

unify notions of lifting from the exact and approximate inference community into a common framework.

2 A Case Study: Markov Logic

The analysis of exchangeability and tractable inference is developed in the context of arbitrary discrete probability distributions, independent of a particular representational formalism. Nevertheless, for the sake of accessibility, we will provide examples and intuitions for Markov logic, a well-known statistical relational language that exhibits several forms of exchangeability. Hence, after the derivation of the theoretical results in each section, we apply the theory to the problem of inference in Markov logic networks. This also allows us to link the theory to existing results from the lifted probabilistic inference literature.

Markov Logic Networks

We first introduce some standard concepts from function-free first-order logic. An *atom* $P(t_1, \dots, t_n)$ consists of a predicate P/n of arity n followed by n argument terms t_i , which are either *constants*, $\{A, B, \dots\}$ or *logical variables* $\{x, y, \dots\}$. A *unary atom* has one argument and a *binary atom* has two. A formula combines atoms with connectives (e.g., \wedge, \Rightarrow). A formula is *ground* if it contains no logical variables. The groundings of a formula are obtained by instantiating the variables with particular constants.

Many statistical relational languages have been proposed in recent years (Getoor and Taskar 2007; De Raedt et al. 2008). We will work with one such language, called *Markov logic networks* (MLN) (Richardson and Domingos 2006). An MLN is a set of tuples (w, f) , where w is a real number representing a weight and f is a formula in first-order logic. Let us, for example, consider the following MLN

$$1.3 \quad \text{Smokes}(x) \Rightarrow \text{Cancer}(x) \quad (1)$$

$$1.5 \quad \text{Smokes}(x) \wedge \text{Friends}(x, y) \Rightarrow \text{Smokes}(y) \quad (2)$$

which states that smokers are more likely to (1) get cancer and (2) be friends with other smokers.

Given a domain of constants \mathbf{D} , a first-order MLN Δ induces a *grounding*, which is the MLN obtained by replacing each formula in Δ with all its groundings (using the same weight). Take for example the domain $\mathbf{D} = \{A, B\}$ (e.g., two people, Alice and Bob), the above first-order MLN represents the following grounding.

$$\begin{aligned} 1.3 \quad & \text{Smokes}(A) \Rightarrow \text{Cancer}(A) \\ 1.3 \quad & \text{Smokes}(B) \Rightarrow \text{Cancer}(B) \\ 1.5 \quad & \text{Smokes}(A) \wedge \text{Friends}(A, A) \Rightarrow \text{Smokes}(A) \\ 1.5 \quad & \text{Smokes}(A) \wedge \text{Friends}(A, B) \Rightarrow \text{Smokes}(B) \\ 1.5 \quad & \text{Smokes}(B) \wedge \text{Friends}(B, A) \Rightarrow \text{Smokes}(A) \\ 1.5 \quad & \text{Smokes}(B) \wedge \text{Friends}(B, B) \Rightarrow \text{Smokes}(B) \end{aligned}$$

This ground MLN contains eight different random variables, which correspond to all groundings of atoms $\text{Smokes}(x)$, $\text{Cancer}(x)$ and $\text{Friends}(x, y)$. This leads to a distribution over 2^8 possible worlds. The weight of each world is the product of the expressions $\exp(w)$, where (w, f) is a ground




Figure 1: An undirected graphical model with 9 finitely exchangeable Bernoulli variables. There are no (conditional) independencies that hold among the variables.

MLN formula and f is satisfied by the world. The probabilities of worlds are obtained by normalizing their weights. Without loss of generality (Jha et al. 2010), we assume that first-order formulas contain no constants.

Lifted Probabilistic Inference

The advent of statistical relational languages such as Markov logic has motivated a new class of *lifted inference* algorithms (Poole 2003). These algorithms exploit the high-level structure and symmetries of the first-order logic formulas to speed up inference (Kersting 2012). Surprisingly, they perform tractable inference even in the absence of conditional independencies. For example, when interpreting the above MLN as an (undirected) probabilistic graphical model, all pairs of random variables in $\{\text{Smokes}(A), \text{Smokes}(B), \dots\}$ are connected by an edge due to the groundings of Formula 2. The model has no conditional or contextual independencies between the Smokes variables. Nevertheless, lifted inference algorithms exactly compute its single marginal probabilities in time linear in the size of the corresponding graphical model (Van den Broeck et al. 2011), scaling up to millions of random variables.

As lifted inference research makes algorithmic progress, the quest for the source of tractability and its theoretical properties becomes increasingly important. For exact lifted inference, most theoretical results are based on the notion of domain-lifted inference (Van den Broeck 2011).

Definition 1 (Domain-lifted). Domain-lifted inference algorithms run in time polynomial in $|\mathbf{D}|$.

Note that domain-lifted algorithms can be exponential in other parameters, such as the number of formulas and predicates. Our current understanding of exact lifted inference is that syntactic properties of MLN formulas permit domain-lifted inference (Van den Broeck 2011; Jaeger and Van den Broeck 2012; Taghipour et al. 2013). We will review these results where relevant. Moreover, the (fractional) automorphisms of the graphical model representation have been related to lifted inference (Niepert 2012b; Bui, Huynh, and Riedel 2012; Noessner, Niepert, and Stuckenschmidt 2013; Mladenov and Kersting 2013). While there are deep connections between automorphisms and exchangeability (Niepert 2012b; 2013; Bui, Huynh, and Riedel 2012; Bui, Huynh, and de Salvo Braz 2012), we refer these to future work.

3 Finite Exchangeability

This section provides some background on the concept of finite partial exchangeability. We proceed by showing that particular forms of finite exchangeability permit tractable inference. For the sake of simplicity and to provide links to statistical relational models such as MLNs, we present the theory for *finite* sets of (upper-case) *binary* random variables $\mathbf{X} = \{X_1, \dots, X_n\}$. However, the theory applies to all distributions over finite valued discrete random variables. Lower-case \mathbf{x} denote an assignments to \mathbf{X} .

We begin with the most basic form of exchangeability.

Definition 2 (Full Exchangeability). A set of variables $\mathbf{X} = \{X_1, \dots, X_n\}$ is *fully exchangeable* if and only if $\Pr(X_1 = x_1, \dots, X_n = x_n) = \Pr(X_1 = x_{\pi(1)}, \dots, X_n = x_{\pi(n)})$ for all permutations π of $\{1, \dots, n\}$.

Full exchangeability is best understood in the context of a finite sequence of binary random variables such as a number of coin tosses. Here, exchangeability means that it is only the number of heads that matters and not their particular order. Figure 1 depicts an undirected graphical model with 9 finitely exchangeable dependent Bernoulli variables.

Finite Partial Exchangeability

The assumption that all variables of a probabilistic model are exchangeable is often too strong. Fortunately, exchangeability can be generalized to the concept of partial exchangeability using the notion of a *sufficient statistic* (Diaconis and Freedman 1980b; Lauritzen et al. 1984; Lauritzen 1988). Particular instances of exchangeability such as full finite exchangeability correspond to particular statistics.

Definition 3 (Partial Exchangeability). Let \mathcal{D}_i be the domain of X_i , and let \mathcal{T} be a finite set. A set of random variables \mathbf{X} is *partially exchangeable* with respect to the statistic $T : \mathcal{D}_1 \times \dots \times \mathcal{D}_n \rightarrow \mathcal{T}$ if and only if

$$T(\mathbf{x}) = T(\mathbf{x}') \text{ implies } \Pr(\mathbf{x}) = \Pr(\mathbf{x}').$$

The following theorem states that the joint distribution of a set of random variables that is partially exchangeable with a statistic T is a unique mixture of uniform distributions.

Theorem 1 (Diaconis and Freedman (1980a)). Let \mathcal{T} be a finite set and let $T : \{0, 1\}^n \rightarrow \mathcal{T}$ be a statistic of a partially exchangeable set \mathbf{X} . Moreover, let $S_t = \{\mathbf{x} \in \mathcal{D}_1 \times \dots \times \mathcal{D}_n \mid T(\mathbf{x}) = t\}$, let U_t be the uniform distribution over S_t , and let $\Pr(S_t) = \Pr(T(\mathbf{x}) = t)$. Then,

$$\Pr(\mathbf{X}) = \sum_{t \in \mathcal{T}} \Pr(S_t) U_t(\mathbf{X}).$$

Hence, a distribution that is partially exchangeable with respect to a statistic T can be parameterized as a unique mixture of uniform distributions. We will see that several instances of partial exchangeability render probabilistic inference tractable. Indeed, the major theme of the present paper can be summarized as finding methods for constructing the above representation and exploiting it for tractable probabilistic inference for a given probabilistic model.

Let $[[\cdot]]$ be the indicator function. The uniform distribution of each equivalence class S_t is $U_t(\mathbf{X}) = [[T(\mathbf{X}) = t]]/|S_t|$;

and the probability of S_t is $\Pr(S_t) = \Pr(\mathbf{x})|S_t|$ for every $\mathbf{x} \in S_t$. Hence, every value of the statistic T corresponds to one equivalence class S_t of joint assignments with identical probability. We will refer to these equivalence classes as *orbits*. We write $\mathbf{x} \sim \mathbf{e}$ when assignments \mathbf{x} and \mathbf{e} agree on the values of their shared variables (Darwiche 2009). The *suborbit* $S_{t,\mathbf{e}} \subseteq S_t$ for some evidence state \mathbf{e} is the set of those states in S_t that are compatible with \mathbf{e} , that is, $S_{t,\mathbf{e}} = \{\mathbf{x} \mid T(\mathbf{x}) = t \text{ and } \mathbf{x} \sim \mathbf{e}\}$.

Partial Exchangeability and Probabilistic Inference

We are now in the position to relate finite partial exchangeability to tractable probabilistic inference, using notions from Theorem 1. The inference tasks we consider are

- *MPE inference*, i.e., finding $\arg \max_{\mathbf{y}} \Pr(\mathbf{y}, \mathbf{e})$ for any given assignment \mathbf{e} to variables $\mathbf{E} \subseteq \mathbf{X}$, and
- *marginal inference*, i.e., computing $\Pr(\mathbf{e})$ for any given \mathbf{e} .

For a set of variables \mathbf{X} , we say that $P(\mathbf{x})$ can be computed *efficiently* iff it can be computed in time polynomial in $|\mathbf{X}|$. We make the following complexity claims

Theorem 2. Let \mathbf{X} be partially exchangeable with statistic T . If we can efficiently

- for all \mathbf{x} , evaluate $\Pr(\mathbf{x})$, and
- for all \mathbf{e} and $t \in \mathcal{T}$ decide whether there exists an $\mathbf{x} \in S_{t,\mathbf{e}}$, and if so, construct it,

then the complexity of MPE inference is polynomial in $|\mathcal{T}|$. If we can additionally compute $|S_{t,\mathbf{e}}|$ efficiently, then the complexity of marginal inference is also polynomial in $|\mathcal{T}|$.

Proof. For MPE inference, we construct an $\mathbf{x}_t \in S_{t,\mathbf{e}}$ for each $t \in \mathcal{T}$, and return the one maximizing $\Pr(\mathbf{x}_t)$. For marginal inference, we return $\sum_{t \in \mathcal{T}} \Pr(\mathbf{x}_t)|S_{t,\mathbf{e}}|$. \square

If the above conditions for tractable inference are fulfilled we say that a distribution is *tractably* partially exchangeable for MPE or marginal inference. We will present notions of exchangeability and related statistics T that make distributions tractably partially exchangeable. Please note that Theorem 2 generalizes to situations in which we can only efficiently compute $\Pr(\mathbf{x})$ up to a constant factor Z , as is often the case in undirected probabilistic graphical models.

Markov Logic Case Study

Exchangeability and independence are not mutually exclusive. Independent and identically distributed (iid) random variables are also exchangeable. Take for example the MLN

$$1.5 \text{ Smokes}(\mathbf{x})$$

The random variables $\text{Smokes}(\text{A}), \text{Smokes}(\text{B}), \dots$ are independent. Hence, we can compute their marginal probabilities independently as

$$\Pr(\text{Smokes}(\text{A})) = \Pr(\text{Smokes}(\text{B})) = \frac{\exp(1.5)}{\exp(1.5) + 1}$$

The variables are also finitely exchangeable. For example, the probability that A smokes and B does not is equal to the probability that B smokes and A does not. The sufficient

statistic $T(\mathbf{x})$ counts *how many* people smoke in the state \mathbf{x} and the probability of a state in which n out of N people smoke is $\exp(1.5n)/(\exp(1.5) + 1)^N$.

Exchangeability can occur without independence, as in the following Markov logic network

$$1.5 \quad \text{Smokes}(\mathbf{x}) \wedge \text{Smokes}(\mathbf{y})$$

This distribution has neither independencies nor conditional independencies. However, its variables are finitely exchangeable and the probability of a state \mathbf{x} is only a function of the sufficient statistic $T(\mathbf{x})$ counting the number of smokers in \mathbf{x} . The probability of a state now increases by a factor of $\exp(1.5)$ with every *pair of smokers*. When n people smoke there are n^2 pairs and, hence, $\Pr(\mathbf{x}) = \exp(1.5n^2)/Z$, where Z is a normalization constant. Let \mathbf{Y} consist of all $\text{Smokes}(\mathbf{x})$ variables except for $\text{Smokes}(A)$, and let \mathbf{y} be an assignment to \mathbf{Y} in which m people smoke. The probability that A smokes given \mathbf{y} is

$$\Pr(\text{Smokes}(A) \mid \mathbf{y}) = \frac{\exp(1.5(m+1)^2)}{\exp(1.5(m+1)^2) + \exp(1.5m^2)},$$

which clearly depends on the number of smokers in \mathbf{y} . Hence, $\text{Smokes}(A)$ is not independent of \mathbf{Y} but the random variables are exchangeable with sufficient statistic n . Figure 1 depicts the graphical representation of the corresponding ground Markov logic network.

4 Exchangeable Decompositions

We now present novel instances of partial exchangeability that render probabilistic inference tractable. These instances generalize exchangeability of single variables to exchangeability of sets of variables. We describe the notion of an exchangeable decomposition and prove that it fulfills the tractability requirements of Theorem 2. We proceed by demonstrating that these forms commonly occur in MLNs.

Variable Decompositions

The notions of independent and exchangeable decompositions are at the core of the developed theoretical results.

Definition 4 (Variable Decomposition). A variable decomposition $\mathcal{X} = \{\mathbf{X}_1, \dots, \mathbf{X}_k\}$ partitions \mathbf{X} into subsets \mathbf{X}_i . We call $w = \max_i |\mathbf{X}_i|$ the *width* of the decomposition.

Definition 5 (Independent Decomposition). A variable decomposition is *independent* if and only if \Pr factorizes as

$$\Pr(\mathbf{X}) = \prod_{i=1}^k Q_i(\mathbf{X}_i).$$

Definition 6 (Exchangeable Decomposition). A variable decomposition is *exchangeable* iff for all permutations π ,

$$\begin{aligned} \Pr(\mathbf{X}_1 = \mathbf{x}_1, \dots, \mathbf{X}_k = \mathbf{x}_k) \\ = \Pr(\mathbf{X}_1 = \mathbf{x}_{\pi(1)}, \dots, \mathbf{X}_k = \mathbf{x}_{\pi(k)}). \end{aligned}$$

Figure 2 depicts an example distribution \Pr with 20 random variables and a decomposition into 5 subsets of width 4. The joint distribution is invariant under permutations of the

$$\Pr \begin{pmatrix} \boxed{1} \boxed{0} \boxed{0} \boxed{0} \\ \boxed{1} \boxed{1} \boxed{0} \boxed{1} \\ \boxed{0} \boxed{1} \boxed{0} \boxed{1} \\ \boxed{1} \boxed{1} \boxed{0} \boxed{1} \\ \boxed{1} \boxed{0} \boxed{0} \boxed{0} \end{pmatrix} = \Pr \begin{pmatrix} \boxed{1} \boxed{0} \boxed{0} \boxed{0} \\ \boxed{1} \boxed{0} \boxed{0} \boxed{0} \\ \boxed{1} \boxed{1} \boxed{0} \boxed{1} \\ \boxed{1} \boxed{1} \boxed{0} \boxed{1} \\ \boxed{0} \boxed{1} \boxed{0} \boxed{1} \end{pmatrix} \quad \begin{array}{c} T(\mathbf{x}) \\ \text{Vector} \quad \text{Count} \\ \boxed{0} \boxed{0} \boxed{0} \boxed{0} \quad 0 \\ \dots \\ \boxed{1} \boxed{0} \boxed{0} \boxed{0} \quad 2 \\ \dots \\ \boxed{1} \boxed{1} \boxed{0} \boxed{1} \quad 2 \\ \dots \\ \boxed{0} \boxed{1} \boxed{0} \boxed{1} \quad 1 \\ \dots \\ \boxed{1} \boxed{1} \boxed{1} \boxed{1} \quad 0 \end{array}$$

Figure 2: An exchangeable decomposition of 20 binary random variables (the boxes) into 5 components of width 4 (the rows). The statistic T counts the occurrences of each unique binary vector.

5 sequences. The corresponding sufficient statistic T counts the number of occurrences of each binary vector of length 4 and returns a tuple of counts.

Please note that the definition of full finite exchangeability (Definition 2) is the special case when the exchangeable decomposition has width 1. Also note that the size of all subsets in an exchangeable decomposition equal the width.

Tractable Variable Decompositions

The core observation of the present work is that variable decompositions that are exchangeable and/or independent result in tractable probabilistic models. For independent decompositions, the following tractability guarantee is used in most existing inference algorithms.

Proposition 3. *Given an independent decomposition of \mathbf{X} with bounded width, and a corresponding factorized representation of the distribution (cf. Definition 5), the complexity of MPE and marginal inference is polynomial in $|\mathbf{X}|$.*

While the decomposition into independent components is a well-understood concept, the combination with finite exchangeability has not been previously investigated as a statistical property that facilitates tractable probabilistic inference. We can now prove the following result.

Theorem 4. *Suppose we can compute $\Pr(\mathbf{x})$ in time polynomial in $|\mathbf{X}|$. Then, given an exchangeable decomposition of \mathbf{X} with bounded width, the complexity of MPE and marginal inference is polynomial in $|\mathbf{X}|$.*

Proof. Following Theorem 2, we have to show that there exists a statistic T so that (a) $|T|$ is polynomial in $|\mathbf{X}|$; (b) we can efficiently decide whether an $\mathbf{x} \in S_{t,\mathbf{e}}$ exists and if so, construct it; and (c) efficiently compute $|S_{t,\mathbf{e}}|$ for all \mathbf{e} and $t \in \mathcal{T}$. Statements (b) and (c) ensure that the assumptions of Theorem 2 hold for exchangeable decompositions, which combined with (a) proves the theorem.

To prove (a), let us first construct a sufficient statistic for exchangeable decompositions. A full joint assignment $\mathbf{X} = \mathbf{x}$ decomposes into assignments $\mathbf{X}_1 = \mathbf{x}_1, \dots, \mathbf{X}_k = \mathbf{x}_k$ in accordance with the given variable decomposition. Each \mathbf{x}_ℓ is a bit string $b \in \{0, 1\}^w$. Consider a statistic $T(\mathbf{x}) = (c_1, \dots, c_{2^w})$, where each c_i has a corresponding unique bit string $b_i \in \{0, 1\}^w$ and $c_i = \sum_{\ell=1}^k [|\mathbf{x}_\ell = b_i|]$. The value c_i

of the statistic thus represents the number of components in the decomposition that are assigned bit string b_i . Hence, we have $\sum_{i=1}^{2^w} c_i = k$, and we prove (a) by observing that

$$|\mathcal{T}| = \binom{k + 2^w - 1}{2^w - 1} \leq k^{2^w - 1} \leq n^{2^w - 1}.$$

To prove statements (b) and (c) we have to find, for each partial assignment $\mathbf{E} = \mathbf{e}$, an algorithm that generates an $\mathbf{x} \in S_{t,\mathbf{e}}$ and that computes $|S_{t,\mathbf{e}}|$ in time polynomial in $|\mathbf{X}|$. To hint at the proof strategy, we give the formula for the orbit without evidence $|S_t|$:

$$|S_{(c_1, \dots, c_{2^w})}| = \binom{k}{c_1} \binom{k - c_1}{c_2} \dots \binom{k - \sum_{i=1}^{2^w-2} c_i}{c_{2^w-1}}.$$

The proof is very technical and deferred to the appendix. \square

Markov Logic Case Study

Let us consider the following MLN

- 1.3 $\text{Smokes}(\mathbf{x}) \Rightarrow \text{Cancer}(\mathbf{x})$
- 1.5 $\text{Smokes}(\mathbf{x}) \wedge \text{Smokes}(\mathbf{y})$

It models a distribution in which every non-smoker or smoker with cancer, that is, every \mathbf{x} satisfying the first formula, increases the probability by a factor of $\exp(1.3)$. Each pair of smokers increases the probability by a factor of $\exp(1.5)$. This model is not fully exchangeable: swapping $\text{Smokes}(\mathbf{A})$ and $\text{Cancer}(\mathbf{A})$ in a state yields a different probability. There are no (conditional) independencies between the $\text{Smokes}(\mathbf{x})$ atoms.

The variables in this MLN do have an exchangeable decomposition whose width is two, namely

$$\mathcal{X} = \{ \{ \text{Smokes}(\mathbf{A}), \text{Cancer}(\mathbf{A}) \}, \\ \{ \text{Smokes}(\mathbf{B}), \text{Cancer}(\mathbf{B}) \}, \\ \{ \text{Smokes}(\mathbf{C}), \text{Cancer}(\mathbf{C}) \}, \dots \}.$$

The sufficient statistic of this decomposition counts the number of people in each of four groups, depending on whether they smoke, and whether they have cancer. The probability of a state only depends on the number of people in each group and swapping people between groups does not change the probability of a state. For example,

$$\begin{aligned} \Pr(\text{Smokes}(\mathbf{A}) = 0, \text{Cancer}(\mathbf{A}) = 1, \\ \text{Smokes}(\mathbf{B}) = 1, \text{Cancer}(\mathbf{B}) = 0) \\ = \Pr(\text{Smokes}(\mathbf{A}) = 1, \text{Cancer}(\mathbf{A}) = 0, \\ \text{Smokes}(\mathbf{B}) = 0, \text{Cancer}(\mathbf{B}) = 1). \end{aligned}$$

Theorem 4 says that this MLN permits tractable inference.

The fact that this MLN has an exchangeable decomposition is not a coincidence. In general, we can show this for MLNs of unary atoms, which are called *monadic* MLNs.

Theorem 5. *The variables in a monadic MLN have an exchangeable decomposition. The width of this decomposition is equal to the number of predicates.*

The proof builds on syntactic symmetries of the MLN, called renaming automorphisms (Bui, Huynh, and Riedel 2012; Niepert 2012a). Please see the appendix for further details.

It now follows as a corollary of Theorems 4 and 5 that MPE and marginal inference in monadic MLNs is polynomial in $|\mathbf{X}|$, and therefore also in the domain size $|\mathbf{D}|$.

Corollary 6. *Inference in monadic MLNs is domain-lifted.*

5 Marginal and Conditional Exchangeability

Many distributions are not decomposable into independent or exchangeable decompositions. Similar to conditional independence, the notion of exchangeability can be extended to conditional exchangeability. We generalize exchangeability to conditional distributions, and state the corresponding tractability guarantees.

Marginal and Conditional Decomposition

Tractability results for exchangeable decompositions on all variables under consideration also extend to subsets.

Definition 7 (Marginal Exchangeability). When a subset \mathbf{Y} of the variables under consideration has an exchangeable decomposition \mathcal{Y} , we say that \mathbf{Y} is *marginally exchangeable*.

This means that \mathcal{Y} is still an exchangeable decomposition when considering the distribution $\Pr(\mathbf{Y}) = \sum_{\mathbf{X} \setminus \mathbf{Y}} \Pr(\mathbf{X})$.

Theorem 7. *Suppose we are given a marginally exchangeable decomposition of \mathbf{Y} with bounded width and let $\mathbf{Z} = \mathbf{X} \setminus \mathbf{Y}$. If computing $\Pr(\mathbf{y}) = \sum_{\mathbf{z}} \Pr(\mathbf{y}, \mathbf{z})$ is polynomial in $|\mathbf{X}|$ for all \mathbf{y} , then the complexity of MPE and marginal inference over variables \mathbf{Y} is polynomial in $|\mathbf{X}|$.*

Proof. Let T be the statistic associated with the given decomposition, and let \mathbf{e} be evidence given on $\mathbf{E} \subseteq \mathbf{Y}$. Then,

$$P(\mathbf{e}) = \sum_{t \in T} |S_{t,\mathbf{e}}| \cdot \sum_{\mathbf{z}} \Pr(\mathbf{y}_{t,\mathbf{e}}, \mathbf{z}), \text{ where } \mathbf{y}_{t,\mathbf{e}} \in S_{t,\mathbf{e}}.$$

By the assumption that \mathbf{Y} is marginally exchangeable and the proof of Theorem 4, we can compute $|S_{t,\mathbf{e}}|$ and an $\mathbf{y}_{t,\mathbf{e}} \in S_{t,\mathbf{e}}$ in time polynomial in $|\mathbf{X}|$. An analogous argument holds for MPE inference on $\Pr(\mathbf{Y})$. \square

We now need to identify distributions $\Pr(\mathbf{Y}, \mathbf{Z})$ for which we can compute $\Pr(\mathbf{y})$ efficiently. This implies tractable probabilistic inference over \mathbf{Y} . Given a particular \mathbf{y} , we have already seen sufficient conditions: when \mathbf{Z} decomposes exchangeably or independently conditioned on \mathbf{y} , Proposition 3 and Theorem 4 guarantee that computing $\Pr(\mathbf{y})$ is tractable. This suggests the following general notion.

Definition 8 (Conditional Decomposability). Let \mathbf{X} be a set of variables with $\mathbf{Y} \subseteq \mathbf{X}$ and $\mathbf{Z} = \mathbf{X} \setminus \mathbf{Y}$. We say that \mathbf{Z} is *exchangeably (independently) decomposable given \mathbf{Y}* if and only if for each assignment \mathbf{y} to \mathbf{Y} , there exists an exchangeable (independent) decomposition $\mathcal{Z}_{\mathbf{y}}$ of \mathbf{Z} .

Furthermore, we say that \mathbf{Z} is decomposable with *bounded width* iff the width w of each $\mathcal{Z}_{\mathbf{y}}$ is bounded. When the decomposition can be computed in time polynomial in $|\mathbf{X}|$ for all \mathbf{y} , we say that \mathbf{Z} is *efficiently decomposable*.

Theorem 8. *Let \mathbf{X} be a set of variables with $\mathbf{Y} \subseteq \mathbf{X}$ and $\mathbf{Z} = \mathbf{X} \setminus \mathbf{Y}$. Suppose we are given a marginally exchangeable decomposition of \mathbf{Y} with bounded width. Suppose further that \mathbf{Z} is efficiently (exchangeably or independently) decomposable given \mathbf{Y} with bounded width. If we can compute $\Pr(\mathbf{x})$ efficiently, then the complexity of MPE and marginal inference over variables \mathbf{Y} is polynomial in $|\mathbf{X}|$.*

Proof. Following Theorem 7, we only need to show that we can compute $\Pr(\mathbf{y}) = \sum_{\mathbf{z}} \Pr(\mathbf{y}, \mathbf{z})$ in time polynomial in $|\mathbf{X}|$ for all \mathbf{y} . When \mathbf{Z} is exchangeably decomposable given \mathbf{Y} , this follows from constructing $\mathcal{Z}_{\mathbf{y}}$ and employing the arguments made in the proof of Theorem 4. The case when \mathbf{Z} is independently decomposable is analogous. \square

Theorems 7 and 8 are powerful results and allow us to identify numerous probabilistic models for which inference is tractable. For instance, we will prove liftability results for Markov logic networks. However, we are only at the beginning of leveraging these tractability results to their fullest extent. Especially Theorem 7 is widely applicable because the computation of $\sum_{\mathbf{z}} \Pr(\mathbf{y}, \mathbf{z})$ can be tractable for many reasons. For instance, conditioned on the variables \mathbf{Y} , the distributions $\Pr(\mathbf{y}, \mathbf{Z})$ could be bounded treewidth graphical models, such as tree-structured Markov networks. Tractability for $\Pr(\mathbf{Y})$ follows immediately from Theorem 7.

Since Theorem 7 only speaks to the tractability of querying variables in \mathbf{Y} , there is the question of when we can also efficiently query the variables \mathbf{Z} . Results from the lifted inference literature may provide a solution by bounding or approximating queries and evidence that includes \mathbf{Z} to maintain marginal exchangeability (Van den Broeck and Darwiche 2013). The next section shows that certain restricted situations permit tractable inference on the variables in \mathbf{Z} .

Markov Logic Case Study

Let us again consider the MLN

- 1.3 $\text{Smokes}(x) \Rightarrow \text{Cancer}(x)$
- 1.5 $\text{Smokes}(x) \wedge \text{Friends}(x, y) \Rightarrow \text{Smokes}(y)$

having the marginally exchangeable decomposition

$$\mathcal{Y} = \{\{\text{Smokes}(A)\}, \{\text{Smokes}(B)\}, \{\text{Smokes}(C)\}, \dots\}$$

whose width is one. To intuitively see why this decomposition is marginally exchangeable, let us consider two states \mathbf{y} and \mathbf{y}' of the $\text{Smokes}(x)$ atoms in which *only* the values of two atoms, for example $\text{Smokes}(A)$ and $\text{Smokes}(B)$, are swapped. There is a symmetry of the MLNs joint distribution that swaps these atoms: the renaming automorphism (Bui, Huynh, and Riedel 2012; Niepert 2012a) that swaps constants A and B in all atoms. For marginal exchangeability, we need that $\sum_{\mathbf{z}} \Pr(\mathbf{y}, \mathbf{z}) = \sum_{\mathbf{z}} \Pr(\mathbf{y}', \mathbf{z})$. But this holds since the renaming automorphism is an automorphism of the set of states $\{\mathbf{yz} \mid \mathbf{z} \in \mathcal{D}_{\mathbf{Z}}\}$ – for every \mathbf{y} , \mathbf{y}' , and \mathbf{z} there exists an automorphism that maps \mathbf{yz} to $\mathbf{y}'\mathbf{z}'$ with $\Pr(\mathbf{y}, \mathbf{z}) = \Pr(\mathbf{y}', \mathbf{z}')$.

The given MLN has several marginally exchangeable decompositions, with the most general one being

$$\mathcal{Y} = \{\{\text{Smokes}(A), \text{Cancer}(A), \text{Friends}(A, A)\}, \\ \{\text{Smokes}(B), \text{Cancer}(B), \text{Friends}(B, B)\}, \\ \{\text{Smokes}(C), \text{Cancer}(C), \text{Friends}(C, C)\}, \dots\}.$$

For that decomposition, the remaining \mathbf{Z} variables

$$\{\text{Friends}(A, B), \text{Friends}(B, A), \text{Friends}(A, C), \dots\}$$

are independently decomposable given \mathbf{Y} . The \mathbf{Z} variables appear at most once in any formula. In a probabilistic graphical model representation, evidence on the \mathbf{Y} variables would therefore decompose the graph into independent components. Thus, it follows from Theorem 8 that we can efficiently answer any query over the variables in \mathbf{Y} .

This insight generalizes to a large class of MLNs, called the *two-variable fragment*. It consists of all MLNs whose formulas contain at most two logical variables.

Theorem 9. *In a two-variable fragment MLN, let \mathbf{Y} and \mathbf{Z} be the ground atoms with one and two distinct arguments respectively. Then there exists a marginally exchangeable decomposition of \mathbf{Y} , and \mathbf{Z} is efficiently independently decomposable given \mathbf{Y} . Each decomposition's width is at most twice the number of predicates.*

The proof of Theorem 9 is rather technical and we refer the reader to the appendix for a detailed proof. It now follows from Theorems 8 and 9 that the complexity of inference over the unary atoms in the two-variable fragment is polynomial in the domain size $|\mathbf{D}|$.

What happens if our query involves variables from \mathbf{Z} – the binary atoms? It is known in the lifted inference literature that we cannot expect efficient inference of general queries that involve the binary atoms. Assignments to the \mathbf{Z} variables break symmetries and therefore break marginal exchangeability. This causes inference to become #P-hard as a function of the query (Van den Broeck and Davis 2012). Nevertheless, if we bound the number of binary atoms involved in the query, we can use the developed theory to show a general liftability result.

Theorem 10. *For any MLN in the two-variable fragment, MPE and marginal inference over the unary atoms and a bounded number of binary atoms is domain-lifted.*

This corresponds to one of the strongest known theoretical results in the lifted inference literature (Jaeger and Van den Broeck 2012). We refer the interested reader to the appendix for the proof. A consequence of Theorem 10 is that we can efficiently compute *all single marginals* in the two-variable fragment, given arbitrary evidence on the unary atoms.

6 Discussion and Conclusion

We conjecture that the concept of (partial) exchangeability has potential to contribute to a deeper understanding of tractable probabilistic models. The important role conditional independence plays in the research field of graphical models is evidence for this hypothesis. Similar to conditional independence, it might be possible to develop a theory of exchangeability that mirrors that of independence.

For instance, there might be a (graphical) structural representations of particular types of partial exchangeability and corresponding logical axiomatizations (Pearl 1988). Moreover, it would be interesting to develop graphical models with exchangeability and independence, and notions like d-separation to detect marginal exchangeability and conditional decomposability from a structural representation. The first author has taken steps in this direction by introducing exchangeable variable models, a class of (non-relational) probabilistic models based on finite partial exchangeability (Niepert and Domingos 2014).

Recently, there has been considerable interest in computing and exploiting the automorphisms of graphical models (Niepert 2012b; Bui, Huynh, and Riedel 2012). There are several interesting connections between automorphisms, exchangeability, and lifted inference (Niepert 2012a). Moreover, there are several group theoretical algorithms that one could apply to the automorphism groups to discover the structure of exchangeable variable decompositions from the structure of the graphical models. Since we presently only exploit renaming automorphisms, there is a potential for tractable inference in MLNs that goes beyond what is known in the lifted inference literature.

Partial exchangeability is related to collective graphical models (Sheldon and Dietterich 2011) (CGMs) and cardinality-based potentials (Gupta, Diwan, and Sarawagi 2007) as these models also operate on sufficient statistics. However, probabilistic inference for CGMs is not tractable and there are no theoretical results that identify tractable CGMs models. The presented work may help to identify such situations. The presented theory generalizes the statistics of cardinality-based potentials.

Lifted Inference and Exchangeability Our case studies identified a deep connection between lifted probabilistic inference and the concepts of partial, marginal and conditional exchangeability. In this new context, it appears that exact lifted inference algorithms (de Salvo Braz, Amir, and Roth 2005; Milch et al. 2008; Jha et al. 2010; Van den Broeck et al. 2011; Gogate and Domingos 2011; Taghipour et al. 2012) can all be understood as performing essentially three steps: (i) construct a sufficient statistic $T(\mathbf{x})$, (ii) generate all possible values of the sufficient statistic, and (iii) count suborbit sizes for a given statistic. For an example of (i), we can show that a compiled first-order circuit (Van den Broeck et al. 2011) or the trace of probabilistic theorem proving (Gogate and Domingos 2011) encode a sufficient statistic in their existential quantifier and splitting nodes. Steps (ii) and (iii) are manifested in all these algorithms through summations and binomial coefficients.

Between Corollary 6 and Theorem 10, we have re-proven almost the entire range of liftability results from the lifted inference literature (Jaeger and Van den Broeck 2012) within the exchangeability framework, and extended these to MPE inference. There is an essential difference though: liftability results make assumptions about the syntax (e.g., MLNs), whereas our exchangeability theorems apply to all distributions. We expect Theorem 8 to be used to show liftability, and more general tractability results for many other repre-

sentation languages, including but not limited to the large number of statistical relational languages that have been proposed (Getoor and Taskar 2007; De Raedt et al. 2008).

Acknowledgments

This work was partially supported by ONR grants #N00014-12-1-0423, #N00014-13-1-0720, and #N00014-12-1-0312; NSF grants #IIS-1118122 and #IIS-0916161; ARO grant #W911NF-08-1-0242; AFRL contract #FA8750-13-2-0019; the Research Foundation-Flanders (FWO-Vlaanderen); and a Google research award to MN.

A Continued Proof of Theorem 4

Proof. To prove statements (b) and (c), we need to represent partial assignments $\mathbf{E} = \mathbf{e}$ with $\mathbf{E} \subseteq \mathbf{X}$. The partial assignments decompose into partial assignments $\mathbf{E}_1 = \mathbf{e}_1, \dots, \mathbf{E}_k = \mathbf{e}_k$ in accordance with \mathcal{X} . Each \mathbf{e}_ℓ corresponds to a string $m \in \{0, 1, *\}^w$ where characters 0 and 1 encode assignments to variables in \mathbf{E}_ℓ and * encodes an unassigned variable in $\mathbf{X}_\ell - \mathbf{E}_\ell$. In this case, we say that \mathbf{e}_ℓ is of type m . Please note that there are 2^w distinct b and 3^w distinct m . We say that \mathbf{x} agrees with \mathbf{e} , denoted by $\mathbf{x} \sim \mathbf{e}$, if and only if their shared variables have identical assignments.

A completion \mathbf{c} of \mathbf{e} to \mathbf{x} is a bijection $\mathbf{c} : \{\mathbf{e}_1, \dots, \mathbf{e}_k\} \rightarrow \{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ such that $\mathbf{c}(\mathbf{e}_i) = \mathbf{x}_j$ implies $\mathbf{e}_i \sim \mathbf{x}_j$. Every completion corresponds to a unique way to assign elements in $\{0, 1\}$ to unassigned variables so as to turn the partial assignment \mathbf{e} into the full assignment \mathbf{x} .

Let $t = (c_1, \dots, c_{2^w}) \in \mathcal{T}$, let $\mathbf{E} \subseteq \mathbf{X}$, and let $\mathbf{e} \in \mathcal{D}_{\mathbf{E}}$. Moreover, let $d_j = \sum_{\ell=1}^k [\mathbf{e}_\ell = m_j]$ for each $m_j \in \{0, 1, *\}^w$. Consider the set of matrices

$$\mathcal{A}_{t,\mathbf{e}} = \left\{ A \in \mathbb{M}(2^w, 3^w) \mid \sum_i a_{i,j} = d_j, \sum_j a_{i,j} = c_i \right. \\ \left. \text{and } a_{i,j} = 0 \text{ if } b_i \not\sim m_j \right\}.$$

Every $A \in \mathcal{A}_{t,\mathbf{e}}$ represents a set of completions from \mathbf{e} to \mathbf{x} for which $T(\mathbf{x}) = t$. The value $a_{i,j}$ indicates that each completion represented by A maps $a_{i,j}$ elements in $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ of type m_j to $a_{i,j}$ elements in $\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ of type b_i . We write $\gamma(A)$ for the set of completions represented by A .

We have to prove the following statements

1. For every $A \in \mathcal{A}_{t,\mathbf{e}}$ and every $\mathbf{c} \in \gamma(A)$ there exists an $\mathbf{x} \in \mathcal{D}_{\mathbf{X}}$ with $T(\mathbf{x}) = t$ and \mathbf{c} is a completion of \mathbf{e} to \mathbf{x} ;
2. For every $\mathbf{x} \in \mathcal{D}_{\mathbf{X}}$ with $T(\mathbf{x}) = t$ and every completion \mathbf{c} of \mathbf{e} to \mathbf{x} there exists an $A \in \mathcal{A}_{t,\mathbf{e}}$ such that $\mathbf{c} \in \gamma(A)$;
3. For all $A, A' \in \mathcal{A}_{t,\mathbf{e}}$ with $A \neq A'$ we have that $\gamma(A) \cap \gamma(A') = \emptyset$;
4. For every $A \in \mathcal{A}_{t,\mathbf{e}}$, we can efficiently compute $|\gamma(A)|$, the size of the set of completions represented by A .

To prove statement (1), let $A \in \mathcal{A}_{t,\mathbf{e}}$ and $\mathbf{c} \in \gamma(A)$. By the definition of $\mathcal{A}_{t,\mathbf{e}}$, \mathbf{c} maps $a_{i,j}$ elements in $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ of type m_j to $a_{i,j}$ elements in $\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ of type b_i . By the conditions $\sum_i a_{i,j} = d_j$ and $\sum_j a_{i,j} = c_i$ of the definition of $\mathcal{A}_{t,\mathbf{e}}$ we have that \mathbf{c} is a bijection. By the condition $a_{i,j} = 0$ if $b_i \not\sim m_j$ of the definition of $\mathcal{A}_{t,\mathbf{e}}$ we have that

$c(\mathbf{e}_i) = \mathbf{x}_j$ implies $\mathbf{e}_i \sim \mathbf{x}_j$ and, therefore, $\mathbf{e} \sim \mathbf{x}$. Hence, c is a completion. Moreover, c completes \mathbf{e} to an \mathbf{x} with $\sum_{\ell=1}^k [[\mathbf{x}_\ell = b_i]] = \sum_j a_{i,j} = c_i$ by the definition of $\mathcal{A}_{t,\mathbf{e}}$. Hence, $T(\mathbf{x}) = t$.

To prove statement (2), let $\mathbf{x} \in \mathcal{D}_{\mathbf{X}}$ with $T(\mathbf{x}) = t$ and let c be a completion of \mathbf{e} to \mathbf{x} . We construct an A with $c \in \gamma(A)$ as follows. Since c is a completion we have that $c(\mathbf{e}_i) = \mathbf{x}_j$ implies $\mathbf{e}_i \sim \mathbf{x}_j$ and, hence, we set $a_{i,j} = 0$ if $b_i \not\sim m_j$. For all other entries in A we set $a_{i,j} = |\{\mathbf{e}_j \mid c(\mathbf{e}_j) = \mathbf{x}_i\}|$. Since c is surjective, we have that $\sum_j a_{i,j} = c_i$ and since c is injective, we have that $\sum_i a_{i,j} = d_j$. Hence, $A \in \mathcal{A}_{t,\mathbf{e}}$.

To prove statement (3), let $A, A' \in \mathcal{A}_{t,\mathbf{e}}$ with $A \neq A'$. Since $A \neq A'$ we have that there exist i, j such that, without loss of generality, $a_{i,j} < a'_{i,j}$. Hence, every $c \in A$ maps fewer elements of type m_j to elements of type b_i than every $c' \in A'$. Hence, $c \neq c'$ for every $c \in A$ and every $c' \in A'$.

To prove statement (4), let $A \in \mathcal{A}_{t,\mathbf{e}}$. Every $c \in A$ maps $a_{i,j}$ elements in $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ of type m_j to $a_{i,j}$ elements in $\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ of type b_i . Hence, the size of the set of completions represented by A is, for each $1 \leq j \leq 3^w$, the number of different ways to place $a_{i,j}$ balls of color i , $1 \leq i \leq 2^w$, into d_j urns. Hence,

$$|\gamma(A)| = \prod_{j=1}^{3^w} \prod_{i=1}^{2^w} \binom{d_j - \sum_{q=1}^{i-1} a_{i,q}}{a_{i,j}}.$$

From the statements (1)-(4) we can conclude that

$$|S_{t,\mathbf{e}}| = \sum_{A \in \mathcal{A}_{t,\mathbf{e}}} |\gamma(A)|.$$

This allows us to prove (b) and (c). We can construct $\mathcal{A}_{t,\mathbf{e}}$ in time polynomial in n as follows. There are $2^w 3^w$ entries in a $2^w \times 3^w$ matrix and each entry has at most k different values. Hence, we can enumerate all $k^{6^w} \leq n^{6^w}$ possible matrices $A \in \mathbb{M}(2^w, 3^w)$. We simply select those A for which the conditions in the definition of $\mathcal{A}_{t,\mathbf{e}}$ hold. For one $A \in \mathcal{A}_{t,\mathbf{e}}$ we can efficiently construct one c and the \mathbf{x} that it completes \mathbf{e} to. This proves (b). Finally, we compute $|S_{t,\mathbf{e}}|$. This proves (c). \square

B Proof of Theorem 5

Proof. Let P_1, \dots, P_N be the N unary predicates of a given MLN and let $\mathbf{D} = \{1, \dots, k\}$ be the domain. After grounding, there are k ground atoms per predicate. We write $P_j(i)$ to denote the ground atom that resulted from instantiating predicate P_j with domain element i . Let $\mathcal{X} = \{\mathbf{X}_1, \dots, \mathbf{X}_k\}$ be a decomposition of the set of ground atoms with $\mathbf{X}_i = \{P_1(i), P_2(i), \dots, P_N(i)\}$ for every $1 \leq i \leq k$. A renaming automorphism (Bui, Huynh, and Riedel 2012; Niepert 2012a) is a permutation of the ground atoms that results from a permutation of the domain elements. The joint distribution over all ground atoms remains invariant under these permutations. Consider the permutation of ground atoms that results from swapping two domain elements $i \leftrightarrow i'$. This permutation acting on the set of ground atoms permutes the components \mathbf{X}_i and $\mathbf{X}_{i'}$ and leaves all other components invariant. Since this is possible for each pair $i, i' \in \{1, \dots, k\}$ it follows that the decomposition \mathcal{X} is exchangeable. \square

C Proof of Theorem 9

Proof. Let P_1, \dots, P_M be the M unary predicates and let Q_1, \dots, Q_N be the N binary predicates of a given MLN and let $\mathbf{D} = \{1, \dots, k\}$ be the domain. After grounding, there are k ground atoms per unary and k^2 ground atoms per binary predicate. We write $P_\ell(i)$ to denote the ground atom that resulted from instantiating unary predicate P_ℓ with domain element i and $Q_\ell(i, j)$ to denote the ground atom that resulted from instantiating binary predicate Q_ℓ with domain elements i and j .

Let \mathbf{X} be the set of all ground atoms, let $\mathbf{Y} = \{P_\ell(i) \mid 1 \leq i \leq k, 1 \leq \ell \leq M\} \cup \{Q_\ell(i, i) \mid 1 \leq i \leq k, 1 \leq \ell \leq N\}$ and let $\mathbf{Z} = \mathbf{X} - \mathbf{Y}$. Moreover, let $\mathcal{Y} = \{\mathbf{Y}_1, \dots, \mathbf{Y}_k\}$ with $\mathbf{Y}_i = \{P_\ell(i) \mid 1 \leq \ell \leq M\} \cup \{Q_\ell(i, i) \mid 1 \leq \ell \leq N\}$. We can make the same arguments as in the proof of Theorem 5 to show that \mathcal{Y} is exchangeable.

Now, we prove that the variables \mathbf{Z} are independently decomposable given \mathbf{Y} . Let $\mathbf{Z}_{i,j} = \{Q_1(i, j), \dots, Q_N(i, j), Q_1(j, i), \dots, Q_N(j, i)\}$ for all $1 \leq i < j \leq k$. Now, let f be any ground formula and let G be the set of ground atoms occurring in both f and \mathbf{Z} . Then, either $G \subseteq \mathbf{Z}_{i,j}$ or $G \cap \mathbf{Z}_{i,j} = \emptyset$ since every formula in the MLN has at most two variables. Hence, $\{\mathbf{Z}_{i,j} \mid 1 \leq i < j \leq k\}$ is a decomposition of \mathbf{Z} with $\binom{k}{2}$ components, width $2N$, and $\Pr(\mathbf{Z}, \mathbf{y})$ factorizes as

$$\Pr(\mathbf{Z}, \mathbf{y}) = \prod_{\substack{i,j \\ i < j}} Q_{i,j}(\mathbf{Z}_{i,j}, \mathbf{y}).$$

By the properties of MLNs, we have that the $Q_{i,j}(\mathbf{Z}_{i,j}, \mathbf{y})$ are computable in time exponential in the width of the decomposition but polynomial in k . \square

D Proof of Theorem 10

Proof. Suppose that the query \mathbf{e} contains a bounded number of binary atoms whose arguments are constants from the set K . Consider the set of variables \mathbf{Q} consisting of all unary atoms whose argument comes from K , and all binary atoms whose arguments both come from K . The unary atoms in \mathbf{Q} are no long marginally exchangeable, because their arguments can appear asymmetrically in \mathbf{e} . We can now answer the query by simply enumerating all states \mathbf{q} of \mathbf{Q} and performing inference in each $\Pr(\mathbf{Y}, \mathbf{Z}, \mathbf{q})$ separately, were all variables \mathbf{Y} have again become marginally exchangeable, and all variables \mathbf{Z} have become independently decomposable given \mathbf{Y} . The construction of \mathbf{Y} and \mathbf{Z} is similar to the proof of Theorem 9, except that some additional binary atoms are now in \mathbf{Y} instead of \mathbf{Z} . These atoms have one argument in K , and one not in K , and are treated as unary. When we bound the number of binary atoms in the query, the size of \mathbf{Q} will not be a function of the domain size, and enumerating over all states \mathbf{q} is domain-lifted. \square

References

Boutilier, C.; Friedman, N.; Goldszmidt, M.; and Koller, D. 1996. Context-specific independence in Bayesian networks. In *Proceedings of the 12th Conference on Uncertainty in artificial intelligence (UAI)*, 115–123.

- Bui, H.; Huynh, T.; and de Salvo Braz, R. 2012. Exact lifted inference with distinct soft evidence on every object. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI)*.
- Bui, H.; Huynh, T.; and Riedel, S. 2012. Automorphism groups of graphical models and lifted variational inference.
- Darwiche, A. 2009. *Modeling and reasoning with Bayesian networks*. Cambridge University Press.
- De Raedt, L.; Frasconi, P.; Kersting, K.; and Muggleton, S., eds. 2008. *Probabilistic inductive logic programming: theory and applications*. Berlin, Heidelberg: Springer-Verlag.
- de Salvo Braz, R.; Amir, E.; and Roth, D. 2005. Lifted first-order probabilistic inference. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1319–1325.
- Diaconis, P., and Freedman, D. 1980a. De Finetti’s generalizations of exchangeability. In *Studies in Inductive Logic and Probability*, volume II.
- Diaconis, P., and Freedman, D. 1980b. Finite exchangeable sequences. *The Annals of Probability* 8(4):745–764.
- Getoor, L., and Taskar, B. 2007. *Introduction to Statistical Relational Learning*. The MIT Press.
- Gogate, V., and Domingos, P. 2011. Probabilistic theorem proving. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI)*, 256–265.
- Gupta, R.; Diwan, A. A.; and Sarawagi, S. 2007. Efficient inference with cardinality-based clique potentials. In *Proceedings of the 24th International Conference on Machine Learning (ICML)*, 329–336.
- Jaeger, M., and Van den Broeck, G. 2012. Liftability of probabilistic inference: Upper and lower bounds. In *Proceedings of the 2nd International Workshop on Statistical Relational AI*.
- Jha, A.; Gogate, V.; Meliou, A.; and Suciu, D. 2010. Lifted inference seen from the other side: The tractable features. In *Proceedings of the 24th Conference on Neural Information Processing Systems (NIPS)*.
- Kersting, K. 2012. Lifted probabilistic inference. In *Proceedings of European Conference on Artificial Intelligence (ECAI)*.
- Koller, D., and Friedman, N. 2009. *Probabilistic Graphical Models*. The MIT Press.
- Lauritzen, S. L.; Barndorff-Nielsen, O. E.; Dawid, A. P.; Diaconis, P.; and Johansen, S. 1984. Extreme point models in statistics. *Scandinavian Journal of Statistics* 11(2).
- Lauritzen, S. L. 1988. *Extremal families and systems of sufficient statistics*. Lecture notes in statistics. Springer-Verlag.
- Lauritzen, S. L. 1996. *Graphical Models*. Oxford University Press.
- Milch, B.; Zettlemoyer, L.; Kersting, K.; Haimes, M.; and Kaelbling, L. 2008. Lifted probabilistic inference with counting formulas. *Proceedings of the 23rd AAAI Conference on Artificial Intelligence* 1062–1068.
- Mladenov, M., and Kersting, K. 2013. Lifted inference via k-locality. In *Proceedings of the 3rd International Workshop on Statistical Relational AI*.
- Niepert, M., and Domingos, P. 2014. Exchangeable variable models. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Niepert, M. 2012a. Lifted probabilistic inference: An MCMC perspective. In *Proceedings of StarAI*.
- Niepert, M. 2012b. Markov chains on orbits of permutation groups. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Niepert, M. 2013. Symmetry-aware marginal density estimation. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI)*.
- Noessner, J.; Niepert, M.; and Stuckenschmidt, H. 2013. RockIt: Exploiting Parallelism and Symmetry for MAP Inference in Statistical Relational Models. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI)*.
- Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.
- Poole, D. 2003. First-order probabilistic inference. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 985–991.
- Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1-2):107–136.
- Robertson, N., and Seymour, P. 1986. Graph minors. II. Algorithmic aspects of tree-width. *Journal of algorithms* 7(3):309–322.
- Sheldon, D., and Dietterich, T. 2011. Collective graphical models. In *Advances in Neural Information Processing Systems (NIPS)*. 1161–1169.
- Taghipour, N.; Fierens, D.; Davis, J.; and Blockeel, H. 2012. Lifted variable elimination with arbitrary constraints. In *Proceedings of the fifteenth international conference on Artificial Intelligence and Statistics*, volume 22, 1194–1202.
- Taghipour, N.; Fierens, D.; Van den Broeck, G.; Davis, J.; and Blockeel, H. 2013. Completeness results for lifted variable elimination. In *Proceedings of the 16th Conference on Artificial Intelligence and Statistics*, 572–580.
- Van den Broeck, G., and Darwiche, A. 2013. On the complexity and approximation of binary evidence in lifted inference. In *Advances in Neural Information Processing Systems 26 (NIPS)*.
- Van den Broeck, G., and Davis, J. 2012. Conditioning in first-order knowledge compilation and lifted probabilistic inference. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI)*.
- Van den Broeck, G.; Taghipour, N.; Meert, W.; Davis, J.; and De Raedt, L. 2011. Lifted probabilistic inference by first-order knowledge compilation. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2178–2185.
- Van den Broeck, G. 2011. On the completeness of first-order knowledge compilation for lifted probabilistic inference. In *Advances in Neural Information Processing Systems 24 (NIPS)*, 1386–1394.