
SEMI-SUPERVISED LEARNING USING DIFFERENTIABLE REASONING

EMILE VAN KRIEKEN
University of Amsterdam
info@emilevankrieken.com

ERMAN ACAR
VU University Amsterdam
erman.acar@vu.nl

FRANK VAN HARMELEN
VU University Amsterdam
Frank.van.Harmelen@vu.nl

Abstract

We introduce Differentiable Reasoning (DR), a novel semi-supervised learning technique which uses relational background knowledge to benefit from unlabeled data. We apply it to the Semantic Image Interpretation (SII) task and show that background knowledge provides significant improvement. We find that there is a strong but interesting imbalance between the contributions of updates from Modus Ponens (MP) and its logical equivalent Modus Tollens (MT) to the learning process, suggesting that our approach is very sensitive to a phenomenon called the Raven Paradox [10]. We propose a solution to overcome this situation.

1 Introduction

Semi-supervised learning is a common class of methods for machine learning tasks where we consider not just labeled data, but also make use of unlabeled data [2]. This can be very beneficial for training in tasks where labeled data is much harder to acquire than unlabeled data.

We thank the anonymous reviewers of NeSy for their valuable feedback. Additionally we thank Luciano Serafini, Jasper Driessens, Fije Overeem and also the other attendees of NeSy for the great discussions.

One such task is *Semantic Image Interpretation (SII)* in which the goal is to generate a semantic description of the objects on an image [7]. This description is represented as a labeled directed graph, which is known as a *scene graph* [13]. An example of a labeled dataset for this problem is VisualGenome [15] which contains 108,077 images to train 156,722 different unary and binary predicates. The binary relations in particular make this dataset very sparse, as there are many different pairs of objects that could be related. However, a far larger, though unfortunately unlabeled, dataset like ImageNet [24] contains over 14 million different pictures. Because it is so much larger, it will have many examples of interactions that are not present in VisualGenome. We show that it is possible to improve the performance of a simple classifier on the SII task significantly by adding the satisfaction of a first-order logic (FOL) knowledge base to the supervised loss function. The computation of this satisfaction uses an unlabeled dataset as its domain.

For this purpose, we introduce a statistical relational learning framework called *Differentiable Reasoning (DR)* in Section 2, as our primary contribution. DR uses simple logical formulas to deduce new training examples in an unlabeled dataset. This is done by adding a differentiable loss term that evaluates the truth value of the formulas.

In the experimental analysis, we find that the gradient updates using the Modus Ponens (MP) and Modus Tollens (MT) rules are disproportionate. That is, MT often strongly dominates MP in the learning process. Such behavior suggests that our approach is highly sensitive to the Raven Paradox [10]. It refers to the phenomenon that the observations obtained from “All ravens are black” are dominated by its logically equivalent “All non-black things are non-ravens”. Indeed, this is closely related to the *material implication* which caused a lot of discussion throughout the history of logic and philosophy [8]. Our second main contribution relies on its investigation in Section 2.4, and our proposal to cope with it. Finally, we show results on a simple dataset in Section 3 and analyze the behavior of the Raven Paradox in Section 4. Related works and conclusion closes the paper.

2 Differentiable Reasoning

2.1 Basics and Notation

We assume a knowledge base \mathcal{K} is given in a relational logic language, where a formula $\varphi \in \mathcal{K}$ is built from predicate symbols $P \in \mathcal{P}$, a finite set \mathcal{D} of objects (also called constants) $o \in \mathbb{R}^m$ with $m \in \mathbb{Z}^+$, and variables $x \in \mathcal{V}$, in the usual way (see [28]). We also assume that every $\varphi \in \mathcal{K}$ is in Skolem normal form. For a vector of objects and variables, we use boldfaced \mathbf{o} and \mathbf{x} , respectively. A ground atom is a formula with no logical connective and no variables, e.g., $\text{partOf}(\text{cushion}, \text{chair})$ where $\text{partOf} \in \mathcal{P}$ and $\text{cushion}, \text{chair} \in \mathcal{D}$. Given a subset $D_i \subseteq \mathcal{D}$, a *Herbrand base* \mathbf{A}_i corresponding to D_i is the set of all ground atoms generated from D_i and \mathcal{P} . A *world* (often called a *Herbrand interpretation*) \mathbf{w}_i for

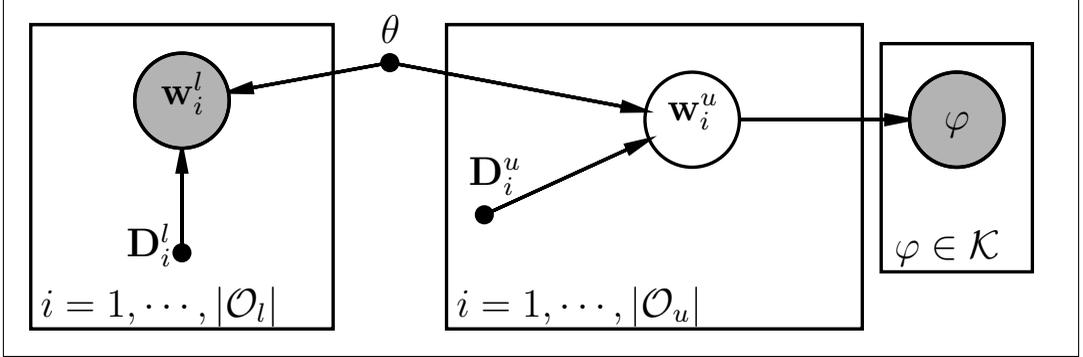


Figure 1: The Bayesian network describing the joint probability $p(\mathcal{W}_l, \mathcal{W}_u, \mathcal{K} | \mathcal{O}_l, \mathcal{O}_u, \theta)$. The left plate is the supervised classification likelihood and the right plates the unsupervised part in which we calculate the probability of the formulas \mathcal{K} . The parameters θ are shared in both parts.

D_i assigns a binary truth value to each ground atom $P(\mathbf{o}) \in \mathbf{A}_i$ i.e., $\mathbf{w}_i(P(\mathbf{o})) \in \{0, 1\}$.

Each predicate P has a corresponding differentiable function $f_P^\theta(\mathbf{o}) \in \mathbb{R}^{\alpha(P) \times m} \rightarrow [0, 1]$ parameterized by θ (a vector of reals) with $\alpha(P)$ being the arity of P , which calculates the probability of $P(\mathbf{o})$. This function could be, for instance, a neural network.

Next, we define a Bernoulli distribution function over worlds as follows

$$p(\mathbf{w}_i | \theta, D_i) = \prod_{P(\mathbf{o}) \in \mathbf{A}_i} f_P^\theta(\mathbf{o})^{\mathbf{w}_i(P(\mathbf{o}))} \cdot (1 - f_P^\theta(\mathbf{o}))^{1 - \mathbf{w}_i(P(\mathbf{o}))} \quad (1)$$

where $\mathbf{w}_i(P(\mathbf{o}))$ (similarly, $1 - \mathbf{w}_i(P(\mathbf{o}))$) refers to the exponent. Given some world \mathbf{w}_i , the *valuation function* $v(\varphi, \mathbf{w}_i)$ is 1 if φ is true in that world, that is, $\mathbf{w}_i \models \varphi$, and 0 otherwise.

Next, we explain the domain we use in this article. We have a dataset \mathcal{D} partitioned into two parts: a labeled dataset $\mathcal{D}_l = \langle \mathcal{O}_l, \mathcal{W}_l \rangle$ and an unlabeled dataset $\mathcal{D}_u = \langle \mathcal{O}_u, \emptyset \rangle$ where both \mathcal{O}_l and \mathcal{O}_u are sets of finite domains D_i , and \mathcal{W}_l is a set containing the correct world \mathbf{w}_i^{l*} for all pictures i .

In Figure 1 we illustrate the Bayesian network associated with this problem. The left plate denotes the usual supervised data likelihood $p(\mathcal{W}_l | \mathcal{O}_l, \theta)$ and the right plates denote the probabilities of the truth values of the formulas $\varphi \in \mathcal{K}$ using $p(\mathcal{K} | \mathcal{O}_u, \theta)$.

It is important to note that the true worlds \mathbf{w}_i^{u*} of the unlabeled dataset are not known, that is, they are latent variables and they have to be marginalized over. The formulas in knowledge base \mathcal{K} are all assumed to be true. We can now obtain the optimization problem

that we can solve using gradient descent as

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} p(\mathcal{W}_l | \mathcal{O}_l, \boldsymbol{\theta}) \cdot p(\mathcal{K} | \mathcal{O}_u, \boldsymbol{\theta}) \quad (2)$$

$$= \arg \max_{\boldsymbol{\theta}} \prod_{i=1}^{|\mathcal{O}_l|} p(\mathbf{w}_i^{l*} | D_i^l, \boldsymbol{\theta}) \cdot \prod_{i=1}^{|\mathcal{O}_u|} \sum_{\mathbf{w}_i^u} p(\mathbf{w}_i^u | D_i^u, \boldsymbol{\theta}) \cdot \prod_{\varphi \in \mathcal{K}} v(\varphi, \mathbf{w}_i^u) \quad (3)$$

$$= \arg \min_{\boldsymbol{\theta}} - \sum_{i=1}^{|\mathcal{O}_l|} \log p(\mathbf{w}_i^{l*} | D_i^l, \boldsymbol{\theta}) - \sum_{i=1}^{|\mathcal{O}_u|} \log \left(\sum_{\mathbf{w}_i^u} p(\mathbf{w}_i^u | D_i^u, \boldsymbol{\theta}) \cdot \prod_{\varphi \in \mathcal{K}} v(\varphi, \mathbf{w}_i^u) \right) \quad (4)$$

where in the last step we take the log and minimize with respect to the negative value. The optimization problem in Equation 4 consists of two terms. The first is the cross-entropy loss for supervised labeled data. The second can be understood as follows: A world entails a (full) knowledge base (i.e., $\mathbf{w} \models \mathcal{K}$) if $\mathbf{w} \models \varphi$ holds for all $\varphi \in \mathcal{K}$ (that is, the product of their valuations is 1). For each domain D_i , we then find the sum of the probabilities of worlds that entail the knowledge base. This is an example of what we call the *differentiable reasoning* loss. The general differentiable reasoning objective is given as

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} - \sum_{i=1}^{|\mathcal{O}_l|} \log p(\mathbf{w}_i^{l*} | D_i^l, \boldsymbol{\theta}) + \mathcal{L}_{DR}(\boldsymbol{\theta}; \mathcal{K}, \mathcal{O}_u). \quad (5)$$

2.2 Differentiable Reasoning Using Product Real Logic

The marginalization over all possible worlds \mathbf{w}_i^u requires $2^{|\mathbf{A}_i|}$ combinations, so it is exponential in the size of the Herbrand base. Therefore, the problem of finding the sum of the probabilities $p(\mathbf{w}_i | \boldsymbol{\theta})$ for all worlds \mathbf{w}_i that entail the knowledge base \mathcal{K} is #P-complete [23]. Instead, we shall perform a much simpler computation defined over logical formulas and the parameters $\boldsymbol{\theta}$ as follows:

$$\mathcal{L}_{DR}(\boldsymbol{\theta}; \mathcal{K}, \mathcal{O}_u) = \sum_{\varphi \in \mathcal{K}} \mathcal{L}(\boldsymbol{\theta}; \varphi, \mathcal{O}_u) \quad (6)$$

$$\mathcal{L}(\boldsymbol{\theta}; \forall \mathbf{x} \phi, \mathcal{O}_u) = - \sum_{D \in \mathcal{O}_u, \mathbf{o} \in D} \log \hat{p}(\phi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) \quad (7)$$

$$\hat{p}(\mathbf{P}(x_1, \dots, x_{\alpha(\mathbf{P})}) | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) = f_{\mathbf{P}}^{\boldsymbol{\theta}}(o_1, \dots, o_{\alpha(\mathbf{P})}) \quad (8)$$

$$\hat{p}(\neg \phi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) = 1 - \hat{p}(\phi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) \quad (9)$$

$$\hat{p}(\phi \wedge \psi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) = \hat{p}(\phi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) \cdot \hat{p}(\psi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) \quad (10)$$

$$\hat{p}(\phi \vee \psi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) = \hat{p}(\neg(\neg \phi \wedge \neg \psi) | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) \quad (11)$$

$$\hat{p}(\phi \rightarrow \psi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) = \hat{p}(\neg \phi \vee \psi | \mathbf{x} = \mathbf{o}, \boldsymbol{\theta}) \quad (12)$$

where $\alpha : \mathcal{P} \rightarrow \mathbb{Z}^+$ is the arity function for each predicate symbol, and ϕ and ψ are subformulas of φ . \hat{p} computes the fuzzy degree of truth of some formula φ using the product norm and the Reichenbach implication [1], which makes our approach a special case of Real Logic [26] that we call *Product Real Logic*. The \forall quantifier is interpreted in Equation 7 by going through all instantiations, which in this case is all n -tuples in the domain D_i , and also looping over all domains D_i in the set of domains (i.e., pictures) \mathcal{O}_i .

Example 1. The loss term associated with the formula $\varphi = \forall x, y \text{ chair}(x) \wedge \text{partOf}(y, x) \rightarrow \text{cushion}(y) \vee \text{armRest}(y)$ is computed as follows:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}; \varphi, \mathcal{O}_u) = - \sum_{D \in \mathcal{O}, o_1, o_2 \in D} & 1 - f_{\text{chair}}^{\boldsymbol{\theta}}(o_1) \cdot f_{\text{partOf}}^{\boldsymbol{\theta}}(o_2, o_1) \cdot \\ & (1 - f_{\text{cushion}}^{\boldsymbol{\theta}}(o_2)) \cdot (1 - f_{\text{armRest}}^{\boldsymbol{\theta}}(o_2)) \end{aligned}$$

Say \mathcal{O}_u contains the picture in Figure 2 whose domain is $\{a, b\}$ and the model predicts the following distribution over worlds:

$$\begin{array}{ll} f_{\text{chair}}^{\boldsymbol{\theta}}(a) = 0.9 & f_{\text{chair}}^{\boldsymbol{\theta}}(b) = 0.4 \\ f_{\text{cushion}}^{\boldsymbol{\theta}}(a) = 0.05 & f_{\text{cushion}}^{\boldsymbol{\theta}}(b) = 0.5 \\ f_{\text{armRest}}^{\boldsymbol{\theta}}(a) = 0.05 & f_{\text{armRest}}^{\boldsymbol{\theta}}(b) = 0.1 \\ f_{\text{partOf}}^{\boldsymbol{\theta}}(a, a) = 0.001 & f_{\text{partOf}}^{\boldsymbol{\theta}}(b, b) = 0.001 \\ f_{\text{partOf}}^{\boldsymbol{\theta}}(a, b) = 0.01 & f_{\text{partOf}}^{\boldsymbol{\theta}}(b, a) = 0.95 \end{array}$$

The model returns high values for $f_{\text{chair}}^{\boldsymbol{\theta}}(a)$ and $f_{\text{partOf}}^{\boldsymbol{\theta}}(b, a)$ but it is not confident of $f_{\text{cushion}}^{\boldsymbol{\theta}}(b)$, even though it is clearly higher than $f_{\text{armRest}}^{\boldsymbol{\theta}}(b)$. We can decrease $\mathcal{L}(\boldsymbol{\theta}; \varphi, \mathcal{O}_u) =$

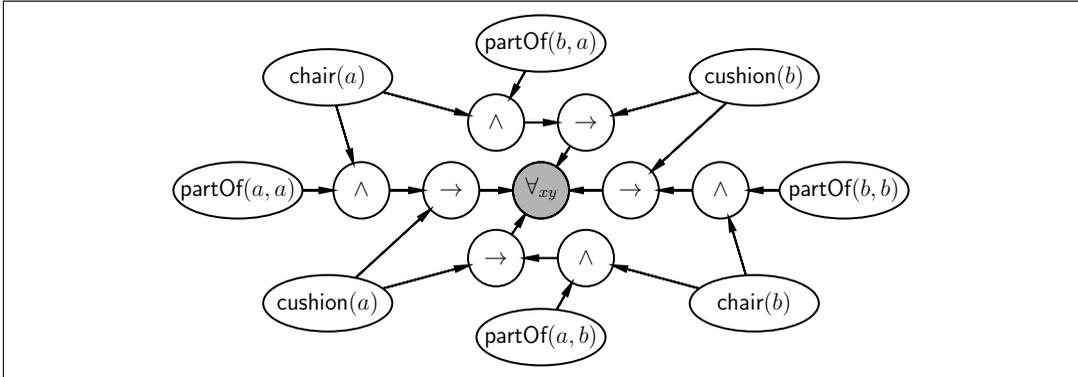


Figure 3: The Bayesian network associated with grounding of the formula $\forall x, y \text{ chair}(x) \wedge \text{partOf}(y, x) \rightarrow \text{cushion}(y)$ on the domain from Figure 2. We treat connectives and quantifiers as binary random variables (which correspond to subformulas through their parents) of which the conditional probabilities are computed using truth tables.

0.612 simply by increasing $f_{\text{cushion}}^\theta(b)$, since $f_{\text{cushion}}^\theta$ is a differentiable function with respect to θ .

This example shows that we can find a new instance of the `cushion` predicate using reasoning on an unlabeled dataset. This process uses both statistical reasoning and symbolic rules. As more data improves generalization, those additional examples could help reducing the sparsity of the SII problem. Furthermore, [7] showed that it is also possible to correct wrong labels due to noisy data when these do not satisfy the formulas.

Figure 3 shows the Bayesian Network for this formula on the picture from Figure 2, illustrating the computation path. We treat each subformula as a binary random variable of which the conditional probabilities are given by truth tables. Because the graph is not acyclic, we can use loopy belief propagation which is empirically shown to often be a good approximation of the correct probability [18]. In fact, Product Real Logic can be seen as performing a single iteration of belief propagation. However, this can be problematic. For example, the degree of truth of the ground formula $\text{chair}(o) \wedge \text{chair}(o)$ would be computed using $f_{\text{chair}}^\theta(o)^2$ instead of the probability of this statement, $f_{\text{chair}}^\theta(o)$ [22]. We show in Appendix A that Product Real Logic computes the correct probability $p(\mathcal{K}|\mathcal{O}_u, \theta)$ for a corpus \mathcal{K} under the strong assumption that, after grounding, each ground atom is used at most once.

An interesting and useful property of our approach is that it can perform multi-hop reasoning in an iterative, yet extremely noisy, manner. In one iteration it might, for instance, increase $f_{\text{cushion}}^\theta(o)$. And since $f_{\text{cushion}}^\theta(o)$ will return higher values in future iterations, it can be used to prove that the probability of other ground atoms that occur in formulas with

cushion(o) should also be increased or decreased.

A convenient property of the SII task is that we consider just binary relations between objects appearing on the same pictures. The Herbrand base then contains $O(|\mathcal{P}| \cdot |D_i|^2)$ ground atoms, which is feasible as there are often not more than a few dozen objects on an image. This property also holds in natural language to some degree in the following way: only the words appearing in the same paragraph can be related. This is in contrast to the knowledge base completion task where we have a single graph with many objects and predicates [27].

2.3 Implementation

We optimize the negative logarithm of the likelihood function given in Equation 4. In particular, we use minibatch gradient descent to decrease the computation time both for the supervised part of the loss and the unsupervised part. We turn the unsupervised loss into minibatch gradient descent by approximating the computation of the \forall quantifier: instead of summing over all n -tuples and all domains, we randomly sample from these n -tuples independently from the domain it belongs to.

2.4 The Material Implication

To provide a better understanding of the inner machinery of our approach, we will elaborate on some interesting partial derivatives. Say, we have a formula φ of the form $\forall \mathbf{x} \phi(\mathbf{x}) \rightarrow \psi(\mathbf{x})$, where $\phi(\mathbf{x})$ is the antecedent and $\psi(\mathbf{x})$ the consequent of φ . First, we write out the partial derivative of $\mathcal{L}(\boldsymbol{\theta}; \varphi, \mathcal{O}_u)$ with respect to the consequent, where we make use of the chain rule:

$$d_{\varphi}^{\text{MP}}(\mathbf{o}) := \frac{\partial \log \hat{p}(\varphi | \mathcal{O}_u, \boldsymbol{\theta})}{\partial \hat{p}(\psi | \mathbf{o}, \boldsymbol{\theta})} = \frac{\partial \sum_{\mathbf{o} \in D, D \in \mathcal{O}_u} \log \hat{p}(\phi \rightarrow \psi | \mathbf{o}, \boldsymbol{\theta})}{\partial \hat{p}(\psi | \mathbf{o}, \boldsymbol{\theta})} \quad (13)$$

$$= \frac{\partial \sum_{\mathbf{o} \in D, D \in \mathcal{O}_u} \log(1 - \hat{p}(\phi | \mathbf{o}, \boldsymbol{\theta}) \cdot (1 - \hat{p}(\psi | \mathbf{o}, \boldsymbol{\theta})))}{\partial \hat{p}(\psi | \mathbf{o}, \boldsymbol{\theta})} \quad (14)$$

$$= \frac{\hat{p}(\phi | \mathbf{o}, \boldsymbol{\theta})}{1 - \hat{p}(\phi | \mathbf{o}, \boldsymbol{\theta}) \cdot (1 - \hat{p}(\psi | \mathbf{o}, \boldsymbol{\theta}))} = \frac{\hat{p}(\phi | \mathbf{o}, \boldsymbol{\theta})}{\hat{p}(\phi \rightarrow \psi | \mathbf{o}, \boldsymbol{\theta})} \quad (15)$$

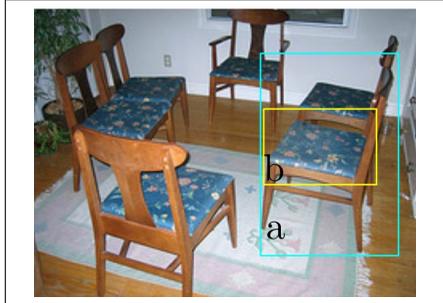


Figure 2: We can deduce that b is a cushion if we are confident about the truth value of $\text{chair}(a)$ and $\text{partOf}(b, a)$ using the formula $\forall x, y \text{ chair}(x) \wedge \text{partOf}(y, x) \rightarrow \text{cushion}(y)$.

$d_{\varphi}^{\text{MP}}(\mathbf{o})$ mirrors the application of the Modus Ponens (MP) rule using the implication $\phi \rightarrow \psi$ for the assignment of \mathbf{o} to \mathbf{x} . The MP rule says that if ϕ is true and $\phi \rightarrow \psi$, then ψ should also be true. Similarly, if $\phi(\mathbf{o})$ is likely and $\phi \rightarrow \psi$, then $\psi(\mathbf{o})$ should also be likely. Indeed, notice that $d_{\varphi}^{\text{MP}}(\mathbf{o})$ grows with $\hat{p}(\phi|\mathbf{o}, \boldsymbol{\theta})$. Also, $d_{\varphi}^{\text{MP}}(\mathbf{o})$ is largest when $\hat{p}(\phi|\mathbf{o}, \boldsymbol{\theta})$ is high and $\hat{p}(\psi|\mathbf{o}, \boldsymbol{\theta})$ is low as it then approaches a singularity in the divisor. We next show the derivation with respect to the negated antecedent:

$$d_{\varphi}^{\text{MT}}(\mathbf{o}) := \frac{\partial \log \hat{p}(\varphi|\mathcal{O}_u, \boldsymbol{\theta})}{\partial \hat{p}(\neg\phi|\mathbf{o}, \boldsymbol{\theta})} = \frac{\hat{p}(\neg\psi|\mathbf{o}, \boldsymbol{\theta})}{\hat{p}(\phi \rightarrow \psi|\mathbf{o}, \boldsymbol{\theta})} \quad (16)$$

Similarly, it mirrors the application of the Modus Tollens (MT) rule which says that if ψ is false and $\phi \rightarrow \psi$, then ϕ should also be false. Again, realize that $d_{\varphi}^{\text{MT}}(\mathbf{o})$ grows with $\hat{p}(\psi|\mathbf{o}, \boldsymbol{\theta})$.

It is easy to see that $d_{\varphi}^{\text{MP}}(\mathbf{o}) > d_{\varphi}^{\text{MT}}(\mathbf{o})$ whenever $\hat{p}(\phi|\mathbf{o}, \boldsymbol{\theta}) > \hat{p}(\neg\psi|\mathbf{o}, \boldsymbol{\theta})$. Furthermore, the global minimum of $\mathcal{L}(\boldsymbol{\theta}; \varphi, \mathcal{O}_u)$ is some parameter value $\boldsymbol{\theta}^*$ so that $\hat{p}(\phi|\mathcal{O}_u, \boldsymbol{\theta}^*) = 0$ and $\hat{p}(\psi|\mathcal{O}_u, \boldsymbol{\theta}^*) = 1$ for all \mathbf{o} , which corresponds to the material implication.

Next, we show how these quantities are used in the updating of the parameters $\boldsymbol{\theta}$ using backpropagation and act as mixing components on the gradient updates:

$$\frac{\log \hat{p}(\varphi|\mathcal{O}_u, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \sum_{\mathbf{o} \in D, D \in \mathcal{O}_u} d_{\varphi}^{\text{MP}}(\mathbf{o}) \cdot \frac{\partial \hat{p}(\psi|\mathbf{o}, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} + d_{\varphi}^{\text{MT}}(\mathbf{o}) \cdot \frac{\partial \hat{p}(\neg\phi|\mathbf{o}, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \quad (17)$$

2.5 The Raven Paradox

In our experiments, we have found that this approach is very sensitive to the *raven paradox* [10]. It is stated as follows: Assuming that observing an example of a statement is evidence for that statement (i.e., the degree of belief in that statement increases), and that evidence for a sentence also is evidence for all the other logically equivalent sentences, then our belief in “ravens are black” increases when we observe non-black non-raven, by the contrapositive “non-ravens are non-black”. Equation 17 shows however that the gradient is equally determined by positive evidence (observing black ravens) as by contrapositive evidence (observing non-black non-ravens). Because in the real world there are far more ravens than non-black objects, optimizing $\hat{p}(\forall o \text{ raven}(o) \rightarrow \text{black}(o)|\mathcal{O}_u, \boldsymbol{\theta})$ amounts to recognizing that something is not a raven when it is not black. However, Machine Learning models tend to be biased when the class distribution is unbalanced during training [30].

Figure 4 shows plots of $d_{\varphi}^{\text{MP}}(\mathbf{o})$ and $d_{\varphi}^{\text{MT}}(\mathbf{o})$ for different values of $\hat{p}(\phi|\mathbf{o}, \boldsymbol{\theta})$ and $\hat{p}(\psi|\mathbf{o}, \boldsymbol{\theta})$. In practice, for many formulas of this form, the most common case will be that the model predicts $\neg\phi(\mathbf{o}) \wedge \neg\psi(\mathbf{o})$. Then, $d_{\varphi}^{\text{MP}}(\mathbf{o})$ approaches 0 and $d_{\varphi}^{\text{MT}}(\mathbf{o})$ will be around 1. For instance, the average value of $d_{\varphi}^{\text{MP}}(\mathbf{o})$ for the problem in Example 1 is 0.214, while the average value of $d_{\varphi}^{\text{MT}}(\mathbf{o})$ is 0.458.

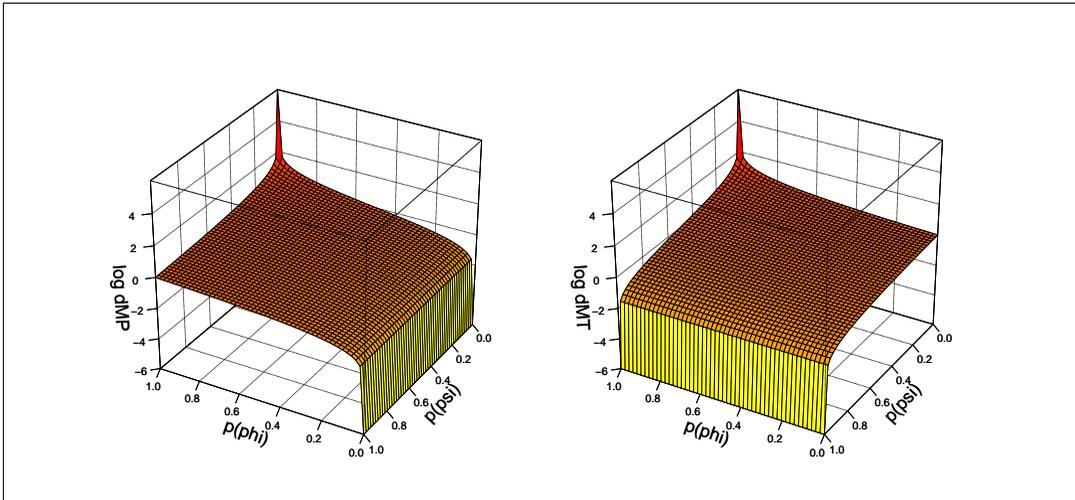


Figure 4: Plots of $d_{\varphi}^{\text{MP}}(\mathbf{o})$ (Equation 15) and $d_{\varphi}^{\text{MT}}(\mathbf{o})$ (Equation 16). Note that the y axis is using a log scale.

We analyze a naive way of dealing with this phenomenon. We normalize the contribution to the total gradient of MP and MT reasoning by replacing the loss function \mathcal{L} of rules of the form $\forall \mathbf{x} \phi(\mathbf{x}) \rightarrow \psi(\mathbf{x})$ as follows:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}; \mathcal{K}, \mathcal{O}_u) = & - \sum_{\varphi \in \mathcal{K}} \sum_{\mathbf{o} \in D, D \in \mathcal{O}_u} \frac{\mu \cdot d_{\varphi}^{\text{MP}}(\mathbf{o})}{\sum_{\mathbf{o}' \in D, D \in \mathcal{O}_u} d_{\varphi}^{\text{MP}}(\mathbf{o}')} \cdot \hat{p}(\psi | \mathbf{o}, \boldsymbol{\theta}) \\ & + \frac{(1 - \mu) \cdot d_{\varphi}^{\text{MT}}(\mathbf{o})}{\sum_{\mathbf{o}' \in D, D \in \mathcal{O}_u} d_{\varphi}^{\text{MT}}(\mathbf{o}')} \cdot \hat{p}(\neg \phi | \mathbf{o}, \boldsymbol{\theta}) \end{aligned} \quad (18)$$

where μ is a hyperparameter that assigns the relative importance of Modus Ponens with respect to Modus Tollens updates. We are then able to control how much either contributes to the training process. We experiment with different values of μ and report our findings in the next section.

3 Experiments

We carried out simple experiments on the PASCAL-PART dataset [3] in which the task is to predict the type of the object in a bounding box and the `partOf` relation which expresses that some bounding box is a part of another. For example, a tail can be a part of a cat. Like in [7], the output softmax layer over the 64 object classes of a Fast R-CNN [9] detector is used for the bounding box features. Note that this makes the problem of recognizing types

very easy as the features correlate strongly with the true output types. Therefore, to get a more realistic estimate, we randomly split the dataset into only 7 labeled pictures for \mathcal{D}_l and 2128 unlabeled pictures for \mathcal{D}_u . Additionally, we only consider 11 (related) types out of 64 due to computational constraints. As there is a large amount of variance associated with randomly splitting in this way, we do all our experiments on 20 random splits of the dataset. The results are evaluated on a held-out validation set of 200 images. We compare the accuracy of prediction of the type of the bounding box and the AUC (area under curve) for the `partOf` relationship.

We model $f_{\text{type}_i}^\theta(o)$ using a single Logic Tensor Network (LTN) layer [7] of width 10 followed by a softmax output layer to ensure mutual exclusivity of types. The term $f_{\text{partOf}}^\theta(o_1, o_2)$ is modeled using an LTN layer of width 2 and a sigmoid output layer. The loss function is then optimized using RMSProp over 6000 iterations. We use the same relational background knowledge as [7] which are rules like the following:

$$\begin{array}{l} \forall x, y \text{ chair}(x) \wedge \text{partOf}(y, x) \rightarrow \text{cushion}(y) \vee \text{armRest}(y) \\ \forall x, y \text{ cushion}(x) \wedge \text{partOf}(x, y) \rightarrow \text{chair}(y) \vee \text{bench}(y) \\ \forall x \neg \text{partOf}(x, x) \\ \forall x, y \text{ partOf}(x, y) \rightarrow \neg \text{partOf}(y, x) \end{array}$$

We compare three methods. In the first one we train without any rules, which forms the *supervised* baseline. In the second, *unnormalized*, we add the rules to the unlabeled data. This does not use any technique for dealing with the raven paradox. In the last one called *normalized*, we normalize MP and MT reasoning using Equation 18 for several different values of μ . The results in Table 1 are statistically significant when using a paired t-test.

	Precision types
Supervised	0.440 \pm 0.0013
Unnormalized	0.455 \pm 0.0014
Normalized $\mu = 0$	0.454 \pm 0.0015
Normalized $\mu = 0.1$	0.505 \pm 0.0014
Normalized $\mu = 0.25$	0.517 \pm 0.0013
Normalized $\mu = 0.5$	0.510 \pm 0.0013
Normalized $\mu = 0.75$	0.496 \pm 0.0012
Normalized $\mu = 1$	0.435 \pm 0.0015

Table 1: Results of the experiments. 20 runs using random splits of the data are averaged alongside 95% confidence intervals. All results are significant.

4 Analysis

Our experiments show that we can significantly improve on the classification of the types of objects for this problem. The normalized method in particular outperforms the unnormalized method, suggesting that explicitly dealing with the raven paradox is essential in this problem.

4.1 Gradient Updates

We analyze how the different methods handle the implication using the quantities d_φ^{MP} and d_φ^{MT} defined in Section 2.4. Figure 5 shows the average magnitude of d_φ^{MP} and d_φ^{MT} in the unnormalized model, which is computed by averaging over all training examples and formulas. This shows that the average MT gradient update is, in this problem, around 100 times larger than the average MP gradient update, i.e., it uses far more contrapositive reasoning. The unnormalized method acts very similar to the normalized one with $\mu \approx 0.01$.

Next, we will analyze how accurate our approach is at reasoning by comparing its 'decisions' to what should have been the correct 'decision'. We sample 2000 pairs of bounding boxes from the PASCAL-PART test set $\langle \mathcal{O}_t, \mathcal{W}_t \rangle$. We consider a pair of bounding boxes \mathbf{o} from an image i in the test set \mathcal{O}_t . $d_\varphi^{\text{MP}}(\mathbf{o})$ is a *correctly reasoned gradient* if both $\phi(\mathbf{o})$ and $\psi(\mathbf{o})$ are true in \mathbf{w}_i^t . Likewise, $d_\varphi^{\text{MT}}(\mathbf{o})$ is a correctly reasoned gradient if $\neg\psi(\mathbf{o})$ and $\neg\phi(\mathbf{o})$ are true in \mathbf{w}_i^t . Furthermore, we say that $d_\varphi^{\text{MP}}(\mathbf{o})$ is a *correctly updated gradient* if at least $\psi(\mathbf{o})$ is true in \mathbf{w}_i^t , and that $d_\varphi^{\text{MT}}(\mathbf{o})$ is correctly updated when $\neg\phi(\mathbf{o})$ is true in \mathbf{w}_i^t . Then the *correctly reasoned ratios* are computed using

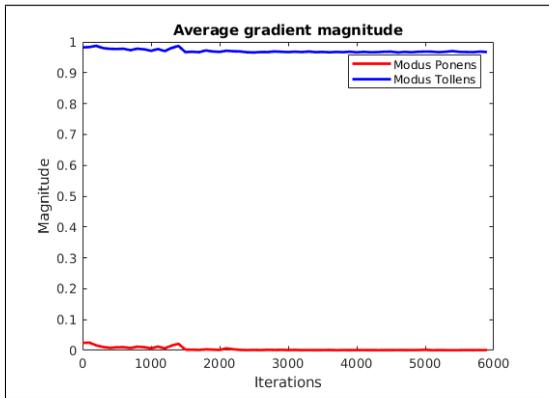


Figure 5: The average magnitude of Modus Ponens and Modus Tollens gradients.

$$c_{\text{I}}^{\text{MP}} = \frac{\sum_{\varphi \in \mathcal{K}} \sum_{\mathbf{o} \in D_i, D_i \in \mathcal{O}_t} v(\phi, \mathbf{w}_i^t) \cdot v(\psi, \mathbf{w}_i^t) \cdot d_\varphi^{\text{MP}}(\mathbf{o})}{\sum_{\varphi \in \mathcal{K}} \sum_{\mathbf{o}' \in D_i, D_i \in \mathcal{O}_t} d_\varphi^{\text{MP}}(\mathbf{o}')} \quad (19)$$

$$c_{\text{I}}^{\text{MT}} = \frac{\sum_{\varphi \in \mathcal{K}} \sum_{\mathbf{o} \in D_i, D_i \in \mathcal{O}_t} v(\neg\phi, \mathbf{w}_i^t) \cdot v(\neg\psi, \mathbf{w}_i^t) \cdot d_\varphi^{\text{MT}}(\mathbf{o})}{\sum_{\varphi \in \mathcal{K}} \sum_{\mathbf{o}' \in D_i, D_i \in \mathcal{O}_t} d_\varphi^{\text{MT}}(\mathbf{o}')} \quad (20)$$

The definition of the *correctly updated ratios* (c_{U}^{MP} and c_{U}^{MT}) are nearly the same. c_{U}^{MP} is found by removing the $v(\phi, \mathbf{w}_i^t)$ term from Equation 19, and c_{U}^{MT} by removing the $v(\neg\psi, \mathbf{w}_i^t)$ term from Equation 20.

Figure 6 shows the value of these ratios during training. The dotted lines that represent MT reasoning shows a convenient property, namely that is nearly always correct because of the large class imbalance. This could be the reason there is a significant benefit to adding contrapositive reasoning. Both normalized and unnormalized at $\mu = 1$ seems to get 'better' at reasoning during training, as the correctly updated ratios go up. After training for some

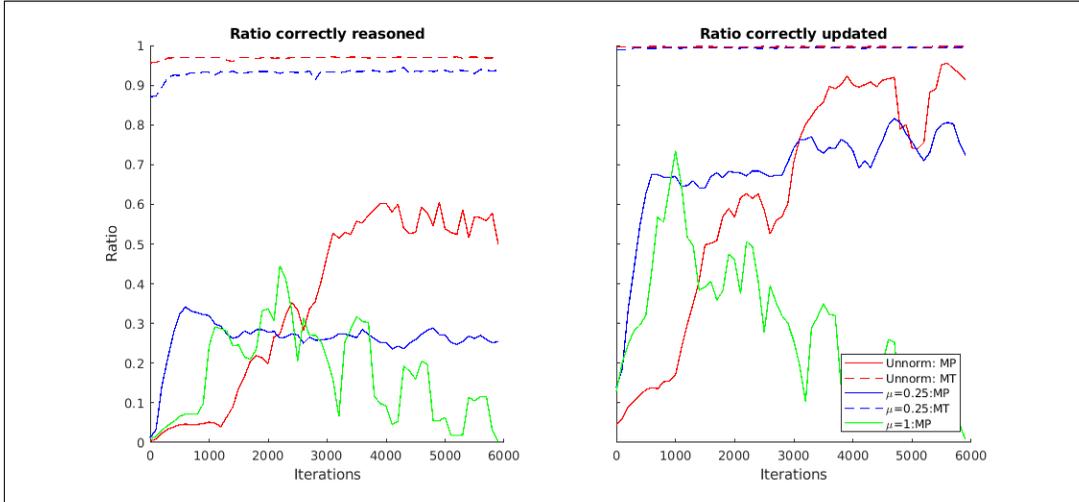


Figure 6: The left plot shows cr^{MP} and cr^{MT} and the right plot cu^{MP} and cu^{MT} for the Unnormalized method (denoted as Unnorm) and the Normalized methods with $\mu = 0.25$ and $\mu = 1$.

time, the unnormalized method seems to be best at reasoning correctly for both MP and MT. Another interesting observation is the difference between cr^{MP} and cu^{MP} . At many points, about half of the gradient magnitude correctly increases $\hat{p}(\psi|\mathbf{o}, \boldsymbol{\theta})$ because the model predicts a high value for $\hat{p}(\phi|\mathbf{o}, \boldsymbol{\theta})$, even though $\phi(\mathbf{o})$ is not actually in the test labels. It is interesting to see that, this kind of faulty reasoning which does lead to the right conclusion is actually beneficial for training.

Furthermore, disabling MT completely by setting μ to 1 seems to destabilize the reasoning. This is also reflected in the validation accuracy that seems to decline when cu^{MP} declines. This suggests that contrapositive reasoning is required to increase the amount of correct gradient updates.

5 Related work

5.1 Injecting Logic into Parameterized Models

Our work follows the recent works on Real Logic [26, 7], and the method we use is a special case of Real Logic with some additional changes. A particular difference is that the logic we employ has no function symbols, which was due to simplicity purposes. Injecting background knowledge into vector embeddings of entities and relations has been studied in [5, 6, 20, 21]. In particular, [22] has some similarities with Real Logic and our method. However, this

method is developed for regularizing vector embeddings instead of any parameterized model. In this sense, it can also be seen as a special case of Real Logic. Semantic Loss [31] is a very similar semi-supervised learning method. This loss is essentially Equation 4, which makes it more accurate than Product Real Logic, but also exponential in runtime. To deal with this, they compile SDD's [4] to make the computation tractable. A recent direction is DeepProbLog [17], a probabilistic version of Prolog with neural predicates that also uses SDD's. [11] also injects rules into a general model with a framework that transfers the logic rules using a so-called teacher network. This model is significantly different from the aforementioned ones, as it does not add a loss for each rule.

5.2 Semi-Supervised Learning

There is a large body of literature on semi-supervised methods [19, 2]. In particular, recent research on graph-based semi-supervised learning [14, 32, 33] relates unlabeled and labeled data through a graph structure. However, they do not use logically structured background knowledge. It is generally used for entity classification, although in [25] it is also used on link prediction. [16] introduced the surprisingly effective method Pseudo-Label that first trains a model using the labeled dataset, then labels the unlabeled dataset using this model and continues training on this newly labeled dataset. Our approach has a similar intuition in that we use the current model to get an estimation about the correct labels of the labeled dataset, and then use those labels to predict remaining labels, but the difference is that we use background knowledge to choose these labels.

6 Conclusion and Future Work

We proposed a novel semi-supervised learning technique and showed that it is possible to find labels for samples in an unlabeled dataset by evaluating them on relational background knowledge. Since implication is at the core of logical reasoning, we analyzed this by inspecting the gradients with respect to the antecedent and the consequent. Surprisingly, we discovered a strong imbalance between the contributions of updates from MP and MT in the induction process. It turned out that our approach is highly sensitive to the Raven paradox [10] requiring us to handle positive and contrapositive reasoning separately. Normalizing these different types of reasoning yields the largest improvements to the supervised baseline. Since it is quite general, we suspect that issues with this imbalance could occur in many systems that perform inductive reasoning.

We would like to investigate this phenomenon with different background knowledge and different datasets such as VisualGenome and ImageNet. In particular, we are interested in other approaches for modelling the implication like different Fuzzy Implications [12] or

by taking inspiration from Bayesian treatments of the Raven paradox [29]. Furthermore, it could be applied to natural language understanding tasks like semantic parsing.

References

- [1] Merrie Bergmann. *An introduction to many-valued and fuzzy logic: semantics, algebras, and derivation systems*. Cambridge University Press, 2008.
- [2] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. *Semi-supervised learning*. 2006.
- [3] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1971–1978, 2014.
- [4] Adnan Darwiche. SDD: A new canonical representation of propositional knowledge bases. *IJCAI International Joint Conference on Artificial Intelligence*, pages 819–826, 2011.
- [5] Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. Lifted rule injection for relation embeddings. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1389–1399. Association for Computational Linguistics, 2016.
- [6] Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. Regularizing relation representations by first-order implications. In *AKBC2016, the Workshop on Automated Base Construction*, pages 1–6, 2016.
- [7] Ivan Donadello, Luciano Serafini, and Artur S. d’Avila Garcez. Logic tensor networks for semantic image interpretation. In *IJCAI*, pages 1596–1602. ijcai.org, 2017.
- [8] Dorothy Edgington. Indicative conditionals. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2014 edition, 2014.
- [9] Ross Girshick. Fast r-cnn. *International Conference on Computer Vision*, pages 1440–1448, 2015.
- [10] Carl G Hempel. Studies in the logic of confirmation (i.). *Mind*, 54(213):1–26, 1945.
- [11] Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, and Eric Xing. Harnessing deep neural networks with logic rules. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2410–2420. Association for Computational Linguistics, 2016.
- [12] Balasubramaniam Jayaram and Michal Baczynski. *Fuzzy Implications*, volume 231. 2008.
- [13] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015.
- [14] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. 2016.
- [15] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting

- language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017.
- [16] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. 2013.
- [17] Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Deepproblog: Neural probabilistic logic programming. *arXiv preprint arXiv:1805.10872*, 2018.
- [18] Kevin Murphy, Yair Weiss, and Michael I. Jordan. Loopy Belief Propagation for Approximate Inference: An Empirical Study. pages 467–476, 2013.
- [19] Avital Oliver, Augustus Odena, Colin Raffel, Ekin D. Cubuk, and Ian J. Goodfellow. Realistic evaluation of semi-supervised learning algorithms. 2018.
- [20] Tim Rocktäschel. Combining representation learning with logic for language processing. *CoRR*, abs/1712.09687, 2017.
- [21] Tim Rocktäschel and Sebastian Riedel. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems*, pages 3791–3803, 2017.
- [22] Tim Rocktäschel, Sameer Singh, and Sebastian Riedel. Injecting logical background knowledge into embeddings for relation extraction. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1119–1129, 2015.
- [23] Dan Roth. On the hardness of approximate reasoning. *Artificial Intelligence*, 82(1-2):273–302, 1996.
- [24] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [25] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling Relational Data with Graph Convolutional Networks. In Aldo Gangemi, Roberto Navigli, Maria-Esther Vidal, Pascal Hitzler, Raphaël Troncy, Laura Hollink, Anna Tordai, and Mehwish Alam, editors, *The Semantic Web*, pages 593–607, Cham, 2018. Springer International Publishing.
- [26] Luciano Serafini and Artur d’Avila Garcez. Logic tensor networks: Deep learning and logical reasoning from data and knowledge. 2016.
- [27] Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. Reasoning with neural tensor networks for knowledge base completion. In *Advances in neural information processing systems*, pages 926–934, 2013.
- [28] Dirk Van Dalen. *Logic and structure*. Springer, 2004.
- [29] Peter B.M. Vranas. Hempel’s raven paradox: A lacuna in the standard Bayesian solution. *British Journal for the Philosophy of Science*, 55(3):545–560, 2004.
- [30] Gm Weiss and Foster Provost. The effect of class distribution on classifier learning: an empirical study. *Rutgers University*, (September 2001), 2001.
- [31] Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. A semantic

- loss function for deep learning with symbolic knowledge. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5502–5511, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- [32] Zhilin Yang, William W. Cohen, and Ruslan Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pages 40–48. JMLR.org, 2016.
- [33] Xiaojin Zhu, Zoubin Ghahramani, and John D Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In *Proceedings of the 20th International conference on Machine learning (ICML-03)*, pages 912–919, 2003.

Appendices

A Conditional Optimality of Product Real Logic

Considering only a single domain D of objects $x \in \mathbb{R}^D$, we have the Herbrand base \mathbf{A} . Let $\varphi \in \mathcal{K}$ be a set of function-free FOL formulas in Skolem-normal form. Furthermore, let $\mathcal{P} = \{\mathbf{P}_1, \dots, \mathbf{P}_K\}$ be a set of predicates which for ease of notation and without loss of generality we assume to all have the arity α .

Each ground atom $\mathbf{P}(\mathbf{o}) \sim \text{Bern}(f_{\mathbf{P}}^{\theta}(\mathbf{o}))$ is a binary random variable that denotes the binary truth value. It is distributed by a Bernoulli distribution with mean $f_{\mathbf{P}}^{\theta} \in \mathbb{R}^{\alpha \times D} \rightarrow [0, 1]$.

For each formula φ , we have the set of ground atoms $\mathbf{A}_{\varphi} \subseteq \mathbf{A}$ appearing in the instantiations of φ . Likewise, the assignment of truth values of \mathbf{A}_{φ} is \mathbf{w}_{φ} , which is a subset of the world \mathbf{w} . We can now express the joint probability, using Equation 1 and the valuation function defined in Section 2.1:

$$p(\mathcal{K}, \mathbf{w} | D, \theta) = p(\mathbf{w} | \theta) \cdot \prod_{\phi \in \mathcal{K}} v(\phi, \mathbf{w}_{\phi}) \quad (21)$$

We will first show that Product Real Logic is equal to this probability with two strong assumptions. The first is that the sets of ground atoms \mathbf{A}_{φ} are disjoint for all formulas in the corpus, i.e. if

$$\bigcup_{\varphi \in \mathcal{K}} \mathbf{A}_{\varphi} = \emptyset \quad (22)$$

The second is that the set of ground atoms used in two children (a direct subformula) of some subformula of a formula in \mathcal{K} are disjoint. If $pa(\phi)$ returns the parent of ϕ and $r(\phi)$ returns the root of ϕ (the formula highest up the tree), then

$$\mathbf{A}_{\phi} \cup \mathbf{A}_{\psi} = \emptyset, \forall \{\phi, \psi | pa(\phi) = pa(\psi) \wedge r(\phi) \in \mathcal{K}\} \quad (23)$$

First, we marginalize over the different possible worlds:

$$p(\mathcal{K}|D, \theta) = \sum_{\mathbf{w}} p(\mathbf{w}|\theta) \cdot \prod_{\varphi \in \mathcal{K}} v(\phi, \mathbf{w}_\varphi) \quad (24)$$

$$= \sum_{\mathbf{w}_{\varphi_1}} p(\mathbf{w}_{\varphi_1}|\theta) \cdot \left(\dots \cdot \sum_{\mathbf{w}_{\varphi_{|\mathcal{K}|}}} p(\mathbf{w}_{\varphi_{|\mathcal{K}|}}|\theta) \cdot \prod_{\varphi \in \mathcal{K}} v(\phi, \mathbf{w}_\varphi) \right) \quad (25)$$

$$= \sum_{\mathbf{w}_{\varphi_1}} p(\mathbf{w}_{\varphi_1}|\theta) \cdot v(\varphi_1, \mathbf{w}_{\varphi_1}) \cdot \left(\dots \cdot \sum_{\mathbf{w}_{\varphi_{|\mathcal{K}|}}} p(\mathbf{w}_{\varphi_{|\mathcal{K}|}}|\theta) \cdot v(\varphi_{|\mathcal{K}|}, \mathbf{w}_{\varphi_{|\mathcal{K}|}}) \right) \quad (26)$$

$$= \prod_{\varphi \in \mathcal{K}} \sum_{\mathbf{w}_\varphi} p(\mathbf{w}_\varphi|\theta) \cdot v(\phi, \mathbf{w}_\varphi) \quad (27)$$

where we make use of Equation 22 to join the summations, the independence of the probabilities of atoms from Equation 1 and marginalization of the atoms other than those in \mathbf{A}_φ .

We denote the set of instantiations of φ by S_φ , and a particular instance by s . $\mathbf{A}_s \subseteq \mathbf{A}_\varphi$ then is the set of ground atoms in s (and respectively for \mathbf{w}_s . Next we show that $\sum_{\mathbf{w}_\varphi} p(\mathbf{w}_\varphi|\theta) \cdot v(\varphi, \mathbf{w}_\varphi) = \prod_{s \in S_\varphi} \hat{p}(\varphi|s, \theta)$. As the formulas are in prenex normal form, $\varphi = \forall x_1, \dots, x_n \phi$. We find that, using Equation 23 and the same procedure as in Equations 24-27

$$\sum_{\mathbf{w}_\varphi} p(\mathbf{w}_\varphi|\theta) \cdot v(\varphi, \mathbf{w}_\varphi) = \sum_{\mathbf{w}_\varphi} p(\mathbf{w}_\varphi|\theta) \cdot \prod_{s \in S_\varphi} v(\phi, \mathbf{w}_s) \quad (28)$$

$$= \prod_{s \in S_\varphi} \sum_{\mathbf{w}_s} p(\mathbf{w}_s|\theta) \cdot v(\phi, \mathbf{w}_s). \quad (29)$$

Then, it suffices to show that $\sum_{\mathbf{w}_s} p(\mathbf{w}_s|\theta) \cdot v(\phi, \mathbf{w}_s) = \hat{p}(\phi|s, \theta)$. This is done using recursion. For brevity, we will only proof it for the \neg and \wedge connectives, as we can proof the others using those.

Assume that $\phi = \mathbf{P}(x_1, \dots, x_n)$. Then if $w_s(\mathbf{P}(x_1, \dots, x_n))$ is the binary random variable of the ground atom $\mathbf{P}(x_1, \dots, x_n)$ under the instantiation s ,

$$\sum_{\mathbf{w}_s} p(\mathbf{w}_s|\theta) \cdot v(\mathbf{P}(x_1, \dots, x_n), \mathbf{w}_s) \quad (30)$$

$$= \sum_{\mathbf{w}_s \setminus \{w_s(\mathbf{P}(x_1, \dots, x_n))\}} p(\mathbf{w}_s \setminus \{w_s(\mathbf{P}(x_1, \dots, x_n))\}|\theta) \cdot \quad (31)$$

$$\sum_{w_s(\mathbf{P}(x_1, \dots, x_n))} p(w_s(\mathbf{P}(x_1, \dots, x_n))|\theta) \cdot w_s(\mathbf{P}(x_1, \dots, x_n)) \quad (32)$$

$$= p(w_s(\mathbf{P}(x_1, \dots, x_n))|\theta) = \hat{p}(\mathbf{P}(x_1, \dots, x_n)|s, \theta). \quad (33)$$

Marginalize out all variables but $w_s(\mathbf{P}(x_1, \dots, x_n))$. $v(\mathbf{P}(x_1, \dots, x_n), \mathbf{w}_s)$ is 1 if $w_s(\mathbf{P}(x_1, \dots, x_n))$ is, and 0 otherwise.

Next, assume $\phi = \neg\psi$. Then

$$\sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\neg\psi, \mathbf{w}_s) \quad (34)$$

$$= \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot (1 - v(\psi, \mathbf{w}_s)) \quad (35)$$

$$= \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) - \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\psi, \mathbf{w}_s) \quad (36)$$

$$= 1 - \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\psi, \mathbf{w}_s) = \hat{p}(\neg\psi | s, \boldsymbol{\theta}) \quad (37)$$

Finally, assume $\varphi = \phi \wedge \psi$. Then

$$\sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\phi \wedge \psi, \mathbf{w}_s) \quad (38)$$

$$= \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\phi, \mathbf{w}_s) \cdot v(\psi, \mathbf{w}_s) \quad (39)$$

$$= \sum_{\mathbf{w}_{\phi_s}} \sum_{\mathbf{w}_{\psi_s}} p(\mathbf{w}_{\phi_s} | \boldsymbol{\theta}) \cdot p(\mathbf{w}_{\psi_s} | \boldsymbol{\theta}) \cdot v(\phi, \mathbf{w}_{\phi_s}) \cdot v(\psi, \mathbf{w}_{\psi_s}) \cdot \quad (40)$$

$$\sum_{\mathbf{w}_s \setminus (\mathbf{w}_{\phi_s} \cup \mathbf{w}_{\psi_s})} p(\mathbf{w}_s \setminus (\mathbf{w}_{\phi_s} \cup \mathbf{w}_{\psi_s}) | \boldsymbol{\theta}) \quad (41)$$

$$= \sum_{\mathbf{w}_{\phi_s}} p(\mathbf{w}_{\phi_s} | \boldsymbol{\theta}) \cdot v(\phi, \mathbf{w}_{\phi_s}) \cdot \sum_{\mathbf{w}_{\psi_s}} p(\mathbf{w}_{\psi_s} | \boldsymbol{\theta}) \cdot v(\psi, \mathbf{w}_{\psi_s}) \quad (42)$$

$$= \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\phi, \mathbf{w}_s) \cdot \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\psi, \mathbf{w}_s) \quad (43)$$

$$= \hat{p}(\phi | s, \boldsymbol{\theta}) \cdot \hat{p}(\psi | s, \boldsymbol{\theta}) = \hat{p}(\phi \wedge \psi | s, \boldsymbol{\theta}) \quad (44)$$

Using this result and equations 27 and 29, we find that

$$p(\mathcal{K} | D, \boldsymbol{\theta}) = \prod_{\varphi \in \mathcal{K}} \sum_{\mathbf{w}_\varphi} p(\mathbf{w}_\varphi | \boldsymbol{\theta}) \cdot v(\phi, \mathbf{w}_\varphi) \quad (45)$$

$$= \prod_{\varphi \in \mathcal{K}} \prod_{s \in S_\varphi} \sum_{\mathbf{w}_s} p(\mathbf{w}_s | \boldsymbol{\theta}) \cdot v(\phi, \mathbf{w}_s) \quad (46)$$

$$= \prod_{\varphi \in \mathcal{K}} \prod_{s \in S_\varphi} \hat{p}(\phi | s, \boldsymbol{\theta}) \quad (47)$$