# Relational Marginal Problems: Theory and Estimation 

Ondřej Kuželka<br>Cardiff University, UK<br>KuzelkaO@cardiff.ac.uk

Yuyi Wang<br>ETH Zurich, Switzerland<br>yuwang@ethz.ch

Jesse Davis<br>KU Leuven, Belgium<br>jesse.davis@cs.kuleuven.be

Steven Schockaert<br>Cardiff University, UK<br>SchockaertS1@cardiff.ac.uk


#### Abstract

In the propositional setting, the marginal problem is to find a (maximum-entropy) distribution that has some given marginals. We study this problem in a relational setting and make the following contributions. First, we compare two different notions of relational marginals. Second, we show a duality between the resulting relational marginal problems and the maximum likelihood estimation of the parameters of relational models, which generalizes a well-known duality from the propositional setting. Third, by exploiting the relational marginal formulation, we present a statistically sound method to learn the parameters of relational models that will be applied in settings where the number of constants differs between the training and test data. Furthermore, based on a relational generalization of marginal polytopes, we characterize cases where the standard estimators based on feature's number of true groundings needs to be adjusted and we quantitatively characterize the consequences of these adjustments. Fourth, we prove bounds on expected errors of the estimated parameters, which allows us to lower-bound, among other things, the effective sample size of relational training data.


## Introduction

Statistical Relational Learning (SRL, Getoor and Taskar, eds., 2007) is concerned with learning probabilistic models from relational data. Many popular SRL frameworks, such as Markov Logic Networks (MLNs, Richardson and Domingos 2006), use weighted logical formulas to encode statistical regularities that hold for the considered problem. Typically, the maximum (pseudo-)likelihood weights of the formulas are estimated from training data, which is usually a single large example (e.g. a social network). This is problematic for two reasons. First, the weights that are learned from this single training example are in general not optimal for examples of different sizes (Jain, Kirchlechner, and Beetz 2007). This turns out to be a fundamental problem, which cannot simply be solved by rescaling the weights (Shalizi and Rinaldo 2013). Second, without making further assumptions, it is difficult to provide any statistical guarantees about the learned weights.

In this paper, we approach parameter estimation in SRL from a novel direction, by introducing the notion of a rela-

[^0]tional marginal problem. In the propositional case (Wainwright and Jordan 2008), marginal problems entail finding a maximum-entropy distribution which has the given marginal probabilities. A well-known property of such problems is that they are dual to the maximum-likelihood estimation of the parameters of an undirected graphical model (where "dual" is in the sense of convex optimization). In relational marginal problems, we are similarly looking for a maximum-entropy distribution which satisfies some given statistics - relational marginals. However, we also need to define what these relational marginals are. Thus, first, we describe two different types of relational marginals, which differ in the kinds of statistics that are provided. ${ }^{1}$ The first type is based on relational marginal distributions (Kuželka, Davis, and Schockaert 2017) and the second is based on Halpern-style random substitution semantics (Bacchus et al. 1992). Second, for both types of statistics, we establish a relational counterpart of the duality between maximumlikelihood estimation and max-entropy marginal problems. Interestingly, for the latter model, the corresponding dual is MLNs.

Third, the relational marginal perspective allows us to learn parameters for domains that have different sizes (i.e., number of constants) than the training data. The basic idea to achieve this is simple. We assume that the training data is a sample of the data that we want to model. For example, imagine trying to model all of Facebook based on a sampled subset of Facebook users along with all relations among them. Assuming the sample is a large enough and was obtained in a suitable way (which is not always the case in practice - we discuss this issue later), the parameters of the marginals estimated from the sample should be close to the respective parameters for the whole network. Then, instead of using a model learned by optimizing the likelihood on the training data, we use a model obtained as a solution of the corresponding relational marginal problem with a domain of the required size. We may end up with estimated parameters for which the relational marginal problem has no solution. Therefore, we propose a method for adjusting the estimated parameters that enables a solution and characterize its effect on the estimates. Then we also introduce the relational

[^1]marginal polytopes, which allows us to provide conditions under which the unbiased unadjusted estimate will be valid ("realizable") for domains of any size.

In addition, the relational marginal view of the parameter learning problem, can be thought of as consisting of two decoupled problems: estimation of the parameters of marginals and optimization to obtain the max-entropy distributions. Thus, to better understand parameter learning from relational data, it is important to characterize how accurate the estimates are. Assuming that all subsamples of the data being modeled are sampled with the same probability, we derive bounds on expected error, that is, the expected difference between the parameters obtained from the subsample and parameters that could be theoretically computed if the whole dataset were accessible (e.g. the whole Facebook). From this, we can also obtain lower-bounds on the effective sample size for relational data.

The paper is structured around addressing the following four questions about relational marginal problems:

1. What should the relational marginals be? (Section Two Types of Relational Marginals)
2. What are the max-entropy distributions with the given relational marginals, and how can we find them? (Section Max-Entropy Models)
3. When are relational marginal problems realizable, and how can we adjust them when they are not? How can we adjust learning to account for differences in domain sizes? (Section Realizability)
4. How accurate are the parameter estimates of relational marginals, and what are the links with realizability? (Sections Relational Marginal Polytopes and Estimation)
Proofs of all propositions stated in the main text and details of the duality derivations are contained in the online appendix. ${ }^{2}$

## Preliminaries

This paper considers a function-free first-order logic language $\mathcal{L}$, which is built from a set of constants Const, variables Var and predicates Rel $=\bigcup_{i} \operatorname{Rel}_{i}$, where $\operatorname{Rel}_{i}$ contains the predicates of arity $i$. We assume an untyped language and use the domain size to refer to the $\mid$ Const $\mid$. For $a_{1}, \ldots, a_{k} \in$ Const $\cup \operatorname{Var}$ and $R \in \operatorname{Rel}_{k}$, we call $R\left(a_{1}, \ldots, a_{k}\right)$ an atom. If $a_{1}, . ., a_{k} \in$ Const, this atom is called ground. A literal is an atom or its negation. We use $\operatorname{vars}(\alpha)$ to denote the variables that appear in a formula $\alpha$. The formula $\alpha_{0}$ is called a grounding of $\alpha$ if $\alpha_{0}$ can be obtained by replacing each variable in $\alpha$ with a constant from Const. A formula is called closed if all variables are bound by a quantifier. A possible world $\omega$ is defined as a set of ground atoms. The satisfaction relation $\vDash$ is defined in the usual way.

A substitution is a mapping from variables to terms. An injective substitution is a substitution which does not map any two variables to the same term. As is commonly done in statistical relational learning, we use the unique names assumption, meaning that $c_{1} \neq c_{2}$ whenever $c_{1}$ and $c_{2}$ are different constants. A first-order universally quantified formula

[^2]$\alpha$ is said to be proper if $\alpha \vartheta$ is trivially true whenever $\vartheta$ is not injective. For instance, the formula $\forall X, Y: f r(X, Y)$ is not proper whereas the formula $\forall X, Y: f r(X, Y) \vee X=Y$ is proper. We sometimes omit the universal quantifiers and simply write, e.g. $f r(X, Y) \vee X=Y$.

A (global) example is a pair $(\mathcal{A}, \mathcal{C})$, with $\mathcal{C}$ a set of constants and $\mathcal{A}$ a set of ground atoms which only use constants from $\mathcal{C}$. Let $\Upsilon=(\mathcal{A}, \mathcal{C})$ be an example and $\mathcal{S} \subseteq \mathcal{C}$. The fragment $\Upsilon\langle S\rangle=(\mathcal{B}, \mathcal{S})$ is defined as the restriction of $\Upsilon$ to the constants in $\mathcal{S}$, i.e. $\mathcal{B}$ is the set of all atoms from $\mathcal{A}$ which only contain constants from $\mathcal{S}$. Two examples $\Upsilon_{1}=\left(\mathcal{A}_{1}, \mathcal{C}_{1}\right)$ and $\Upsilon_{2}=\left(\mathcal{A}_{2}, \mathcal{C}_{2}\right)$ are isomorphic, denoted as $\Upsilon_{1} \approx \Upsilon_{2}$, if there exists a bijection $\sigma: \mathcal{C}_{1} \rightarrow \mathcal{C}_{2}$ such that $\sigma\left(\mathcal{A}_{1}\right)=\mathcal{A}_{2}$, where $\sigma$ is extended to ground atoms in the usual way. When $\mathcal{C}$ is a set of constants and $\Phi_{0}$ a set of closed formulas, $\Pi\left(\mathcal{C}, \Phi_{0}\right)$ denotes the set of all $\Upsilon=(\mathcal{A}, \mathcal{C})$ such that $\Upsilon \models \Phi_{0}$ (we can think of $\Phi_{0}$ as a set of constraints).

A Markov logic network (MLN, Richardson and Domingos 2006) is a set of weighted formulas $(\alpha, w), w \in \mathbb{R}$ and $\alpha$ a function-free and quantifier-free first-order formula. The semantics are defined w.r.t. the groundings of the first-order formulas, relative to some finite set of constants $\mathcal{C}$, called the domain. An MLN is seen as a template that defines a Markov random field. Specifically, an MLN $\mathcal{M}$ induces the following probability distribution on the set of global examples ${ }^{3}$ $\Upsilon$ :

$$
\begin{equation*}
p_{\mathcal{M}}(\Upsilon)=\frac{1}{Z} \exp \left(\sum_{(\alpha, w) \in \mathcal{M}} w \cdot n(\alpha, \Upsilon)\right) \tag{1}
\end{equation*}
$$

where $n(\alpha, \Upsilon)$ is the number of groundings of $\alpha$ that are satisfied in $\Upsilon$, and $Z$ is a normalization constant to ensure that $p_{\mathcal{M}}$ is a probability distribution.

Weights of MLNs are typically estimated using maximum (pseudo)likelihood from a given global example. When the MLN with weights learned in this way is used as a probabilistic model with a domain of different size, there are no guarantees regarding the induced distribution. This is most obvious when the MLN contains formulas having different numbers of variables. Then, keeping the weights fixed, the formulas with the highest number of variables often completely dominate the others if we increase the domain size. While some simple cases could be solved by normalizing the counts $n(\alpha, \Upsilon)$, in general this is not the case. Shalizi and Rinaldo (2013) list the example of modelling homophily in social networks; we refer to their paper for details.

## Two Types of Relational Marginals

Typically, parameters for a statistical relational model are estimated from a single example of a relational structure that consists of a large set of ground atoms $\mathcal{A}$. Intuitively, the goal is to learn a probability distribution of such relational structures. The challenge is how to estimate the distribution from a single example. One solution is based on the assumption that the relational structure has repeated regularities. Then, statistics about these regularities can be computed for

[^3]small substructures of the train example and used to construct a distribution over large relational structures. Thus, the next issue is how to construct the fragments and compute statistics on them. Next, we discuss two possible ways to do so, which we will refer to as Model A and Model B.

## Model A

The first approach to constructing fragments is from (Kuželka, Davis, and Schockaert 2017). It repeatedly samples subsets $\mathcal{S} \subseteq \mathcal{C}$ of the constants from the given example $\Upsilon=(\mathcal{A}, \mathcal{C})$ and then builds one training example $\Upsilon\langle\mathcal{S}\rangle$ for each $\mathcal{S}$. However, the training examples must consider isomorphic classes of constants to account for the fact that each fragment will contain different constants. This is accomplished by using the notion of local examples.
Definition 1 (Local example). Let $k \in \mathbb{N}$. A local example of width $k$ is a pair $\omega=(\mathcal{A},\{1, \ldots, k\})$, where $\mathcal{A}$ is a set of ground atoms that contain only constants from the set $\{1,2, \ldots, k\}$. For an example $\Upsilon=(\mathcal{A}, \mathcal{C})$ and $\mathcal{S} \subseteq \mathcal{C}$, we write $\Upsilon[\mathcal{S}]$ for the set of all local examples of width $|\mathcal{S}|$ which are isomorphic to $\Upsilon\langle\mathcal{S}\rangle$.
To distinguish local examples from global examples, we will denote them using lower case Greek letters such as $\omega$ instead of upper case letters such as $\Upsilon$.
Example 1. Let $\Upsilon=(\{f r($ alice, bob $), f r(b o b$, alice $)$, fr(bob, eve), fr(eve, bob), sm(alice) \}, \{alice, bob, eve\}), i.e. the only smoker is alice and the friendship structure is:


Then $\Upsilon\langle\{$ alice, bob $\}\rangle=(\{$ sm(alice $), \quad$ fr $($ alice, bob $)$, $f r($ bob, alice $)\},\{$ alice, bob $\}), \Upsilon[\{$ alice, bob $\}]=\{(\{f r(1,2)$, $f r(2,1), \operatorname{sm}(1)\},\{1,2\}),\{(\{f r(2,1), f r(1,2), \operatorname{sm}(2)\}$,
$\{1,2\})\}$.
This leads to a natural definition of a probability distribution over local examples of width $k$.
Definition 2 (Relational marginal distribution of a global example). Let $\Upsilon=(\mathcal{A}, \mathcal{C})$ be an example and $k \in \mathbb{N}$. The relational marginal distribution of $\Upsilon$ of width $k$ is a distribution $P_{\Upsilon, k}$ over local examples, where $P_{\Upsilon, k}(\omega)$ is defined as the probability that $\omega$ is sampled by the following process: (i) Uniformly sample a subset $\mathcal{S}$ of $k$ constants from $\mathcal{C}$. (ii) Uniformly sample a local example $\omega$ from the set $\Upsilon[\mathcal{S}]$. For a closed formula $\alpha$ without constants, we also define, its probability: $P_{\Upsilon, k}(\alpha)=\sum_{\omega: \omega \models \alpha} P_{\Upsilon, k}(\omega)$.
We will call a pair $(\alpha, p)$, where $\alpha$ is a constant-free closed formula and $p \in[0 ; 1]$, a relational marginal constraint. We may also interpret the probability $P_{\Upsilon, k}(\alpha)$ of a closed constant-free formula $\alpha$ as the probability that $\alpha$ is true in a restriction $\Upsilon\langle\mathcal{S}\rangle$ of $\Upsilon$ to a randomly sampled subset $\mathcal{S}$ of $k$ constants from $\Upsilon$. Thus, if we are only interested in the probabilities of closed constant-free formulas, we do not have to refer to local examples. Local examples are important because relational marginal distributions defined using them are themselves probability distributions on possible worlds, which is both nice conceptually and convenient (as it means that we can model relational marginals of Model A using
any standard propositional probabilistic model). Arguably, this property also makes Model A more interpretable and hence easier to explain to non-specialists.

Global examples may also be assumed to be sampled from some distribution and we define the corresponding marginal distributions induced by such distributions accordingly. When $P(\Upsilon)$ is a distribution over finite global examples from a possibly countably infinite set $\Omega$, then the marginal distribution of width $k$ is a distribution $P_{k}$ over local examples where $P_{k}(\omega)$ is defined as $P_{k}(\omega)=\sum_{\Upsilon \in \Omega} P(\Upsilon)$. $P_{\Upsilon, k}(\omega)$. For a closed formula $\alpha$ without constants, we also analogically define: $P_{k}(\alpha)=\sum_{\omega: \omega \models \alpha} P_{k}(\omega)$. In other words, a relational marginal distribution is a mixture of (possibly countably many) relational marginal distributions of global examples.
Proposition 2. Let $P(\Upsilon)$ be a distribution on domain size $n$ and $k \leq n$ be an integer. Let $\Omega_{\alpha}=\{\Upsilon\langle\mathcal{S}\rangle: \mathcal{S} \subseteq \mathcal{C},|\mathcal{S}|=$ $k, \Upsilon\langle\mathcal{S}\rangle \vDash \alpha\}$ where $\Upsilon=(\mathcal{A}, \mathcal{C})$ is sampled according to the distribution $P(\Upsilon)$. Then, for a closed constant-free formula $\alpha, \widehat{p}_{\alpha}=\left|\Omega_{\alpha}\right| \cdot\binom{n}{k}^{-1}$ is an unbiased estimate of $P_{k}(\alpha)$ for Model $A$.

## Model B

The second approach is to consider random substitutions, which is in the spirit of existing works (Bacchus et al. 1992; Schulte et al. 2014). Here, the statistics that we collect about $\Upsilon$ are defined as follows.
Definition 3 (Probability of formulas under Model B). Let $\Upsilon=(\mathcal{A}, \mathcal{C})$ be a global example and $\alpha$ be a universally quantified formula. Let $P_{\vartheta}$ be a uniform distribution on injective substitutions from the set $\Theta_{\alpha}=\{\vartheta \mid \vartheta: \operatorname{vars}(\alpha) \rightarrow$ $\mathcal{C}$ and $\vartheta$ is injective $\}$. Then the probability $Q(\alpha)$ of the formula $\alpha$ under model $B$ is defined as

$$
Q_{\Upsilon}(\alpha)=\sum_{\vartheta \in \Theta_{\alpha}} \mathbb{1}(\Upsilon \models \alpha \vartheta) P_{\vartheta}(\vartheta)=\frac{1}{|\Theta|} \sum_{\vartheta \in \Theta_{\alpha}} \mathbb{1}(\Upsilon \models \alpha \vartheta) .
$$

Just like for Model A, we extend the definition of the probability of formulas straightforwardly to the case where $\Upsilon$ is not fixed but sampled from some distribution over a countable set $\Omega: Q(\alpha)=\sum_{\Upsilon \in \Omega} Q_{\Upsilon}(\alpha) \cdot P(\Upsilon)$.
Example 3. Let $\Upsilon$ be as in Example 1. Let

$$
\begin{array}{r}
\alpha=\forall X, Y: \neg f r(X, Y) \vee \operatorname{sm}(Y), \\
\beta=\forall X, Y: \neg f r(X, Y) \vee \operatorname{sm}(X) \vee \operatorname{sm}(Y) .
\end{array}
$$

Assume that the relation $f r(.,$.$) , "friends", is symmetric.$ Then, the formula $\alpha$ is classically true if all people who are friends with someone, smoke. The formula $\beta$ is classically true if for every pair of people $A$ and $B$ who are friends, at least one of them smokes. Computing the respective probabilities, we get $P_{\Upsilon, 2}(\alpha)=\frac{1}{3}, P_{\Upsilon, 2}(\beta)=\frac{2}{3}$, $Q_{\Upsilon}(\alpha)=\frac{1}{2}, Q_{\Upsilon}(\beta)=\frac{2}{3}$, which illustrates that in general the "marginal" probabilities given by the two models will differ. The first model might be slightly easier to interpret as the marginal probabilities $P(\gamma)$ correspond to the fraction of the width- $k$ fragments of $\Upsilon$ in which $\gamma$ is true as a classical logic formula.

We note that the straightforward analogue of Proposition 2 also holds for Model B.

## Max-Entropy Models

In this section we show how to compute models of given relational marginals under Model A and Model B.
Definition 4 (Model of relational marginals). Let us have a set of pairs $\Phi=\left\{\left(\alpha_{1}, \theta_{1}\right), \ldots,\left(\alpha_{h}, \theta_{h}\right)\right\}$ of relational marginals, where $\alpha_{i}$ is a closed formula and $\theta_{i} \in[0 ; 1]$. Let $\Phi_{0}$ be a set of formulas, called hard rules (i.e. formulas that are supposed to always hold). We say that a probability distribution $P(\Upsilon)$ over worlds satisfying the hard rules from $\Phi_{0}$ is a width- $k$ model of $\Phi$ iff $P_{k}\left(\alpha_{i}\right)=\theta_{i}\left(Q_{k}\left(\alpha_{i}\right)=\theta_{i}\right.$, respectively) for all $\left(\alpha_{i}, \theta_{i}\right) \in \Phi$.

We will use standard duality arguments from convex optimization (Boyd and Vandenberghe 2004), essentially following (Singh and Vishnoi 2014). For both models, the result will in both cases an exponential-family model and the one for Model B will turn out to be equivalent to MLNs.

## Model A

Let $\mathcal{C}$ be a set of constants, $\mathcal{A}$ be the set of all atoms over $\mathcal{C}$ based on some given first-order language and let $\mathcal{C}_{k}$ denote the set of all $k$-element subsets of $\mathcal{C}$. Next, let $\Phi$ be a set of relational marginals and $\Phi_{0}$ be a set of hard constraints, i.e. formulas $\alpha$ such that if $\Upsilon \not \vDash \alpha$ then $P(\Upsilon)=0$. We assume that there exists at least one distribution $P^{\prime}$ which is a model of $\Phi$ and which satisfies the following positivity condition: $P^{\prime}(\Upsilon)>0$ for all $\Upsilon$ satisfying the hard constraints (i.e. those for which $\Upsilon \models \Phi_{0}$ ).

$$
\begin{gather*}
\sup _{\left\{P_{\Upsilon}: \Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)\right\}} \sum_{\Upsilon \in \Pi_{n}\left(\mathcal{C}, \Phi_{0}\right)} P_{\Upsilon} \log \frac{1}{P_{\Upsilon}} \text { s.t. }  \tag{2}\\
\forall\left(\alpha_{i}, \theta_{i}\right) \in \Phi: \\
\frac{1}{\left|\mathcal{C}_{k}\right|} \sum_{\mathcal{S} \in \mathcal{C}_{k}} \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} \mathbb{1}\left(\Upsilon\langle\mathcal{S}\rangle \models \alpha_{i}\right) \cdot P_{\Upsilon}=\theta_{i}  \tag{3}\\
\forall \Upsilon \in \Pi_{n}\left(\mathcal{C}, \Phi_{0}\right): P_{\Upsilon} \geq 0, \quad \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} P_{\Upsilon}=1 \tag{4}
\end{gather*}
$$

Next we define $\#_{k}(\alpha, \Upsilon)=\left|\left\{\mathcal{S} \in \mathcal{C}_{k}: \Upsilon\langle\mathcal{S}\rangle \models \alpha\right\}\right|$ as the number of sets $\mathcal{S} \in \mathcal{C}_{k}$ such that the formula $\alpha$ is classically true in the restriction of $\Upsilon$ to constants in $\mathcal{S}$. In the appendix, we show that the distribution, if it exists, that is a solution to the optimization problem has the following form:

$$
\begin{equation*}
P_{\Upsilon}=\frac{1}{Z} \exp \left(\sum_{\alpha_{i} \in \Phi} \lambda_{i} \frac{\#_{k}\left(\alpha_{i}, \Upsilon\right)}{\left|\mathcal{C}_{k}\right|}\right) \tag{5}
\end{equation*}
$$

The Lagrangian dual problem of the maximum entropy problem is to maximize (where $\lambda_{i} \in \mathbb{R}$ ):

$$
\begin{equation*}
\sum_{\alpha_{i} \in \Phi} \lambda_{i} \theta_{i}-\log \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} e^{\sum_{\alpha_{i} \in \Phi} \lambda_{i} \frac{\# k_{k}\left(\alpha_{i}, \Upsilon\right)}{\left|\mathcal{C}_{k}\right|}} \tag{6}
\end{equation*}
$$

Due to the positivity assumption, Slater's condition (Boyd and Vandenberghe 2004) is satisfied and strong duality holds. Consequently, instead of solving the original problem, which has an intractable number of constraints and variables (one variable for each world $\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)$ ), we can solve the dual problem, which only has $|\Phi|$ variables. On the other hand, the optimization criterion of the dual problem may still be computationally hard to solve as it requires weighted counting over worlds in $\Pi\left(\mathcal{C}, \Phi_{0}\right)$. However, in many restricted, but non-trivial, cases, we can exploit lifted weighted model counting techniques in the same way as they were used for maximum-likelihood estimation in Van Haaren et al. (2016).

Let us perform a change of variables $w_{i}:=\lambda_{i} /\left|\mathcal{C}_{k_{i}}\right|$. This gives us

$$
\begin{equation*}
P(\Upsilon)=\frac{1}{Z} \exp \left(\sum_{\alpha_{i} \in \Phi} w_{i} \cdot \#_{k}\left(\alpha_{i}, \Upsilon\right)\right) \tag{7}
\end{equation*}
$$

for the distribution and

$$
\begin{equation*}
\sum_{\alpha_{i} \in \Phi} w_{i} \theta_{i}\left|\mathcal{C}_{k_{i}}\right|-\log \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} e^{\sum_{\alpha_{i} \in \Phi} w_{i} \cdot \#_{k}\left(\alpha_{i}, \Upsilon\right)} \tag{8}
\end{equation*}
$$

for the optimization criterion of the dual problem. Assuming that the marginals were estimated from a global example $\widehat{\Upsilon} \in \Pi\left(\mathcal{C}, \Phi_{0}\right)$ (note that here the domain $\mathcal{C}$ is the same as the domain of the global examples over which the distribution is defined) and that they still satisfy the positivity assumption, we can also rewrite (8) as

$$
\sum_{\alpha_{i} \in \Phi} w_{i} \cdot \#_{k}\left(\alpha_{i}, \widehat{\Upsilon}\right)-\log \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} e^{\sum_{\alpha_{i} \in \Phi} w_{i} \cdot \#_{k}\left(\alpha_{i}, \Upsilon\right)}
$$

It is straightforward to check that this is the same as directly optimizing the log-likelihood of $\widehat{\Upsilon}$. Thus, this is a relational analogue of the well-known duality of maximum likelihood and maximum entropy (Wainwright and Jordan 2008). Note that it is important for the duality of maximum likelihood and maximum entropy that both the $\widehat{\Upsilon}$, from which we estimated the parameters, and the global examples over which the distribution is computed have the same domain size. The Section Realizability will address cases where the domain sizes of the training and testing data differ.

## Model B

Like for Model A, we can construct a convex optimization problem to obtain a maximum-entropy distribution with the given relational marginals under Model B. This problem's optimization criterion is the same as (2). To obtain the constraints enforcing the marginals, we can replace the summation over subsets of constants in $\mathcal{C}$ in Equation (3) by a summation over substitutions from $\Theta_{\alpha_{i}}$, where $\Theta_{\alpha_{i}}$ is defined as in Definition 3. This yields the following set of constraints for all $\left(\alpha_{i}, \theta_{i}\right) \in \Phi$ :

$$
\begin{equation*}
\sum_{\vartheta \in \Theta_{\alpha_{i}}} \frac{1}{\left|\Theta_{\alpha_{i}}\right|} \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} \mathbb{1}\left(\Upsilon \models \alpha_{i} \vartheta\right) \cdot P_{\Upsilon}=\theta_{i} \tag{9}
\end{equation*}
$$

Using basically the same reasoning as for Model A , we arrive at the following form of the probability distribution

$$
\begin{equation*}
P(\Upsilon)=\frac{1}{Z} \exp \left(\sum_{\alpha_{i} \in \Phi} w_{i} \cdot n\left(\alpha_{i}, \Upsilon\right)\right) \tag{10}
\end{equation*}
$$

where $n\left(\alpha_{i}, \Upsilon\right)$ is the number of groundings $\alpha_{i} \vartheta$ of the formula $\alpha_{i}$, where all $\vartheta$ 's are injective, which are true in $\Upsilon$. This distribution is identical to the one for MLNs which only contain proper formulas (because of the injectivity requirement in the definition of Model B). The only difference with the distribution of Model A is the use of $n\left(\alpha_{i}, \Upsilon\right)$ instead of $\#_{k}\left(\alpha_{i}, \Upsilon\right)$. The dual optimization criterion for Model B then becomes to maximize

$$
\begin{equation*}
\sum_{\alpha_{i} \in \Phi} w_{i} \theta_{i}\left|\Theta_{\alpha_{i}}\right|-\log \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} e^{\sum_{\alpha_{i} \in \Phi} w_{i} \cdot n\left(\alpha_{i}, \Upsilon\right)} \tag{11}
\end{equation*}
$$

which can be rewritten, when $\theta_{i}$ 's are estimated from some $\widehat{\Upsilon}=(\widehat{\mathcal{A}}, \widehat{\mathcal{C}})$ with $|\widehat{\mathcal{C}}|=|\widehat{\mathcal{C}}|$, as $\sum_{\alpha_{i} \in \Phi} w_{i} \cdot n\left(\alpha_{i}, \widehat{\Upsilon}\right)-$ $\log \sum_{\Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)} e^{\sum_{\alpha_{i} \in \Phi} w_{i} \cdot n\left(\alpha_{i}, \Upsilon\right)}$. This is the same as loglikelihood of $\widehat{\Upsilon}$ w.r.t. the MLN given by (11), assuming that all the formulas in the MLN are proper. Thus, when the size of the domain of the training example $\widehat{\Upsilon}$ is the same as the cardinality of the domain of the modeled distribution, the max-entropy relational marginal problem in Model B is the same as maximum-likelihood estimation in MLNs.

## Realizability

Not all relational marginals have a model for a given domain size (or even any model at all). This is a problem if we want to estimate relational marginals on some given global example $\widehat{\Upsilon}$ and then use them to obtain a distribution on global examples that have a different domain size (which we address in Section Estimation). The duality between maximum-likelihood estimation and max-entropy relational marginal problems discussed in the previous section only holds when the training data and the distribution that we want to model have the same domain size.

In this section, we will show how to obtain relational marginals for any domain size. Accomplishing this requires replacing the consistency of relational marginal estimators with a weaker notion. However, we can still bound the difference between the unbiased and the adjusted estimates. More importantly, the difference will tend to zero as the domain size of the examples from which the respective marginals were estimated increases.

First, it is easy to see that if a model exists for a given set of marginals and domain size, then there is also such a model for smaller domain sizes, as the next proposition asserts. ${ }^{4}$
Proposition 4. For Model A, if there is a width-k model of $\Phi$ on domain of size $n$ then there is also a width- $k$ model of $\Phi$ on domain of size $m$ for any $m$ satisfying $n \geq m \geq k$. For Model B, if there is a model of $\Phi$ on domain of size $n$

[^4]then there is also a model of $\Phi$ on domain of size $m$ for any $m$ satisfying $n \geq m \geq \max _{\alpha \in \Phi}|\operatorname{vars}(\alpha)|$.
Next, we give an example of a relational marginal distributions that does not have a model for arbitrary domain cardinalities.
Example 5. Let $\mathcal{L}$ consist of predicate symbols $e / 2, r / 1, g / 1, b / 1$ and $\Upsilon=(\mathcal{A}, \mathcal{C})$ be a global example where
\[

$$
\begin{gathered}
\mathcal{A}=\left\{e\left(v_{1}, v_{2}\right), e\left(v_{2}, v_{1}\right), e\left(v_{2}, v_{3}\right), e\left(v_{3}, v_{2}\right), e\left(v_{1}, v_{3}\right),\right. \\
\left.e\left(v_{3}, v_{1}\right), r\left(v_{1}\right), g\left(v_{2}\right), b\left(v_{3}\right)\right\}, \quad \mathcal{C}=\left\{v_{1}, v_{2}, v_{3}\right\}
\end{gathered}
$$
\]

Let $k=2$ be the width of local examples. And let $F\left(X_{1}, X_{2}\right) \stackrel{\text { def }}{=} X_{1} \neq X_{2} \wedge \neg e\left(X_{1}, X_{1}\right) \wedge e\left(X_{1}, X_{2}\right) \wedge$ $e\left(X_{2}, X_{1}\right) \wedge \neg e\left(X_{2}, X_{2}\right)$. Then we can estimate, for instance, the following marginals from $\Upsilon$ under Model $A$ :

$$
\begin{array}{r}
P\left[\exists X_{1}, X_{2}: F\left(X_{1}, X_{2}\right) \wedge r\left(X_{1}\right) \wedge \neg g\left(X_{1}\right) \wedge \neg b\left(X_{1}\right) \wedge\right. \\
\left.\wedge \neg r\left(X_{2}\right) \wedge g\left(X_{2}\right) \wedge \neg b\left(X_{2}\right)\right]=\frac{1}{3} \\
P\left[\exists X_{1}, X_{2}: F\left(X_{1}, X_{2}\right) \wedge r\left(X_{1}\right) \wedge \neg g\left(X_{1}\right) \wedge \neg b\left(X_{1}\right) \wedge\right. \\
\left.\wedge \neg r\left(X_{2}\right) \wedge \neg g\left(X_{2}\right) \wedge b\left(X_{2}\right)\right]=\frac{1}{3} \\
P\left[\exists X_{1}, X_{2}: F\left(X_{1}, X_{2}\right) \wedge \neg r\left(X_{1}\right) \wedge g\left(X_{1}\right) \wedge \neg b\left(X_{1}\right) \wedge\right. \\
\left.\wedge \neg r\left(X_{2}\right) \wedge \neg g\left(X_{2}\right) \wedge b\left(X_{2}\right)\right]=\frac{1}{3}
\end{array}
$$

(which can also be rewritten as probabilities of universally quantified formulas by negating the existentially quantified conjunctions). The global example $\Upsilon$ can be imagined as a complete directed graph (without self-loops) on 3 vertices $v_{1}, v_{2}, v_{3}$ where each of the vertices is colored by one of the "colors" r, $g$, $b$.

We claim that no distribution on global examples satisfies the above marginal probabilities for domain size greater than 3. This can be shown as follows. Using the intuitive view of $\Upsilon$ as a colored directed graph, the distribution on global examples of domain size e.g. 4 would be a distribution on graphs with 4 vertices. Such graphs would have to contain either two vertices not connected by an edge or two vertices connected by an edge but labeled with the same color or a vertex with no color. However, two such vertices would correspond to a local example which would otherwise have zero probability under the marginals estimated from $\Upsilon$, which are shown above. While the above reasoning is for Model A, a similar argument can be used to show the same issues for Model B.

One of the consequences of the above example is that the unbiased estimates of relational marginals from Proposition 2 cannot always be used for defining distributions of arbitrary domain sizes. The section Estimation shows under which such unbiased estimates do exist, using the concept of relational marginal polytopes.

In order to construct distributions for arbitrary domain sizes, which have relational marginals that are close to the relational marginals given by some global example $\Upsilon$, we will rely on the following construction which we call expansion of global example.

Definition 5 (Expansion of a global example). Let $\Upsilon=$ $(\mathcal{A}, \mathcal{C})$ be a global example where $\mathcal{C}=\left\{c_{1}, c_{2}, \ldots, c_{n}\right\}$, and let $l$ be a positive integer. Then the $l$-level expansion of $\Upsilon$ is a global example $\Upsilon^{\prime}=\left(\mathcal{A}^{\prime}, \mathcal{C}^{\prime}\right)$ given by: $\mathcal{A}^{\prime}=\{a \theta$ : $a \in \mathcal{A}, \theta \in \Theta\}, \mathcal{C}^{\prime}=\left\{c_{1}, c_{2}, \ldots, c_{n}, c_{n+1}, \ldots, c_{l \cdot n}\right\}$. Here, constants $c_{i}, c_{j}$ are said to be congruent if $i \equiv j \bmod n$. Here, $c_{n+1}, \ldots, c_{l \cdot n}$ are some arbitrary new constants and $\Theta$ is a set of all substitutions $\theta$ which satisfy the requirement that $c \theta$ is congruent with $c$ for each $c \in \mathcal{C}$.
Next we illustrate the notion of expansion of examples.
Example 6. Let $\Upsilon=(\mathcal{A}, \mathcal{C})$ be given by

$$
\mathcal{A}=\left\{e\left(c_{1}, c_{2}\right), e\left(c_{2}, c_{3}\right)\right\}, \mathcal{C}=\left\{c_{1}, c_{2}, c_{3}\right\}
$$

Interpreting the predicate $e$ as edge, the corresponding graph looks like:


Then the 2-level expansion of $\Upsilon$ is $\Upsilon^{\prime}=\left(\mathcal{A}^{\prime}, \mathcal{C}^{\prime}\right)$ where

$$
\begin{aligned}
& \mathcal{A}=\left\{e\left(c_{1}, c_{2}\right), e\left(c_{2}, c_{3}\right), e\left(c_{4}, c_{5}\right), e\left(c_{5}, c_{6}\right), e\left(c_{1}, c_{5}\right),\right. \\
& \left.\quad e\left(c_{2}, c_{6}\right), e\left(c_{4}, c_{2}\right), e\left(c_{5}, c_{3}\right)\right\}, \mathcal{C}=\left\{c_{1}, c_{2}, c_{3}, c_{4}, c_{5}, c_{6}\right\} .
\end{aligned}
$$

Interpreting the predicate e again as edge, the expansion corresponds to the following graph:


The width-2 marginal probabilities on $\Upsilon$ and $\Upsilon^{\prime}$ are:

$$
\begin{array}{rrrr}
P_{\Upsilon, 2}\left(\omega_{1}\right)=\frac{1}{3} & P_{\Upsilon^{\prime}, 2}\left(\omega_{1}\right)=\frac{7}{15} & \omega_{1}=(\{ \},\{1,2\}) \\
P_{\Upsilon, 2}\left(\omega_{2}\right)=\frac{1}{3} & P_{\Upsilon^{\prime}, 2}\left(\omega_{2}\right)=\frac{4}{15} & \omega_{2}=(\{e(1,2)\},\{1,2\}) \\
P_{\Upsilon, 2}\left(\omega_{3}\right)=\frac{1}{3} & P_{\Upsilon^{\prime}, 2}\left(\omega_{3}\right)=\frac{4}{15} & \omega_{3}=(\{e(2,1)\},\{1,2\}) \\
P_{\Upsilon, 2}(\omega)=0 & P_{\Upsilon^{\prime}, 2}(\omega)=0 & \omega \notin\left\{\omega_{1}, \omega_{2}, \omega_{3}\right\}
\end{array}
$$

The differences between the marginal probabilities given by $\Upsilon$ and $\Upsilon^{\prime}$ are at most $\frac{2}{15}$ in this case, which is quite high. However, it follows from what we show in turn that this is mostly because of the small size of $\Upsilon$. For larger global examples, the difference between the marginals obtained from them and from their expansions will tend to be smaller.

Importantly, it is possible to bound the difference of the parameters obtained on expansions of global examples and the unbiased estimates obtained on the original examples.
Proposition 7. Let $\Upsilon=(\mathcal{A}, \mathcal{C})$ be a global example and $\Upsilon^{\prime}$ its l-level expansion and let $n=|\mathcal{A}|$ and $k$ be the width of local examples. Then for any formula $\alpha$ :

$$
\left|P_{\Upsilon, k}(\alpha)-P_{\Upsilon^{\prime}, k}(\alpha)\right| \leq 1-\left(\frac{n-k+1}{n}\right)^{k-1}
$$

for Model A, and

$$
\left|Q_{\Upsilon}(\alpha)-Q_{\Upsilon^{\prime}}(\alpha)\right| \leq 1-\left(\frac{n-k_{\alpha}+1}{n}\right)^{k_{\alpha}-1}
$$

for Model B, where $k_{\alpha}=|\operatorname{vars}(\alpha)|$.

Note that the difference between the true and modeled probabilities of a fixed formula $\alpha$ decays as $O\left(\frac{1}{n}\right)$.

The techniques described in this section have the following limitation. If we have a set of hard rules $\Phi_{0}$ which are satisfied by a given $\Upsilon$, these rules may not be satisfied in an expansion of $\Upsilon$. This is not just a limitation of our method though. There are cases where it is not possible to extend a given $\Upsilon$ while satisfying the constraints (this is because we allow the use of equality in the formulas and because we use the unique names assumption). However, if the example $\Upsilon$ is large enough and it satisfies the hard rules, then the number of violations of these rules will be small, which follows again from Proposition 7.

In fact, it seems to be a desirable property that formulas $\alpha$ satisfying $P_{\Upsilon, k}(\alpha)=1$ do not have to be treated as completely hard rules but as rules that "mostly" hold if they are learned from $\Upsilon$, since it may be that they are not really rules that should always hold. Yet, if we actually took them as hard rules we would be forced to assign probability 0 to any example that violates them. It is possible to use the idea of expansions to obtain a distribution in which any formula $\alpha$ has nonzero probability and all properties of expansions are still preserved (such as those from Proposition 7). This can be achieved by randomly sampling additional atoms containing only congruent constants and adding them to the respective expansion. If we use a sufficiently high level of the expansion (at least $k$ for Model A and at least $\max _{\alpha \in \Phi}|\operatorname{vars}(\alpha)|$ for Model B) then the probability of any formula will be nonzero and not equal to one w.r.t. the distribution induced by the expansions with the sampled atoms.

## Relational Marginal Polytopes

In this section, we define another important concept called relational marginal polytope, which will be used in the next section where we deal with estimation errors.
Definition 6 (Relational marginal polytope for Model A). Let $k, m \in \mathbb{N}$ and $\Phi=\left\{\alpha_{1}, \ldots, \alpha_{l}\right\}$ be a set of formulas and $\Phi_{0}$ be a set of hard rules. Let $\mathcal{C}=\{1, \ldots, m\}$ and $\mathcal{C}_{k}$ be the set of size- $k$ subsets of $\mathcal{C}$. Then, for Model $A$, the relational marginal polytope of $\Phi$ of width $k$ and cardinality $m$ w.r.t. the hard rules from $\Phi_{0}$ is the convex hull of the set

$$
\left\{\left.\left(\frac{\#_{k}\left(\alpha_{1}, \Upsilon\right)}{\left|\mathcal{C}_{k}\right|}, \ldots, \frac{\#_{k}\left(\alpha_{l}, \Upsilon\right)}{\left|\mathcal{C}_{k}\right|}\right) \right\rvert\, \Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)\right\}
$$

Let $\Theta_{\alpha_{i}}$ be the set of all injective substitutions from variables of $\alpha_{i}$ to constants from $\mathcal{C}$. Then, for Model B, the relational marginal polytope of $\Phi$ of cardinality $m$ w.r.t. the hard rules from $\Phi_{0}$ is the convex hull of the set

$$
\left\{\left.\left(\frac{n\left(\alpha_{1}, \Upsilon\right)}{\left|\Theta_{\alpha_{1}}\right|}, \ldots, \frac{n\left(\alpha_{l}, \Upsilon\right)}{\left|\Theta_{\alpha_{l}}\right|}\right) \right\rvert\, \Upsilon \in \Pi\left(\mathcal{C}, \Phi_{0}\right)\right\}
$$

Any realizable set of relational marginals for Model A and Model B naturally corresponds to a point in the respective polytope. In the remainder of this paper, we only consider the cases when the relational marginal polytope is full-dimensional, that is, it does not live in a lower dimensional subspace which could happen if some of the relational marginals that define it were linearly dependent.

We will also need the concept of $\eta$-interior of a relational marginal. We say that a point $y$ is in the $\eta$-interior of a relational marginal polytope $P$ if $P$ contains a ball of radius $\eta$ with center in $y$. Using Proposition 4 and Proposition 7, we can show the following both for Model A and Model B.
Proposition 8. Let $\theta$ be a vector representing the values of a set of relational marginals given by formulas from a set $\Phi=\left\{\alpha_{1}, \ldots, \alpha_{l}\right\}$. Let $k$ be the width of the relational marginals of Model A or $k=\max _{\alpha_{i}}\left|\operatorname{vars}\left(\alpha_{i}\right)\right|$ for Model B. Let the set of hard rules $\Phi_{0}$ be empty. If $\theta$ is in the $\left(\eta+\sqrt{l}\left(1-\left(\frac{m-k+1}{m}\right)^{k-1}\right)\right)$-interior of the relational marginal polytope of $\Phi$ of domain-size $m$ then it is also in the $\eta$-interior of the relational marginal polytope of $\Phi$ for any domain size $m^{\prime}$.

## Estimation

In this section, we present error bounds for the estimation of relational marginals. We start by defining the learning setting. Clearly, we need some assumptions on the training and test data and their relationship (otherwise one could always come up with a setting in which the error can be arbitrarily large). In order to stay close to realistic settings we assume that there is some large global example $\aleph=\left(\mathcal{A}_{\aleph}, \mathcal{C}_{\aleph}\right)$ that is not available and that represents the ground truth. That is what we essentially want to estimate for a given formula $\alpha$ is $P_{\aleph, k}(\alpha)$, but we do not have access to whole $\aleph$. Imagine for instance that $\aleph$ is the human gene regulatory network or Facebook. We assume that there is a process that samples size- $m$ subsets of $\mathcal{C}_{\aleph}$ uniformly and that we have access to one such sample $\mathcal{C}_{\Upsilon}$ and also to the respective induced $\Upsilon=\aleph\left\langle\mathcal{C}_{\Upsilon}\right\rangle$. So, for a given formula $\alpha$, we need to estimate $P_{\aleph, k}(\alpha)$ using the available example $\Upsilon$ and the estimate needs to be realizable (otherwise the optimization problem would have no solution and the duality would also not hold). This is a reasonably realistic setting ${ }^{5}$ as in practice we also do not have a distribution over different Facebooks but there is one Facebook and we want to model it based on a sample that is available to us.

We now provide theoretical upper bounds for the expected error of the estimates of $P_{\aleph, k}(\alpha)$ assuming the just described learning setting. However, we first need to describe the estimators. Based on the results from the previous sections, the estimator works as follows. Given a global example $\Upsilon=\left(\mathcal{A}_{\Upsilon}, \mathcal{C}_{\Upsilon}\right)$ and an integer $n$, which is the size of the domain of the modelled distribution (e.g. $n$ can be size of $\aleph ' s ~ d o m a i n ~ i f ~ i t ~ i s ~ k n o w n), ~ w e ~ c o n s t r u c t ~ t h e ~ l-l e v e l ~ e x p a n-~$ sion $\Upsilon^{(l)}$ of $\Upsilon$, where $l=\left\lceil n /\left|\mathcal{C}_{\Upsilon}\right|\right\rceil$, and we use it to estimate the parameters as $\widehat{P}_{\alpha}=P_{\Upsilon^{(l)}, k}(\alpha)$ for Model A and as $\widehat{Q}_{\alpha}=Q_{\Upsilon^{(l)}}(\alpha)$ for Model B. The following proposition introduces an upper bound for the expected error of the estimated parameters.

Proposition 9. Let $m$ and $n$ be positive integers, $\alpha$ a closed formula and let $k$ be the width of local examples. Let

[^5]$\aleph=\left(\mathcal{A}_{\aleph}, \mathcal{C}_{\aleph}\right)$ be a global example, $\mathcal{C}_{\Upsilon}$ be sampled uniformly among all size-m subsets of $\mathcal{C}_{\aleph}$ and $\Upsilon=\aleph\left\langle\mathcal{C}_{\Upsilon}\right\rangle$. Let $\widehat{A}_{\aleph}=P_{\aleph, k}(\alpha)$. Let $\widehat{B}_{\Upsilon}$ be an estimate computed from the l-level expansion of $\mathcal{C}_{\Upsilon}$. Then
$\mathbb{E}\left[\left|\widehat{A}_{\aleph}-\widehat{B}_{\Upsilon}\right|\right] \leq 1-\left(\frac{m-k+1}{m}\right)^{k-1}+\sqrt{\frac{1+2 \log 2}{4\lfloor m / k\rfloor}}$.
In the case of model B, the same upper bound holds if we choose $k=|\operatorname{vars}(\alpha)|$ and $\widehat{A}_{\aleph}=Q_{\aleph}(\alpha)$.
The proof of this proposition, which is contained in the online appendix, is based on a deviation bound for a randomized estimator of $\widehat{B}_{\Upsilon}$.

It is possible to improve the estimation in some cases. If the vector corresponding to the marginals $\Phi$ estimated from $\Upsilon$ is guaranteed to be in the $\left(\eta+\sqrt{l}\left(1-\left(\frac{m-k+1}{m}\right)^{k-1}\right)\right)$ interior of the relational marginal polytope of domain of size $m=\left|\mathcal{C}_{\Upsilon}\right|$ for some $\eta>0$, where $k$ is the width of local examples for Model A (or $\max _{\alpha \in \Phi}|\operatorname{vars}(\alpha)|$ for Model B), then, by Proposition 8, we can estimate the parameters directly from $\Upsilon$ without constructing its expansion. We then have the following improved bound:

$$
\mathbb{E}\left[\left|\widehat{A}_{\aleph}-\widehat{B}_{\Upsilon}\right|\right] \leq \sqrt{(1+2 \log 2) /(4\lfloor m / k\rfloor)}
$$

It is interesting to note that the lower-bound on effective sample size obtained from these bounds is $\lfloor m / k\rfloor$, which is also the maximum number of non-overlapping size- $k$ subsets of $\mathcal{C}_{\Upsilon}$. A consequence for learning the parameters of models such as MLNs (which corresponds to relational marginal problems in Model B ) is that this bound is inversely proportional to the number of variables in the used formulas, which also suggests an explanation for why learning with longer formulas is difficult.

## Related Work

The relational marginals from Model A were recently introduced (Kuželka, Davis, and Schockaert 2017). However, they were only studied in a possibilistic setting, which differs substantially from the probabilistic maximum-entropy setting that we considered. The idea of using random substitutions (Model B) goes back to (Bacchus et al. 1992) who, however, only considered unary predicates. Schulte et al. (2014) used the random substitutions semantics to define a relational Bayesian network model for population statistics. However, their model is, not based on any underlying ground model, and it is unclear whether the distributions are always realizable by a ground model.

In the more restricted setting of exponential random graph models (ERGMs, Chatterjee and Diakonis 2013), a formally similar duality, based on densities of graph homomorphisms, has previously been established. To the best of our knowledge, however, such a duality has never been established in an SRL setting. In fact, even for ERGMs this duality has not yet been exploited for estimating parameters for models of different domain sizes, which is one of the key contributions of our work.

Certain statistical properties of learning have been already studied for SRL models. Xiang and Neville (2011) studied consistency but postulated rather strong assumptions ${ }^{6}$, as a result of which their results are not comparable with ours. Their approach also differs in that it only considers distributions of labels conditioned on the underlying graph structure. It is interesting that a statistical estimation problem equivalent to the estimation of parameters in Model A has also been studied in the literature. In Nandi and Sen (1963), the variance of an estimator, equivalent to the unbiased estimator for Model A, was given. However, we are not aware of any work showing a deviation bound in the same setting, which was needed to establish the bound on expected error in Proposition 9. Interestingly, the effective sample size $m / k$ stemming from the work (Nandi and Sen 1963) for variance of the estimator is almost the same as the effective sample size $\lfloor m / k\rfloor$ stemming from our deviation and expected-error bounds. This actually suggests that our bounds are rather tight. There were many works on U-statistics (Hoeffding 1948) which are related as well but they rely on assumptions that are generally not realistic for SRL and, in particular, are not applicable to our setting; the work of Nandi and Sen (1963) is an exception.

## Conclusions

In this paper, we have introduced and studied relational marginal problems. Interestingly, this perspective enables learning a model that is applicable to data sets whose domain sizes differ from that of the training data. We established a relational counterpart of the classical duality between maximum-likelihood and max-entropy marginal problems. Then, we showed how to estimate and adjust parameters of the marginals in order to guarantee their realizability. We complemented these results by providing bounds on the expected errors of the estimates in a reasonable setting.

We believe that due to the simplicity and transparency of the learning setting that we introduced, this setting could play a similar role for SRL as the standard i.i.d. statistical learning setting plays for learning in propositional domains (Vapnik 1995). That is, as an idealized setting that is suitable for theoretical study, but that is not too far from settings that one encounters in reality. Still, it would be possible to extend the learning setting to make it more realistic. In particular, the sampling process that creates the training examples could be replaced by another sampling process that would take into account the structure of the relational data. That would probably make estimation of parameters and derivation of error bounds significantly more complex, and hence arguably less illuminating, which is why we leave it for future work.

## Acknowledgment

The authors would like to thank the anonymous reviewers for their helpful comments. This work was supported by a Leverhulme Trust grant (RPG-2014-164) and ERC Starting Grant 637277. JD is partially supported by the KU Leuven

[^6]Research Fund (C14/17/070,C22/15/015,C32/17/036), and FWO-Vlaanderen (G.0356.12, SBO-150033). YW is partially supported by Guangdong Shannon Intelligent Tech. co., Ltd.

## References

Bacchus, F.; Grove, A. J.; Koller, D.; and Halpern, J. Y. 1992. From statistics to beliefs. In Proceedings of the 10th National Conference on Artificial Intelligence, 602-608.
Boyd, S., and Vandenberghe, L. 2004. Convex optimization. Cambridge university press.
Chatterjee, S., and Diaconis, P. 2013. Estimating and understanding exponential random graph models. The Annals of Statistics 41(5):2428-2461.
Getoor, L., and Taskar, B. 2007. Introduction to statistical relational learning. MIT press.
Hoeffding, W. 1948. A class of statistics with asymptotically normal distribution. The annals of mathematical statistics 293-325.
Jain, D.; Kirchlechner, B.; and Beetz, M. 2007. Extending Markov logic to model probability distributions in relational domains. In KI 2007: Advances in Artificial Intelligence, 30th Annual German Conference on AI, KI 2007, 129-143.
Kuželka, O.; Davis, J.; and Schockaert, S. 2017. Induction of interpretable possibilistic logic theories from relational data. In Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, 11531159.

Nandi, H., and Sen, P. 1963. On the properties of u-statistics when the observations are not independent: Part two unbiased estimation of the parameters of a finite population. Calcutta Statistical Association Bulletin 12(4):124-148.
Richardson, M., and Domingos, P. 2006. Markov logic networks. Machine Learning 62(1-2):107-136.
Schulte, O.; Khosravi, H.; Kirkpatrick, A. E.; Gao, T.; and Zhu, Y. 2014. Modelling relational statistics with Bayes nets. Machine Learning 94(1):105-125.
Shalizi, C. R., and Rinaldo, A. 2013. Consistency under sampling of exponential random graph models. The Annals of Statistics 41(2):508-535.
Singh, M., and Vishnoi, N. K. 2014. Entropy, optimization and counting. In Proceedings of the forty-sixth annual ACM symposium on Theory of computing, 50-59. ACM.
Van Haaren, J.; Van den Broeck, G.; Meert, W.; and Davis, J. 2016. Lifted generative learning of Markov logic networks. Machine Learning 103(1):27-55.
Vapnik, V. N. 1995. The Nature of Statistical Learning Theory. New York, NY, USA: Springer-Verlag New York, Inc. Wainwright, M. J., and Jordan, M. I. 2008. Graphical models, exponential families, and variational inference. Foundations and Trends in Machine Learning 1(1-2):1-305.
Xiang, R., and Neville, J. 2011. Relational learning with one network: An asymptotic analysis. In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2011, 779-788.


[^0]:    Copyright © $C^{2018}$, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ All statistics in this paper are based on universally quantified formulas; we do not consider formulas with existential quantifiers.

[^2]:    ${ }^{2}$ https://arxiv.org/abs/1709.05825

[^3]:    ${ }^{3}$ What we call a global example in this paper is usually called a possible world in the MLN literature.

[^4]:    ${ }^{4}$ Proposition 4 is one of the results which would not hold for Model B if we did not require injectivity of the randomly sampled substitutions in the definition of Model B.

[^5]:    ${ }^{5}$ What might differ in realistic settings is the sampling process. We briefly discuss this in section Conclusions.

[^6]:    ${ }^{6}$ The assumptions used in their work were weak dependency and bounded degree of graph nodes.

