

MRFs from Knowledge Graphs

Sargur N. Srihari
srihari@buffalo.edu

Topics

- Knowledge Graphs (KGs)
- Statistical Relation Learning (SRL) for KGs
- Latent Feature Models
 - RESCAL, ER-MLP, Latent distance
 - Training SRL
- Markov Random Fields from KGs
- References

Markov Random Fields

- We drop the assumption that the random variables y_{ijk} in Y are conditionally independent
- However, in the case of relational data and without the conditional independence assumption, each y_{ijk} can depend on any of the other $N_e \times N_e \times N_r - 1$ random variables in Y
- Due to this enormous number of possible dependencies, it is intractable to estimate the joint distribution $P(Y)$ without further constraints

MRFs are Dependency graphs

- Each random variable (in our case, a possible fact y_{ijk}) is represented as a node in the graph, while each dependency between random variables is represented as an edge
- To distinguish MRFs from knowledge graphs, we refer to them as *dependency graphs*
 - While knowledge graphs encode the existence of facts, dependency graphs encode statistical dependencies between random variables

Definition of MRF

- MRF has the form

$$P(\underline{\mathbf{Y}}|\boldsymbol{\theta}) = \frac{1}{Z} \prod_c \psi(\mathbf{y}_c|\boldsymbol{\theta})$$

- where $\psi(\mathbf{y}_c|\boldsymbol{\theta}) \geq 0$ is a *potential function* on the c^{th} clique in the dependency graph, and

$$Z = \sum_{\mathbf{y}} \prod_c \psi(\mathbf{y}_c|\boldsymbol{\theta})$$

- is the partition function to ensure distribution sums to one
- potential functions capture local correlations between variables in each clique c
- Defines a distribution over “possible worlds”
 - i.e., over random variables \mathbf{Y}
- Structure of MRF derived using Markov logic
a template language based on logical formulae

Markov Logic Networks

- Combining probability and first-order logic in a single representation
 - PGMs enable efficient handling of uncertainty
 - First-order logic enables compact representation of a wide variety of knowledge

First Order Knowledge Base

- A first-order knowledge base(KB) is a set of sentences or formulas in first order logic
- Formulas are constructed using four types of symbols: constants, variables, functions, and predicates.
- Constant symbols represent objects in the domain of interest (e.g., people: Anna, Bob, Chris, etc.).
- Variable symbols range over objects in the domain
- Function symbols (e.g.,MotherOf) represent mappings from tuples of objects to objects.
- Predicate symbols represent relations among objects in the domain (e.g.,Friends) or attributes of objects(e.g.,Smokes)

First-Order Logic

- An atomic formula or atom is a predicate symbol applied to a tuple of terms e.g., $\text{Friends}(x; \text{MotherOf}(\text{Anna}))$
- Formulas are recursively constructed from atomic formulas using logical connectives and quantifiers.
- If F_1 and F_2 are formulas, the following are also formulas $\neg F_1$ (negation), which is true iff F_1 is false;
- $F_1 \wedge F_2$ (conjunction), which is true iff both F_1 and F_2 are true
- $F_1 \vee F_2$ (disjunction), which is true iff F_1 or F_2 is true; F_1)
- $F_1 \rightarrow F_2$ (implication), which is true iff F_1 is false or F_2 is true
- $F_1 \leftrightarrow F_2$ (equivalence), true iff F_1 and F_2 have same truth value
- $\forall x F_1$ (universal quantification), which is true iff F_1 is true for every object in the domain; and
- $\exists x F_1$ (existential quantification), which is true iff F_1 is true for at least one object x in the domain.

Definition of MLN

- An MLN L is a set of pairs $(F_i; w_i)$, where F_i is a formula in first-order logic and w_i is a real no.
- Together with constants $C = \{c_1; c_2; \dots; c_{|C|}\}$, it defines MRF $M_{L;C} \left(\boxed{P(X=x) = \frac{1}{Z} \exp \left(\sum_j w_j f_j(x) \right)} \right)$ as follows:
 - 1. $M_{L;C}$ contains one binary node for each possible grounding of each predicate appearing in L
 - Value of node is 1 if ground atom is true, and 0 otherwise
 - 2. $M_{L;C}$ contains one feature for each possible grounding of each formula F_i in L
 - The value of this feature is 1 if the ground formula is true, and 0 otherwise. The weight of the feature is the w_i associated with F_i in L

Dependency Graph from Markov Logic

- Given formulae $\mathcal{F} = \{F_i\}_{i=1}^L$, create an edge between nodes if corresponding facts occur in at least one grounded formula
 - A grounding of formula F_i is given by the assignment of entities to the variables in F_i
- Furthermore, we define $\psi(y_c|\theta)$ such that

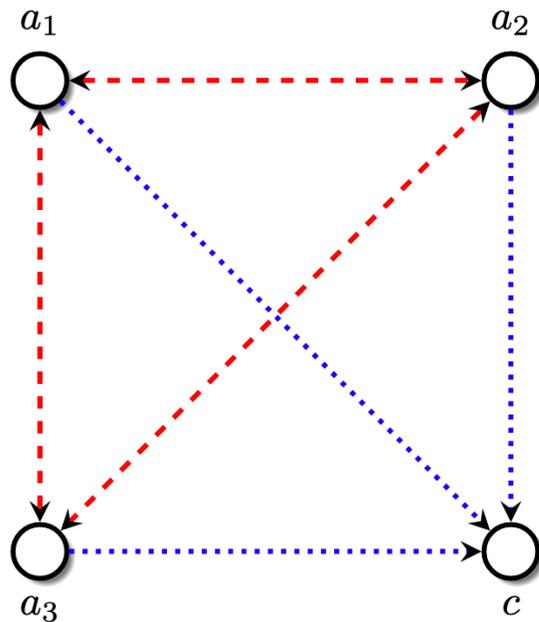
$$P(\mathbf{Y}|\boldsymbol{\theta}) = \frac{1}{Z} \prod_c \exp(\theta_c x_c)$$

- where x_c denotes no. of true groundings of F_c in Y , and θ_c denotes the weight for formula F_c
 - If $\theta_c > 0$, we prefer worlds where formula F_c is satisfied
 - If $\theta_c < 0$, we prefer worlds where formula F_c is violated
 - If $\theta_c = 0$, then formula F_c is ignored

An example

- Two types of entities: adults and children
- Two types of relations: parentOf and marriedTo
 - These relations (edges) are correlated
 - since people who share a common child are often married, while people rarely marry their own children
- In Markov logic, we represent these dependencies using formulae such as:
 - $F_1: (x, \text{parentOf}, z) \wedge (y, \text{parentOf}, z) \Rightarrow (x, \text{marriedTo}, y)$
 - $F_2: (x, \text{marriedTo}, y) \Rightarrow \neg (y, \text{parentOf}, x)$

Simple Knowledge Graph



There are 4 entities (circles): 3 adults (a_1, a_2, a_3) and 1 child c

There are 2 types of edges:

adults may or may not be married to each other, as indicated by the red dashed edges,

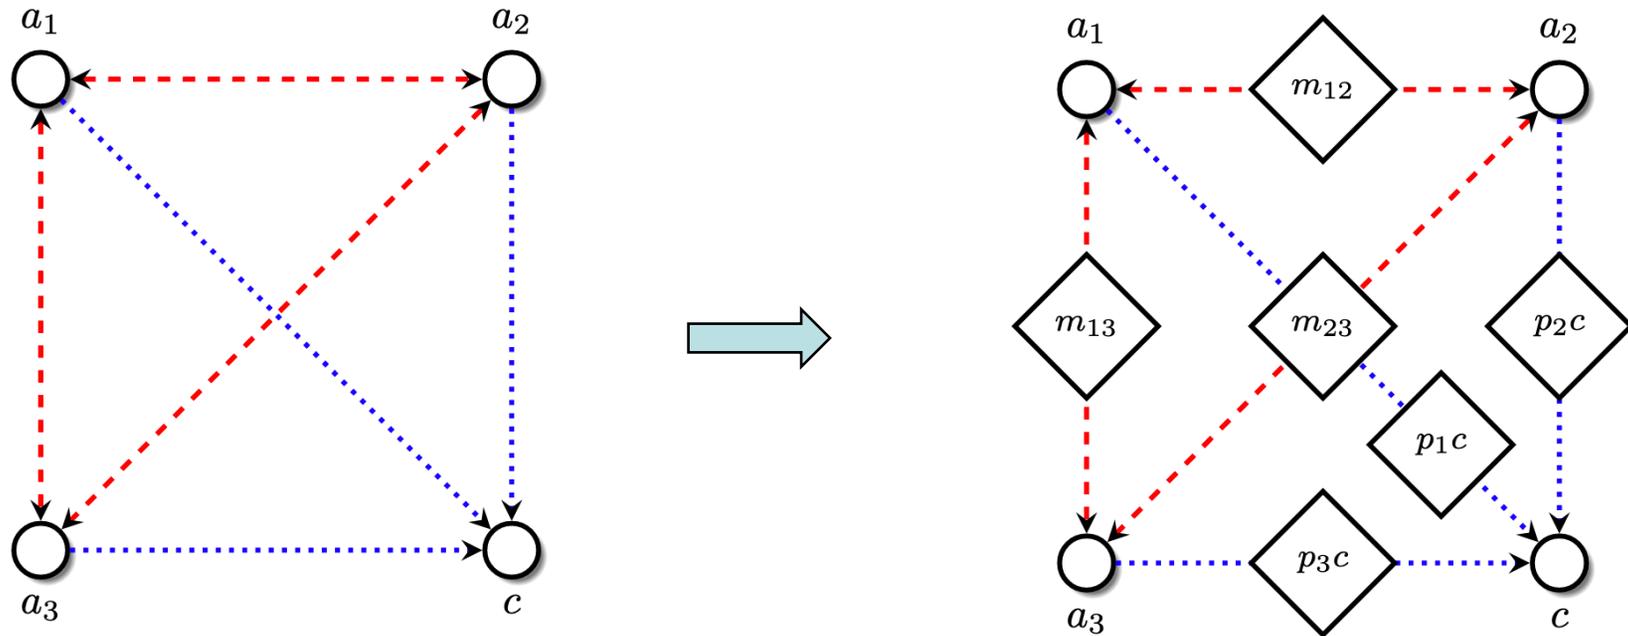
and the adults may or may not be parents of the child, as indicated by the blue dotted edges.

$$F_1: (x, \text{parentOf}, z) \wedge (y, \text{parentOf}, z) \Rightarrow (x, \text{marriedTo}, y)$$

$$F_2: (x, \text{marriedTo}, y) \Rightarrow \neg (y, \text{parentOf}, x)$$

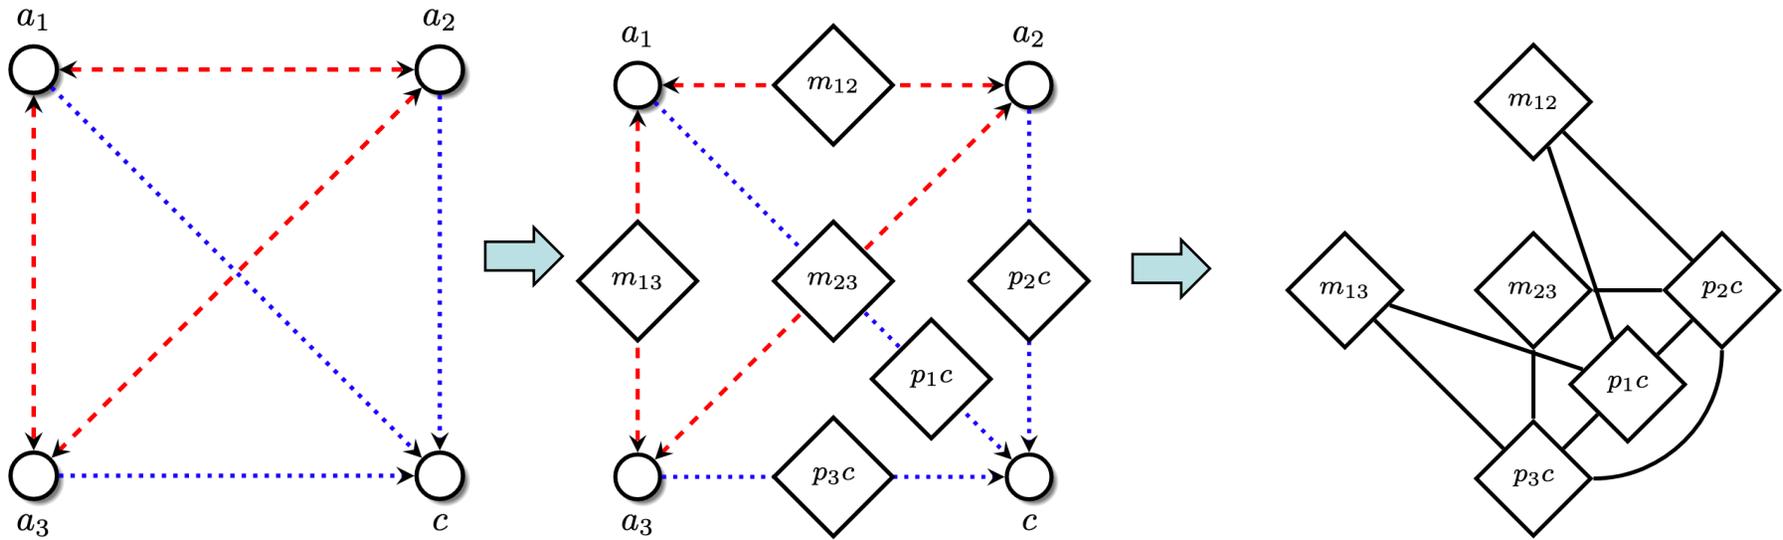
Six possible facts corresponding to the edges

Knowledge Graph relations



- We add binary random variables (represented by diamonds) to each KG edge
 - m : married, p : parent

Knowledge Graph to Dependency Graph



Create edges between these nodes if the corresponding facts occur in grounded formulae

$$F_1: (x, \text{parentOf}, z) \wedge (y, \text{parentOf}, z) \Rightarrow (x, \text{marriedTo}, y)$$

$$F_2: (x, \text{marriedTo}, y) \Rightarrow \neg (y, \text{parentOf}, x)$$

- For instance, grounding F_1 with $x=a_1$, $y=a_3$, and $z=c$, creates the edges $m_{13} \rightarrow p_{1c}$, $m_{13} \rightarrow p_{3c}$, and $p_{1c} \rightarrow p_{3c}$

MRF graph differs from KG

- The process of generating the MRF graph by applying templated rules to a set of entities is known as grounding or instantiation
- The topology of the resulting graph is quite different from the original KG
 - In particular, we have one node per possible KG edge, and these nodes are densely connected
 - This can cause computational difficulties

MRF Inference Tasks

- The inference problem consists of estimating:
 1. The most probable configuration $\mathbf{y}^* = \arg \max_{\mathbf{y}} p(\mathbf{y}|\theta)$
 2. The posterior marginals $p(\mathbf{y}_i|\theta)$
- Both problems are computationally intractable so heuristic approximations must be used
 - Solutions
 - For MAP Task: *Max Product Linear Programming*
 - For posterior marginals: *Gibbs sampling*
 - If potential functions are restricted to disjunctions (OR and NOT, but no AND), then we get a hinge loss MRF (HL-MRF)
 - For which efficient convex algorithms can be applied, based on a continuous relaxation of the binary random variables

MRF Learning Task

- Deals with specifying:
 - Form of potential functions (called *structure learning*)
 - Values for the numerical parameters θ
- In the case of MRFs for KGs, potential functions are often specified in the form of logical rules
 - In this case, structure learning is equivalent to rule learning
 - The parameter estimation problem (which is usually cast as MAP estimation), although convex, is in general quite expensive, since it needs to call inference as a subroutine

Conclusion on MRFs

- Approaches based on MRFs are very flexible
- But harder to make scalable inference and devise learning algorithms
 - Compared to methods based on observable or latent feature models

References

1. M. Nickel, K. Murphy, V. Tresp, E. Gabrilovich, “A Review of Relational Machine Learning for Knowledge Graphs”
Proceedings of the IEEE 104(1): 11-33 (2016)
2. M. Richardson, P. Domingos, “Markov logic networks”, *Machine Learning* 62, 107–136(2006)