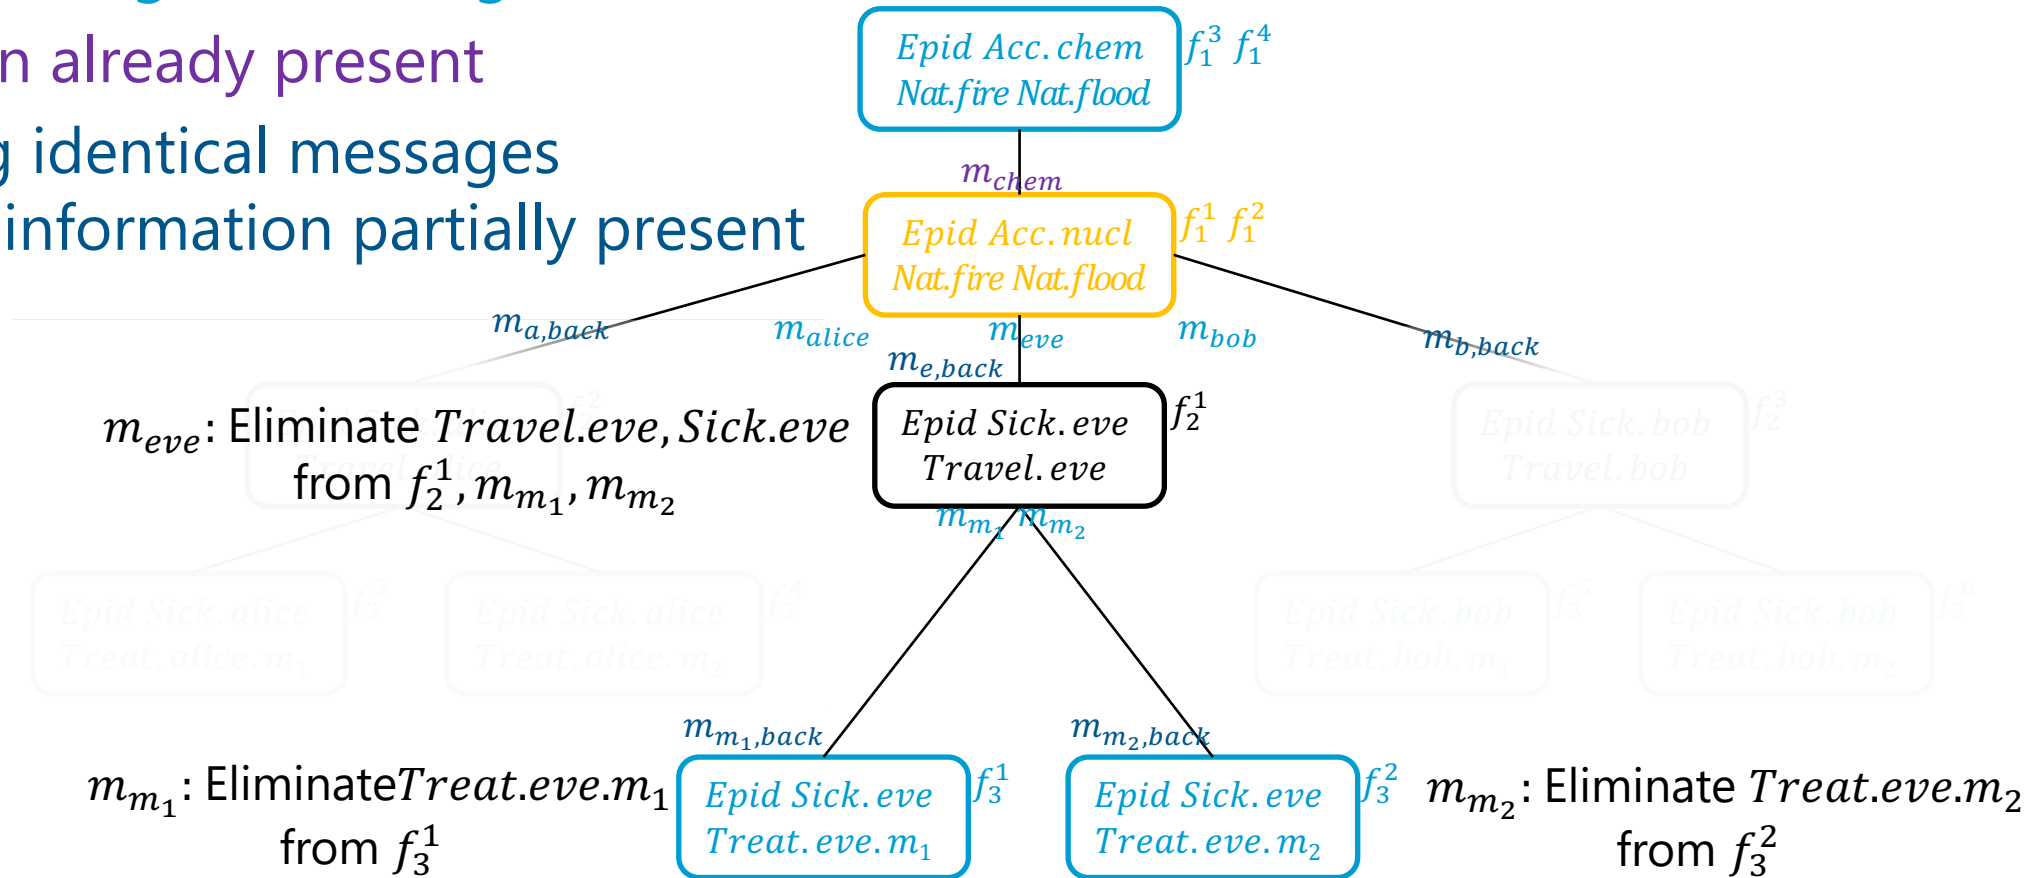
Junction Tree: Messages

- From periphery to centre and back
 - Inbound
 - Outbound



Junction Tree: Symmetry → Inefficiency

- Identical messages incoming
- Information already present
- Calculating identical messages + sending information partially present




Message Calculation Strategies

There is also a lifted version of Hugin using a lifted division operator

[Hoffmann et al., 2022]

- Strategy used so far: so-called **Shafer-Shenoy architecture** [Shafer and Shenoy, 1989]
 - Disadvantage: many operations (multiplications) duplicated
 - Especially in (FO) jtrees with high degree
 - Even if only one factor per parcluster and message
 - Example right: for each outgoing message, only one incoming message changes
- Alternative: **Hugin architecture** [Jensen et al., 1989]
 - Hugin factor $h_i = \phi_i(\mathbf{C}_i)$ per parcluster \mathbf{C}_i as a product of G_i
 - Incoming messages m_{ji} multiplied into h_i (and stored): $h_i \leftarrow h_i \cdot m_{ji}$
 - When calculating message m_{ij} back: $\text{VE-JT}(f_i / m_{ji}, S_{ij}, \emptyset, \cdot)$
 - Divide h_i by message m_{ji} , then sum out non-separators
 - One multiplication and one division instead of multiple multiplications

$$\begin{aligned}
 m_{i1} &\leftarrow \text{VE-JT}(\{f_i, m_{2i}, m_{3i}, m_{4i}, m_{5i}, m_{6i}\}, \dots) \\
 m_{i2} &\leftarrow \text{VE-JT}(\{f_i, m_{1i}, m_{3i}, m_{4i}, m_{5i}, m_{6i}\}, \dots) \\
 m_{i3} &\leftarrow \text{VE-JT}(\{f_i, m_{1i}, m_{2i}, m_{4i}, m_{5i}, m_{6i}\}, \dots) \\
 m_{i4} &\leftarrow \text{VE-JT}(\{f_i, m_{1i}, m_{2i}, m_{3i}, m_{5i}, m_{6i}\}, \dots) \\
 m_{i5} &\leftarrow \text{VE-JT}(\{f_i, m_{1i}, m_{2i}, m_{3i}, m_{4i}, m_{6i}\}, \dots) \\
 m_{i6} &\leftarrow \text{VE-JT}(\{f_i, m_{1i}, m_{2i}, m_{3i}, m_{4i}, m_{5i}\}, \dots)
 \end{aligned}$$




$h_i \leftarrow h_i \cdot m_{1i}$	$h_i / m_{1i} \rightarrow \text{VE-JT}$
$h_i \leftarrow h_i \cdot m_{2i}$	$h_i / m_{2i} \rightarrow \text{VE-JT}$
$h_i \leftarrow h_i \cdot m_{3i}$	$h_i / m_{3i} \rightarrow \text{VE-JT}$
$h_i \leftarrow h_i \cdot m_{4i}$	$h_i / m_{4i} \rightarrow \text{VE-JT}$
$h_i \leftarrow h_i \cdot m_{5i}$	$h_i / m_{5i} \rightarrow \text{VE-JT}$
$h_i \leftarrow h_i \cdot m_{6i}$	$h_i / m_{6i} \rightarrow \text{VE-JT}$

In terms of Lifting: Is it that simple?

- Algorithm-induced **groundings** due to message passing
 - For message calculation, non-separator PRVs are eliminated
 - Separator as the query terms containing *logical variables*
 - Non-separator PRVs have to fulfil sum-out preconditions
 - Preconditions 1 + 3 fulfilled by construction
 - May be that *Precondition 2 is not fulfilled* → can cause groundings
 - E.g., logical variable **E** added to PRVs *Epid*, *Sick(X)*, *Treat(X, M)*
 - When calculating m_{23} , one has to eliminate *Travel(X)*
 - But: does not contain both *X* and **E**, count conversion does not apply (**E** occurs in two PRVs) → **ground E**

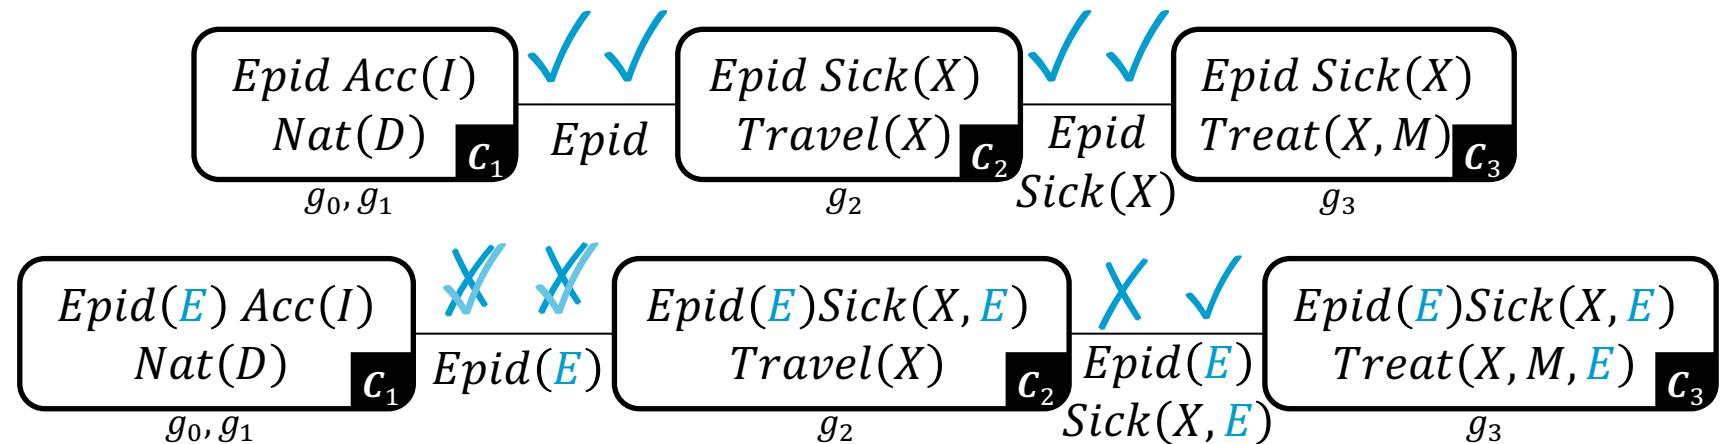
Preconditions:

- $\forall B \in rv(G \setminus \{g\}) : gr(B|_C) \cap gr(A_i|(x, c_x)) = \emptyset$
- $\forall X \in \{X \mid |\pi_X(C_x)| > 1\} : X \in lv(A_i)$
- $X^{excl} = lv(A_i) \setminus (X \setminus lv(A_i))$ count-normalised w.r.t.



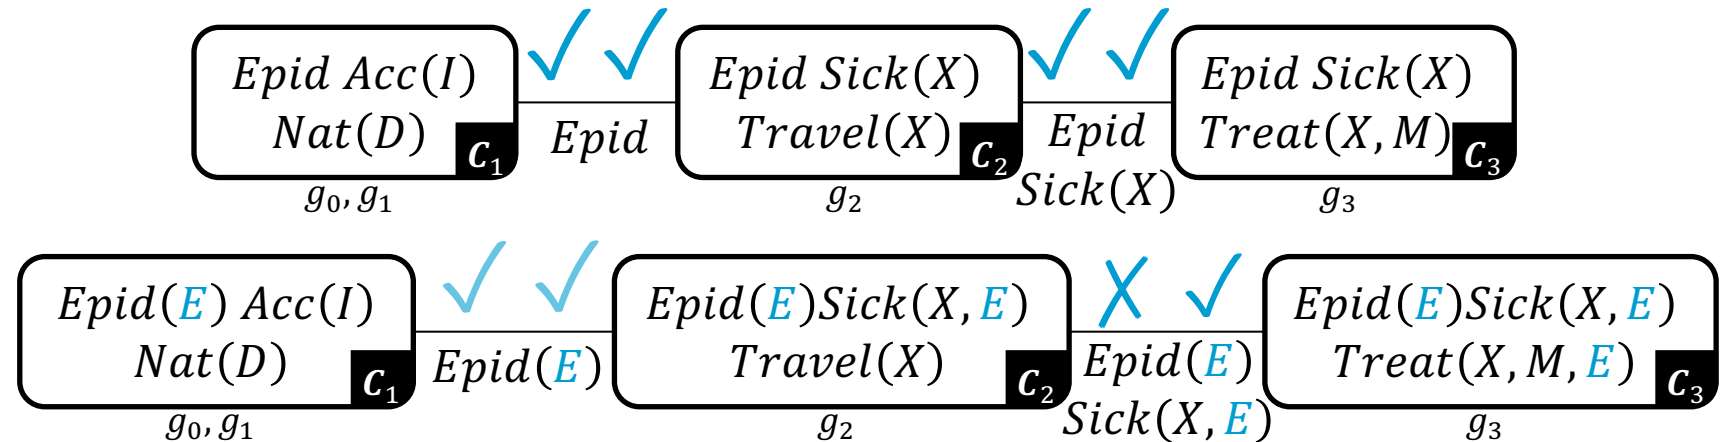
Conditions on Groundings

- For a lifted calculation of message m_{ij} , it necessarily has to hold that
 - for each PRV $A \in (\mathcal{C}_i \setminus \mathcal{S}_{ij})$, i.e., A has to be eliminated:
 - for each separator PRV $S \in \mathcal{S}_{ij} : lv(S) \subseteq lv(A)$ (Cond. 1)
- If Cond. 1 does not hold, i.e., $lv(S) \not\subseteq lv(A)$, one may induce Cond. 1 by count conversion
 - If $lv(S) \setminus lv(A)$ are countable in G_{ij} (Cond. 2)



Conditions on Groundings

- Problem with Cond. 1 induced using count conversions on logical variables $lv(S) \setminus lv(A)$:
 - Logical variables that were previously not counted are now counted
 - All receiving parclusters need to be able to handle counted versions, which needs to be checked
 - If newly counted logical variable arrives at parcluster C_k , it does not directly lead to groundings in G_k as well (Cond. 3)
 - For further calculations, since they refer to the same set of randvars, they have to occur in the same form, i.e., at one point the logical variable has to be counted in G_k as well

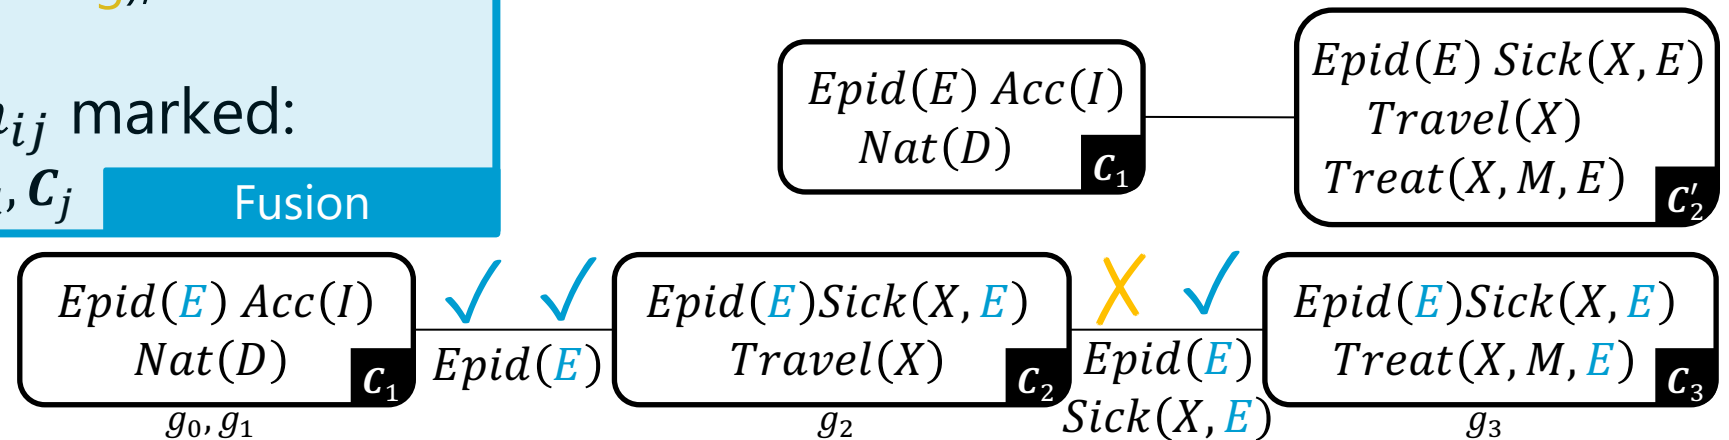


Fusion

- Test each message m_{ij} for each PRV A to eliminate and each separator PRV S
 - If Cond. 1 holds: continue (no groundings)
 - Else if Cond. 2 and Cond. 3 holds: continue
 - Else: mark m_{ij} (grounding); continue with next m_{ij}
- For each message m_{ij} marked:
 - Merge parclusters C_i, C_j

Fusion

- Fusion an additional step after construction to guarantee lifted calculations for liftable models
- E.g., testing marks m_{23}
 - merge C_2, C_3 (as in minimisation)



LJT: Complexity

- Uses also the notion of lifted width $w_T = (w_g, w_{\#})$
 - w_g largest ground width
 - $w_{\#}$ largest counting width
 - As FO jtree constructed from FO dtree, w_T identical between LVE and LJT
 - Fusion may change w_T in terms of the FO jtree
 - But in terms of the LVE calculations in the merged parcluster, w_T is still the same with multiple nodes being combined into one
 - For simplicity, let us consider models that all fulfil Cond. 1 in fusion such that w_T is identical for both LJT and LVE

LJT: Complexity

- LJT complexity based on complexity of LVE:

$$O(n_T \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$$

- Complexity of individual steps

- Construction: linear in number of nodes, no calculations; negligible compared to later steps

- Evidence entering: $O(n_J \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$

- Absorbing evidence complexity: $O(\log_2(n) \cdot r^{w_g-1} \cdot n^{r_{\#}w_{\#}})$

- Visits $\frac{1}{r} \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}}$ lines, possibly exponentiates the potentials

- At each node $\rightarrow n_J \cdot O(\log_2(n) \cdot r^{w_g-1} \cdot n^{r_{\#}w_{\#}})$

- n_J number of nodes in FO jtree J

- For each e evidence parfactors $\rightarrow e \cdot O(n_J \cdot \log_2(n) \cdot r^{w_g-1} \cdot n^{r_{\#}w_{\#}})$

- Assuming $e \ll n_J \rightarrow O(n_J \cdot \log_2(n) \cdot r^{w_g-1} \cdot n^{r_{\#}w_{\#}})$

- First two steps accumulated: $O(n_J \cdot \log_2(n) \cdot r^{w_g-1} \cdot n^{r_{\#}w_{\#}})$

LJT: Complexity

- Complexity of individual steps
 - First two steps accumulated: $O(n_J \cdot \log_2(n) \cdot r^{w_g-1} \cdot n^{r_{\#}w_{\#}})$
 - Message passing: $O(n_J \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - Calculating one message = answering one query on a parcluster
 - Worst-case parfactor size at parcluster: $r^{w_g} \cdot n^{r_{\#}w_{\#}}$
 - Elimination of $|\mathcal{C}_i \setminus \mathcal{S}_{ij}|$ PRVs goes through each line, potentials may be exponentiated $\rightarrow O(\log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - Two messages per edge, $n_J - 1$ edges in $J \rightarrow n_J \cdot O(\log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - Query answering: $O(m \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - Each query answered in one parcluster $\rightarrow O(\log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - With m query terms $\rightarrow m \cdot O(\log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
- All four steps accumulated:
$$O\left((n_J + m) \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}}\right)$$

Comparison to LVE

- LVE complexity of one query = LJT complexity of message passing
 - $O(n_T \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$ vs. $O(n_J \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - Actual number of calculations:
 - In LVE: c_{LVE}
 - For message pass: $2 \cdot c_{LVE}$
- For m queries
 - LVE: $O(m \cdot n_T \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - LJT: $O((n_J + m) \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$
 - Difference in $m \cdot n_T$ vs. $(n_J + m)$
 - LVE has complexity of $O(n_T \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$ for one query
 - LJT only complexity of $O(\log_2(n) \cdot r^{w_g} \cdot n^{r_{\#}w_{\#}})$ for one query

LJT only pays off if $m > 1$, most likely, starting with third query (two queries in LVE = one message pass)

LJT: Implementation

- Available at:
 - <https://www.ifis.uni-luebeck.de/index.php?id=518&L=2>
 - Based on the LVE implementation by Taghipour
 - Available at:
 - <https://dtai.cs.kuleuven.be/software/gcfove>
 - Includes an implementation of the propositional junction tree algorithm for comparison
- Input: BLOG files
 - Based on Bayesian Logic Programming Language
 - <https://bayesianlogic.github.io>

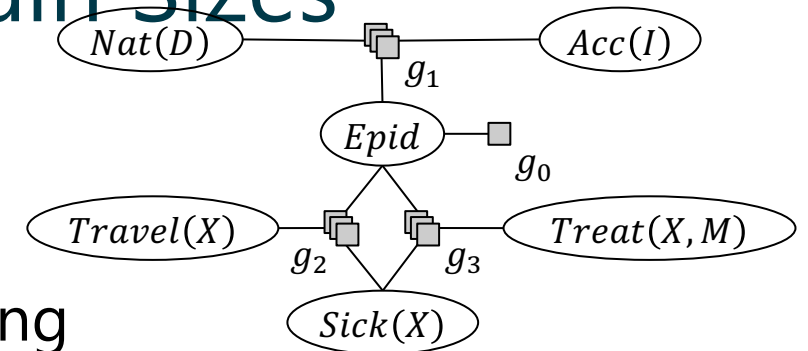
Runtimes: Increasing Domain Sizes

- Example model with all domain sizes $\in \{2,4, \dots, 20, 30, \dots, 100, 200, \dots, 1000\}$
- No evidence
- Queries:
 - $P(\text{Travel}(x_1))$
 - $P(\text{Sick}(x_1))$
 - $P(\text{Treat}(x_1, m_1))$
 - $P(\text{Nat}(d_1))$
 - $P(\text{Man}(w_1))$
 - $P(\text{Epid})$

- Test

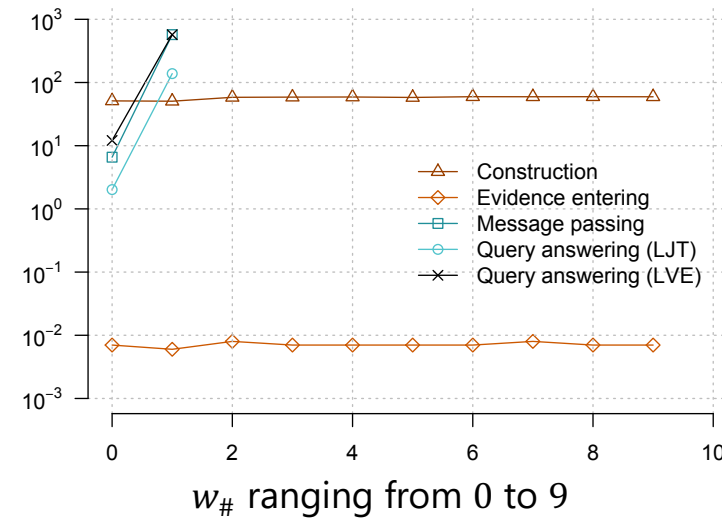


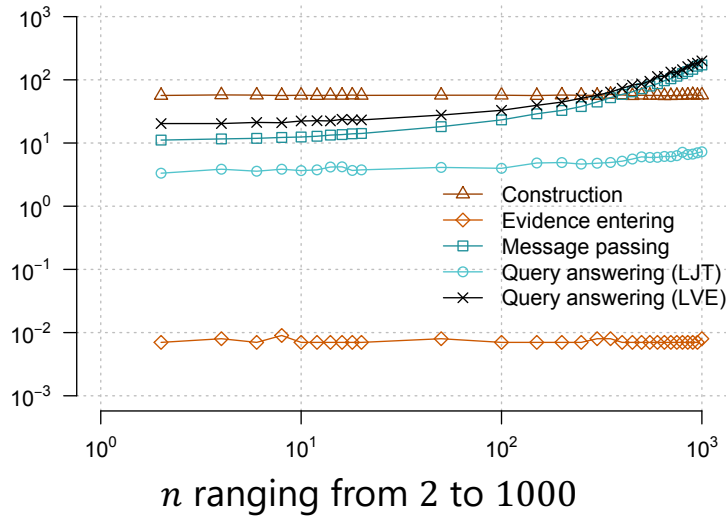
- Increasing

- Ground width w_g
 - Default: 3
 - Counting width $w_\#$
 - Default: 1
 - Number of nodes n_j
 - Default: 3
 - Domain size n
 - Default: 1000
 - Based on $O(n_j \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r \cdot w_\#})$



Step-wise

$$O(n_J \cdot \log_2(n) \cdot r^{w_g} \cdot n^{r \cdot w_{\#}})$$




Runtimes in milliseconds
 Default: $n = 1000, n_J = 3, w_g = 3, w_{\#} = 1$



Changing Inputs

- Known as **adaptive inference**
 - Goal: do not start from scratch
- New queries $\{Q'_i\}_{i=1}^m$
 - Restart query answering step: Answer queries in J
 - JLT supports *online* query answering
 - Queries not known beforehand \rightarrow Stream of queries
- Changed evidence e'
 - Restart with evidence handling: take original local models, handle e' , pass messages, answer queries
- Changed model G'
 - Restart at beginning: Build new FO jtree, ...

If only local changes in e or G , proceed adaptively




Evidence:
~~sick(eve)~~

$\neg travel(eve), \neg sick(eve)$

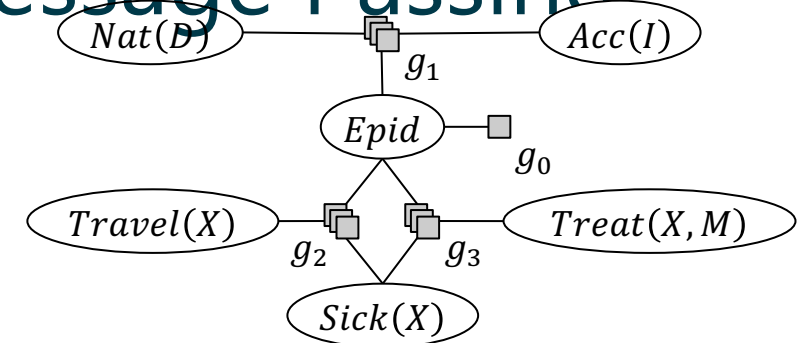
Queries:

~~{ {Epid}, {Travel(eve), Treat(eve, m1)} }~~
~~{ {sick(alice)}, {Treat(eve, m2)} }~~



Changing Inputs and Adaptive Message Passing


- Local changes in G
 - Different potentials in parfactors
 - Changes in domain sizes (*special to relational modelling*)
 - Parfactors are removed or added
 - Maintain FO jtree properties!
 - Only worth it given local changes, otherwise build anew
- Local changes in e
 - Only reset local models of parclusters covered by evidence
- Adaptive message passing
 - If changes in local model or incoming message, calculate new message
 - Otherwise: send empty message
 - Save up to half of the messages



Evidence:
 $sick(eve)$

Queries:

$\{\{Epid\}, \{Travel(eve), Treat(eve, m_1)\}\}$



Interim Summary

- Motivation: Find clusters that are enough for query answering
- FO jtree: From FO dtree clusters to FO jtree parclusters
- LJT algorithm
 - Evidence handling before message passing
 - Propagation/message passing: Dynamic programming
 - Query answering: Find subgraph covering the query terms
- Runtime behaviour
 - Overhead for construction, message passing
 - Savings during query answering
 - Trade-off between those two
- Adaptive inference for local changes: adaptive message passing

Outline: 4. Lifted Inference

Exact Inference

- i. Lifted Variable Elimination for Parfactor Models
 - Idea, operators, algorithm, complexity
- ii. Lifted Junction Tree Algorithm
 - Idea, helper structure: junction tree, algorithm
- iii. First-order Knowledge Compilation for MLNs
 - Idea, helper structure: circuit, algorithm