

Differentiable Logic Programming for Distant Supervision

Akihiro Takemura^{a,*} and Katsumi Inoue^a

^aNational Institute of Informatics, Tokyo, Japan

ORCID (Akihiro Takemura): <https://orcid.org/0000-0003-4130-8311>, ORCID (Katsumi Inoue): <https://orcid.org/0000-0002-2717-9122>

Abstract. We introduce a new method for integrating neural networks with logic programming in Neural-Symbolic AI (NeSy), aimed at learning with distant supervision, in which direct labels are unavailable. Unlike prior methods, our approach does not depend on symbolic solvers for reasoning about missing labels. Instead, it evaluates logical implications and constraints in a differentiable manner by embedding both neural network outputs and logic programs into matrices. This method facilitates more efficient learning under distant supervision. We evaluated our approach against existing methods while maintaining a constant volume of training data. The findings indicate that our method not only matches or exceeds the accuracy of other methods across various tasks but also speeds up the learning process. These results highlight the potential of our approach to enhance both accuracy and learning efficiency in NeSy applications.



1 Introduction

Neural-Symbolic AI (NeSy) [14, 15] is a field of research aimed at combining neural networks with symbolic reasoning. While deep learning is capable of learning complex representations from input-output pairs, it requires a large amount of training data and struggles with tasks that require logical reasoning. On the other hand, learning with symbolic reasoning can be done with small amounts of data, but it is sensitive to noise and unable to handle non-symbolic data. In NeSy, it is crucial to combine the roles of neural networks and symbolic reasoning in a way that leverages their respective strengths.

There are various methods for implementing NeSy, including associating the continuous-valued parameters of neural networks (NN) with logical language and using the results of logical reasoning as the value of the loss function (e.g., semantic loss [27]). Also known are methods that combine symbolic solvers with neural networks, e.g., DeepProbLog [16] uses Problog for probabilistic logic programming, and NeurASP [28] uses clingo for answer set programming (ASP). These methods that internally call solvers often encapsulate computationally expensive problems such as weighted model counting or enumerating stable models in each iteration during learning.

Alternative methods have been proposed that embed inference traditionally done by symbolic reasoning solvers into vector spaces and perform symbolic reasoning using linear algebra [20]. One such method embeds logic programs into vector spaces and designs appropriate loss functions based on the semantics of non-monotonic reasoning to compute the results of reasoning in a differentiable manner

[2, 23]. However, these methods have issues such as not being able to directly handle logical constraints and not being applicable to neural network learning as is. Thus, in this paper, we propose a method that enables learning in neural networks for NeSy tasks using logical programs that include constraints.

Distant supervision is a method of generating labeled data for learning using rules, external data, or knowledge bases and was proposed by Mintz et al. [17] as a method to train classifiers for relation extraction based on information from knowledge bases. In NeSy, tasks where label information is provided through symbolic reasoning are commonly used, with MNIST Addition [16] being a representative task. In this task, pairs of handwritten digits are input, and the goal is to learn the classification of handwritten digits with the sum of the digits provided as the label (e.g.,  +  = 5). Unlike the usual MNIST classification, in MNIST Addition, the labels are not given for each image individually. In this case, the relationship between the sum given as the label and the digits corresponding to the images is expected to be provided through symbolic reasoning.

In this paper, we propose a novel architecture for NeSy systems [16, 28] that integrates differentiable logic programming [2, 23] and neural networks. This paper makes the following contributions:

1. We propose a novel architecture that integrates neural networks with logic programming through a differentiable approach. This method facilitates the direct evaluation of logical implications and constraints using differentiable operations, thus enabling effective learning under distant supervision without relying on symbolic solvers for reasoning about missing labels.
2. We demonstrate through experiments with a constant volume of training data that our proposed method not only matches but, in some cases, exceeds the accuracy of existing approaches that utilize symbolic solvers. Moreover, we achieved a significant reduction in the training time for neural networks, highlighting substantial gains in computational efficiency.

The structure of this paper is as follows. After the preliminaries in Section 2, Section 3 introduces the logic programming semantics in vector spaces. Section 4 presents our proposed method for using differentiable logic programming for distant supervision. Section 5 presents the results of experiments and comparison to the state of the art NeSy methods. Section 6 covers the related works in the literature. Finally, Section 7 presents the conclusion.

* Corresponding Author. Email: atakemura@nii.ac.jp

2 Preliminaries

A normal logic program P is a set of rules of the form:

$$A \leftarrow A_1 \wedge \cdots \wedge A_m \wedge \neg A_{m+1} \wedge \cdots \wedge \neg A_n \quad (1)$$

where A and $A_i (n \geq m \geq 0)$ are atoms. In this paper, the terms ‘normal logic program’, ‘logic program’, and ‘program’ are used interchangeably. An *atom* is a predicate with some arity, e.g., $p(X, Y)$, where variables are represented by upper case characters, and predicates and constants are represented by lower case characters. A literal is either an atom p , or its negation $\neg p$. The atom A in (1) is the *head* and $\{A_1, \dots, A_n\}$ is the *body* of a rule. For each rule R_i of the form (1), define $\text{head}(R_i) = A$, $\text{body}^+(R_i) = \{A_1, \dots, A_m\}$ and $\text{body}^-(R_i) = \{A_{m+1}, \dots, A_n\}$.

The *Herbrand universe* of a logic program P is the set of all ground terms in the language of P , i.e., terms composed of function symbols and constants that appear in P . The *Herbrand base* B_P is the set of atoms that can be formed from the relations of the program and terms in the Herbrand universe. We assume that the Herbrand base B_P of a program to be lexicographically ordered.

A rule with an empty body is a *fact*. A program P is *definite* if no rule in P contains negation as failure. A program, a rule, or an atom is ground if it is variable free. A program P is semantically identified with its ground instantiation, $\text{ground}(P)$, by substituting variables in P by elements of its Herbrand universe in every possible way.

An *interpretation* $I \subseteq B_P$ satisfies a rule R_i of the form (1) if $\text{body}^+(R_i) \subseteq I$ and $\text{body}^-(R_i) \cap I = \emptyset$ imply $A \in I$. An interpretation that satisfies every rule in a program P is a *model* of the program. A model of a program is *supported* if for each atom $p \in I$, there exists a ground rule such that I satisfies its body [1]. A model M is minimal if there is no model J of P such that $J \subset M$. A definite program has a unique minimal model, which is the least model.

Given a normal logic program P and an interpretation I , the *reduct* P^I , which is a ground definite program, is constructed as follows: a ground rule $A \leftarrow L_1, \dots, L_m$ is in P^I iff there is a ground rule of the form (1) such that $\text{body}^-(R_i) \cap I = \emptyset$. If the least model of P^I is identical to I , then I is a *stable model* of P [13]. For a definite program, the stable model coincides with the least model. A stable model is always supported, but the converse does not hold in general.

Supported models can be computed as the models of *Clark’s completion* [5]. Let $\text{heads}(P, a)$ be the set of rules in P whose head is a . The completion of P , denoted $\text{comp}(P)$, is the set of clauses

$$a \leftrightarrow \bigvee_{R_i \in \text{heads}(P, a)} \text{body}(R_i) \quad (2)$$

for all $a \in B_P$. A model of $\text{comp}(P)$ is a supported model of P [1].

Let $I \subseteq B_P$ be an interpretation of P . The relation \models is defined as follows: for a rule R_i of the form (1), I *satisfies* R_i if $\text{head}(R_i) \cap I \neq \emptyset$ whenever $\text{body}(R_i) \subseteq I$, and denoted as $I \models R_i$; for a program P , I satisfies P if $I \models R_i$ for all $R_i \in P$; for a formula $F = F_1 \vee \cdots \vee F_k$ ($k \geq 0$), $I \models F$ iff there is a F_i ($k \geq i \geq 1$) such that $I \models F_i$, i.e., the empty disjunction is false. Let $\text{comp}(R_p)$ denote the completed rule ($p \leftrightarrow \text{body}(R_{p1}) \vee \cdots \vee \text{body}(R_{pj})$) for the atom p , then $p \in I$ iff $I \models \text{comp}(R_p)$.

3 Semantics

In this section, we consider the semantics of ground normal logic programs in vector spaces. First, we introduce the necessary notations. Matrices are denoted using bold uppercase letters (\mathbf{M}), and

vectors are denoted using bold lowercase letters (\mathbf{v}). The element in the i -th row and j -th column of a matrix is denoted by \mathbf{M}_{ij} , and the i -th element of a vector is denoted by \mathbf{v}_i . The slice of the i -th row of a matrix is denoted by $\mathbf{M}_{i:}$, and the slice of the j -th column is denoted by $\mathbf{M}_{:j}$. Variables are denoted by upper case letters, and constants and predicates are denoted by lower case letters; e.g., in $\text{sum}(L)$, L is a variable and $\text{sum}/1$ is a predicate with arity 1.

3.1 Embedding Normal Logic Programs

Given a ground normal logic program P , we introduce two matrices that jointly represent the program. The program matrix represents the bodies of the rules in the program, and the head matrix represents their disjunctions. This is an alternative formulation to the embedding approach described by Sakama et al. [20].

Definition 1 (Program Matrix). *Let P be a ground normal logic program with R rules and the size of its Herbrand base be $|B_P| = N$. Then P is represented by a binary matrix $\mathbf{Q} \in \{0, 1\}^{R \times 2N}$ such that i -th row corresponds to the body of the i -th rule R_i : $\mathbf{Q}_{ij} = 1$ if $a_j \in \text{body}^+(R_i)$, $\mathbf{Q}_{i(N+j)} = 1$ if $a_j \in \text{body}^-(R_i)$, and $\mathbf{Q}_{ij} = 0$ otherwise.*

Definition 2 (Head Matrix). *Let $\mathbf{D} \in \{0, 1\}^{(N \times R)}$ be the head matrix associated with P . Then the element $\mathbf{D}_{ji} = 1$ if the head of rule R_i ($1 \leq i \leq R$) is a_j ($1 \leq j \leq N$), and 0 otherwise.*

Example 1. Consider the following program P_1 with 3 rules:

$$(R_1) a \leftarrow c \wedge \neg b \quad (R_2) a \leftarrow a \quad (R_3) b \leftarrow \neg a \quad (3)$$

P_1 is encoded into a pair of matrices (\mathbf{Q}, \mathbf{D}) :

$$\mathbf{Q} = \begin{matrix} & a & b & c & \neg a & \neg b & \neg c \\ \begin{matrix} R_1 \\ R_2 \\ R_3 \end{matrix} & \begin{pmatrix} 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix} \end{matrix} \quad \mathbf{D} = \begin{matrix} & a & b & c \\ \begin{matrix} R_1 \\ R_2 \\ R_3 \end{matrix} & \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \end{matrix} \quad (4)$$

\mathbf{Q} represents the bodies of the rules, which are the conjunctions of the literals appearing in the bodies. For example, $\mathbf{Q}_{1:}$ represents the body of R_1 , $(c \wedge \neg b)$. \mathbf{D} represents the disjunctions of the bodies of the rules sharing the same head. For example, $\mathbf{D}_{1:}$ represents the disjunction $\text{body}(R_1) \vee \text{body}(R_2) = (c \wedge \neg b) \vee a$. Together, \mathbf{Q} and \mathbf{D} represent the logic program P .

3.2 Evaluating Embedded Normal Logic Programs

We consider the conjunction appearing in the bodies of the rules as the negation of disjunctions of negated literals using De Morgan’s law, i.e., $L_1 \wedge \cdots \wedge L_n = \neg(\neg L_1 \vee \cdots \vee \neg L_n)$. This means that when evaluating the body of a rule, instead of checking whether all literals hold (as in [23]), we can count the number of false literals and check whether the count exceeds 1. To this end, we introduce a piecewise linear function $\min_1(x) = \min(x, 1) = \text{ReLU}(1 - x)$, which gives 1 for $x \geq 1$. This function is almost everywhere differentiable (except at $x = 1$), which allows gradient-based optimization to be applied effectively.

To evaluate normal logic programs in vector spaces, we introduce the vectorized counterparts of interpretation and model.

Definition 3 (Interpretation Vector). *Let P be a ground normal logic program. An interpretation $I \subseteq B_P$ is represented by a binary vector $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_N)^T \in \mathbb{Z}^N$ where each element \mathbf{v}_i ($1 \leq i \leq N$)*

represents the truth value of the proposition a_i such that $\mathbf{v}_i = 1$ if $a_i \in I$, otherwise $\mathbf{v}_i = 0$. We assume propositional variables share the common index such that \mathbf{v}_i corresponds to a_i , and we write $\text{idx}(\mathbf{v}_i) = a_i$.

Definition 4 (Complementary Interpretation Vector). *The complementary interpretation vector $\mathbf{w} \in \mathbb{Z}^{2N}$ is a binary vector, which is a concatenation of the interpretation vector \mathbf{v} and its complement: $\mathbf{w} = [\mathbf{v}; \mathbf{1}_N - \mathbf{v}]$.*

Proposition 1. (Embedding Models of Normal Logic Programs) *Let $\mathbf{P} = (\mathbf{Q}, \mathbf{D})$ be an embedding of a ground normal logic program P , $\text{dist}(\cdot, \cdot)$ be a distance function in a metric space, \mathbf{v} be an interpretation vector representing $I \subseteq B_P$, and \mathbf{w} be its complementary interpretation vector. Then, for an interpretation vector \mathbf{v} ,*

$$I \models \text{comp}(P) \text{ iff } \text{dist}\left(\mathbf{v}, \min_1\left(\mathbf{D}\left(\mathbf{1} - \min_1(\mathbf{Q}(\mathbf{1} - \mathbf{w}))\right)\right)\right) = 0$$

Proof. (Sketch; full proof in the Appendix [24].) A row slice of the program matrix \mathbf{Q}_i corresponds to the body of a rule R_i , so the matrix-vector products $\mathbf{Q}_i \mathbf{w}$ and $\mathbf{Q}_i(\mathbf{1} - \mathbf{w})$ computes the number of true and false literals in I , respectively. The conjunctions can be computed as the negation of disjunctions of negated literals using De Morgan's law, i.e., $\mathbf{1} - \min_1(\mathbf{Q}_i(\mathbf{1} - \mathbf{w}))$.

Let $\{R^{a_i}\} = \{a_i \leftarrow B_j, \dots, a_i \leftarrow B_t\}$ be the set of rules that share the same head atom a_i , where B_j denote the rule bodies, and $B_j \vee \dots \vee B_t$ be the disjunction of the rule bodies. By construction of the head matrix \mathbf{D} , $\mathbf{D}_i(\mathbf{1} - \min_1(\mathbf{Q}(\mathbf{1} - \mathbf{w})))$ computes the number of true rule bodies that share the same head. Thus, $\mathbf{h}_i = \min_1\left(\mathbf{D}_i(\mathbf{1} - \min_1(\mathbf{Q}(\mathbf{1} - \mathbf{w})))\right) = 1$ if there is at least one rule body that is true in I and in the disjunction $B_j \vee \dots \vee B_t$, and 0 otherwise. Then computing \mathbf{h}_i corresponds to the evaluation of $I \models \text{comp}(R_{a_i})$. This can be generalized to the entire matrix.

Let $\mathbf{h}^I = \min_1\left(\mathbf{D}(\mathbf{1} - \min_1(\mathbf{Q}(\mathbf{1} - \mathbf{w})))\right)$, then the second part of the iff relation is simplified to $\text{dist}(\mathbf{v}, \mathbf{h}^I) = 0$.

- If $I \models \text{comp}(P)$, then $\text{dist}(\mathbf{v}, \mathbf{h}^I) = 0$.
Suppose $I \models \text{comp}(R_{a_i})$, then there is at least one rule body that is true in I , so $\mathbf{h}_i^I = 1$. Otherwise, when we have $I \not\models \text{comp}(R_{a_i})$, $\mathbf{h}_i^I = 0$. Therefore, it holds that $\mathbf{h}_i^I = \mathbf{v}_i$, and since the index i is arbitrary, we have $\mathbf{v} = \mathbf{h}_i^I$, i.e., $\text{dist}(\mathbf{v}, \mathbf{h}^I) = 0$.
- If $\text{dist}(\mathbf{v}, \mathbf{h}^I) = 0$, then $I \models \text{comp}(P)$.
Consider $\text{dist}(\mathbf{v}, \mathbf{h}^I) = 0$. For $\mathbf{h}_i^I = 1$, there is at least one rule body that is true in I , and for $\mathbf{h}_i^I = 0$, there is no rule body that is true in I . Since we have $\mathbf{v}_i = \mathbf{h}_i^I$, for $\mathbf{v}_i = \mathbf{h}_i^I = 1$, $a_i \leftrightarrow \bigvee_{R_j \in \text{heads}(P, a_i)} \text{body}(R_j)$ is satisfied and denote $I \models \text{comp}(R_{a_i})$, and for $\mathbf{v}_i = \mathbf{h}_i^I = 0$ denote $I \not\models \text{comp}(R_{a_i})$. Since the index i is arbitrary, we conclude $I \models \text{comp}(P)$.

□

Example 2. (Example 1 contd.) Consider the program P_1 and corresponding matrices \mathbf{Q} and \mathbf{D} from Example 2. This program has 2 supported models $\{\{a\}, \{b\}\}$. Let $\mathbf{v}^{\{a\}} = (1 \ 0 \ 0)^\top$ represent the interpretation $\{a\}$, and \mathbf{w} be its complementary interpretation vector. We have $\mathbf{h}^I = (\mathbf{1} \ 0 \ 0)^\top = \mathbf{v}$, and take the Euclidean distance: $\text{dist}(\mathbf{v}, \mathbf{h}^I) = \sqrt{\sum_{i=1}^3 (\mathbf{v}_i - \mathbf{h}_i^I)^2} = 0$. Therefore, according to Proposition 1, $I \models \text{comp}(P_1)$.

The vector $\mathbf{h} = \min_1\left(\mathbf{D}(\mathbf{1} - \min_1(\mathbf{Q}(\mathbf{1} - \mathbf{w})))\right)$ serves as the head vector, which is an indicator vector representing true atoms following the evaluation of rule bodies in the logic program. This will be used later to define the loss function in Section 4.1.4.

3.3 Embedding and Evaluating Constraints

A constraint is a rule with an empty head, e.g., $\leftarrow a \wedge b$ represents a constraint where a and b must not both be true simultaneously. Since constraints are rules in a program, we embed them into a constraint matrix \mathbf{C} in the same manner as the program matrix \mathbf{P} . Note that we do not require the head matrix because constraints have empty heads.

Definition 5 (Constraint Matrix). *Let $C = \{C_1, \dots, C_k\} = \{\leftarrow B_1, \dots, \leftarrow B_k\}$ be the set of constraints in a program P with $|B_P| = N$. Then the matrix corresponding to the constraints is $\mathbf{C} \in \{0, 1\}^{(k \times 2N)}$ such that i -th row corresponds to the body of the i -th constraint C_i : $\mathbf{C}_{ij} = 1$ if $a_j \in \text{body}^+(C_i)$, $\mathbf{C}_{i(N+j)} = 1$ if $a_j \in \text{body}^-(C_i)$ and $\mathbf{C}_{ij} = 0$ otherwise.*

To evaluate the constraints, we check whether the bodies of the constraint rules are in I : given a constraint R_i , if $\text{body}(R_i) \subseteq I$ then the constraint is violated; otherwise it is satisfied.

Proposition 2. (Evaluating Constraints) *Let \mathbf{C} be an embedding of constraints C , $\text{dist}(\cdot, \cdot)$ be a distance function in a metric space, \mathbf{v} be an interpretation vector representing $I \subseteq B_P$, and \mathbf{w} be its complementary interpretation vector. Then, for an interpretation vector \mathbf{v} , it holds that $I \not\models C$ iff $\text{dist}(\mathbf{1}, \min_1(\mathbf{C}(\mathbf{1} - \mathbf{w}))) = 0$.*

Proof. Proved similarly to Proposition 1. Let $\mathbf{c}^I = \min_1(\mathbf{C}(\mathbf{1} - \mathbf{w}))$. Consider the i -th constraint C_i and the corresponding row slice \mathbf{C}_i . The existence of at least one false literal in the body is computed by $\mathbf{c}_i^I = \min_1(\mathbf{C}_i(\mathbf{1} - \mathbf{w}))$, where $\mathbf{c}_i^I = 1$ if there is a false literal and $\mathbf{c}_i^I = 0$ otherwise, i.e., when $\mathbf{c}_i^I = 0$, the body is satisfied and the constraint is violated.

Suppose $I \not\models C_i$, then there is at least one false literal in the body of C_i , so $\min_1(\mathbf{C}_i(\mathbf{1} - \mathbf{w})) = 1$. Repeat this for all $C_i \in C$, we obtain a 1-vector, which means there is at least one false literal in the bodies of all constraints. By definition, $\text{dist}(\mathbf{1}, \mathbf{1}) = 0$. The converse can be proved similarly. □

Example 3. Consider the constraint $\leftarrow a \wedge b$. Then we have:

$$\begin{array}{cccccc} & a & b & c & \neg a & \neg b & \neg c \\ \mathbf{C} = \mathbf{C}_1 & (1 & 1 & 0 & 0 & 0 & 0) \end{array} \quad (5)$$

We examine two scenarios: in the first, the constraint is violated, and in the second, it is not. Let $\mathbf{c}^I = \min_1(\mathbf{C}(\mathbf{1} - \mathbf{w}))$ and $\text{dist}(\cdot, \cdot)$ denote the Euclidean distance, then, for the following cases,


- $\mathbf{v}^{\{a,b\}} = (1 \ 1 \ 0)$: we obtain $\mathbf{c}^{\{a,b\}} = (0)$, and $\text{dist}(\mathbf{1}, \mathbf{c}^{\{a,b\}}) = 1$, so we conclude that the constraint is violated.
- $\mathbf{v}^{\{a\}} = (1 \ 0 \ 0)$: we obtain $\mathbf{c}^{\{a\}} = (1)$, and $\text{dist}(\mathbf{1}, \mathbf{c}^{\{a,b\}}) = 0$, so we conclude $I \not\models C$.

For later use in the loss function (Section 4.1.4), we define the constraint violation vector \mathbf{c}' as $\mathbf{c}' = \mathbf{1} - \mathbf{c} = \mathbf{1} - \min_1(\mathbf{C}(\mathbf{1} - \mathbf{w}))$. Intuitively, this modification turns \mathbf{c} into an indicator vector, where each element $\mathbf{c}'_i = 1$ means that the i -th constraint is violated.

4 Learning with Differentiable Logic Program

In this section, we show how the aforementioned differentiable logic program semantics can be used to train neural networks. Although our method supports both implication and constraint rules, it is not always necessary to use both of them for learning, as we shall show later in the experimental section. Specifically, for the NeSy tasks we studied, using exclusively either one of implication or constraint rules is enough to achieve competitive accuracy. On the other hand, in NeurASP [28] for example, the observation atoms are typically given as integrity constraints in ASP rules to compute stable models, and implication rules are defined similarly to ours. Consequently, we included a combination of both implication rules and constraints in our experiments to provide a thorough evaluation.

4.1 Example: MNIST Addition

The MNIST digit addition problem [16] is a simple distant supervision task that is commonly used in neural-symbolic literature. The input consists of two images, and the output is their sum (e.g., , 9). The goal is to learn digit classification from the sum of digits, rather than from images with individually labeled digits, as in the usual MNIST digit classification. For brevity, we shall focus on the single digit additions in this section.

We first introduce *neural predicates* [16], which act as an interface between the neural and logic programming parts. More concretely, a neural predicate is a predicate which can take references to tensorized objects as arguments. A *neural atom* is an instance of a neural predicate, representing a specific combination of values or variables. For example, in MNSIT Addition, we define two neural predicates, *obs*/4 and *label*/3. *obs*(i_1, D_1, i_2, D_2) represents a situation where images i_1 and i_2 were classified as digits D_1 and D_2 ranging from 0 to 9, respectively. *label*(i_1, i_2, S) represents the two images and their sum S ranging from 0 to 18. Thus, we obtain 100 *obs* and 19 *label* neural atoms.

4.1.1 Implication rules

In general, we expect the label atoms to appear in the heads of the implication rules so that we can compare the output of the logic programming component with the label using a loss function such as binary cross entropy. For a single digit MNIST Addition, the label is the integer values between 0 and 18 represented by the *label*/3 predicate. As for the individual rules, we enumerate the possible combinations of digits that sum to the given label, e.g.,

$$\begin{aligned} \text{label}(i_1, i_2, 0) &\leftarrow \text{obs}(i_1, 0, i_2, 0). \\ \text{label}(i_1, i_2, 1) &\leftarrow \text{obs}(i_1, 0, i_2, 1). \\ \text{label}(i_1, i_2, 1) &\leftarrow \text{obs}(i_1, 1, i_2, 0). \\ &\dots \\ \text{label}(i_1, i_2, 18) &\leftarrow \text{obs}(i_1, 9, i_2, 9). \end{aligned} \quad (6)$$

In this way, 100 rules with *label*/3 in the heads can be instantiated, which results in the embedded program $\mathbf{P}_{impl.} = (\mathbf{Q}_{impl.}, \mathbf{D}_{impl.})$ where $\mathbf{Q}_{impl.} \in \{0, 1\}^{(100 \times 200)}$ and $\mathbf{D}_{impl.} \in \{0, 1\}^{(19 \times 100)}$.

4.1.2 Constraints

In the case of MNIST Addition, constraints can be represented with smaller number of rules compared to the implication rules, e.g.,

$$\begin{aligned} &\leftarrow \text{label}(i_1, i_2, 0) \wedge \neg \text{obs}(i_1, 0, i_2, 0). \\ &\leftarrow \text{label}(i_1, i_2, 1) \wedge \neg \text{obs}(i_1, 0, i_2, 1) \wedge \neg \text{obs}(i_1, 1, i_2, 0). \\ &\dots \\ &\leftarrow \text{label}(i_1, i_2, 18) \wedge \neg \text{obs}(i_1, 9, i_2, 9). \end{aligned} \quad (7)$$

Intuitively, one can read the first rule as “when the label is 0, both of the digits must be 0”. This essentially amounts to enumerating all possible combinations of digits that sum to the label and adding them as negative literals to the rule bodies. Preparing constraints for each label results in a constraint matrix $\mathbf{C} \in \{0, 1\}^{(19 \times 238)}$.

4.1.3 Handling Neural Network Outputs

Here, the outputs of the neural network combined with facts evident from the problem and labels are treated as a continuous interpretation. Facts evident from the problem refer to, for example, the three digits (d_1, d_2, d_3) in the Apply2x2 task or the given digit number in the Member task. The details of the Apply2x2 and Member tasks will be explained later in the experiment section. The label information is also incorporated into this continuous-valued interpretation vector.

Definition 6 (Continuous-valued Interpretation Vector). Let $\mathbf{x} \in [0, 1]^N$ be the output passed through the last layer of the neural network. Let $\mathbf{f} \in \{0, 1\}^N$ represent facts from the problem that are not dependent on NN’s output, and $\mathbf{l} \in \{0, 1\}^N$ represent the label information available from the instance. If the number of elements in \mathbf{x} , \mathbf{f} or \mathbf{l} is less than N , pad appropriately with zeros. The continuous-valued interpretation vector \mathbf{z} is computed as follows: $\mathbf{z} = \mathbf{x} + \mathbf{f} + \mathbf{l}$.

In MNIST Addition, facts are not available from the problem settings, so we only focus on the neural network outputs and labels. In this task, the inputs to the (convolutional) neural network are two images (i_1, i_2). After passing through the Softmax activation, we obtain two output vectors $\mathbf{x}_1, \mathbf{x}_2 \in [0, 1]^{10}$ as (probabilistic) outputs. To map these vectors to 100 *obs* neural atoms, we compute their Cartesian product and obtain $\mathbf{x} \in [0, 1]^{100}$: $\mathbf{x} = \mathbf{x}_1 \times \mathbf{x}_2$. Depending on the problem, the dimension of the continuous-valued interpretation vector may be different for implication and constraints. In MNIST Addition, for implication rules it is not necessary to have *label* in the interpretation vector, as they are the heads of the rules. On the other hand, it is necessary to include *label* in the interpretation vector for evaluating constraints, as they are present in the rule bodies. It is imperative that the indexing remain consistent across all vectors and matrices; otherwise we risk neural network outputs being mapped to incorrect neural atoms.

The learning pipeline is shown in Figure 1. Firstly, the input images are classified using the CNN, and we associate its probabilistic output with neural atoms *obs*/4. Then, using the neural atoms and label information, we obtain the logical loss which can be used to train the CNN with gradient backpropagation.

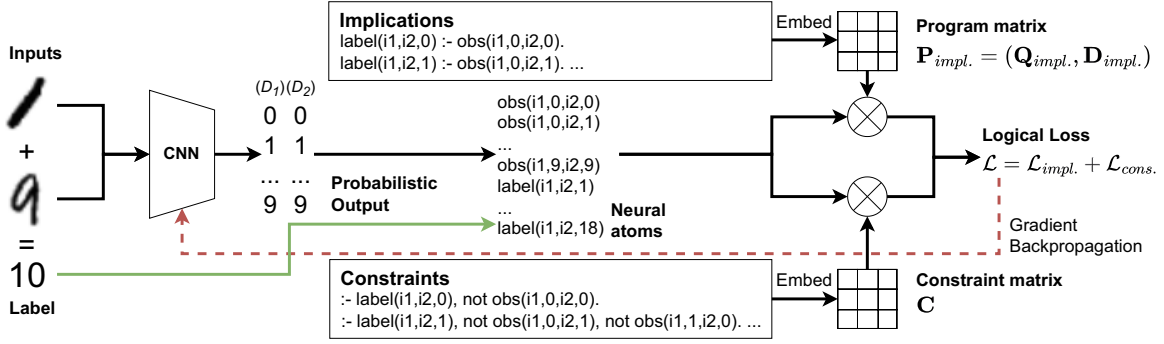


Figure 1. The learning pipeline for MNIST Addition.

4.1.4 Loss Function

Using the embedded program, continuous-valued interpretation vector and label information, the loss function is defined as follows:

$$\mathbf{h} = \min_1 \left(\mathbf{D} \left(\mathbf{1} - \min_1 (\mathbf{Q}(\mathbf{1} - \mathbf{w}_z)) \right) \right) \quad (8)$$

$$\mathbf{c}' = \mathbf{1} - \min_1 (\mathbf{C}(\mathbf{1} - \mathbf{w}_z)) \quad (9)$$

$$\mathcal{L} = \mathcal{L}_{impl.} + \mathcal{L}_{cons.} = BCE(\mathbf{h}, \mathbf{lb}) + BCE(\mathbf{c}', \mathbf{0}) \quad (10)$$

where BCE stands for binary cross-entropy, \mathbf{lb} corresponds to the label vector, and $\mathbf{0}$ is a zero vector with the same dimension as \mathbf{c}' . Note that \mathbf{w}_z here is the complementary interpretation vector based on the continuous-valued interpretation vector \mathbf{z} which contains neural atoms, facts and labels: $\mathbf{w}_z = [\mathbf{z}; \mathbf{1}_N - \mathbf{z}]$. When the continuous valued interpretation vector is a binary one, \mathbf{h} corresponds to the interpretation I . If all atoms that are supposed to be true in I are true, then the BCE loss will be 0. \mathbf{c}' corresponds to an indicator vector for the violated constraints. Thus, \mathbf{c}' should be all 0 when all constraints are satisfied. Combining the aforementioned BCE's, the loss function will be 0 iff all implied neural atoms are true, and all constraints are satisfied.

The difference between programs (6) and (7) lies in their structure: implication rules may contain label neural atoms in the head, whereas constraints always contain label neural atoms in the body. Implication rules present a more straightforward approach for representing partial information in the form of logical rules, especially in scenarios where employing intermediate predicates is necessary or enumerating constraints is time-consuming. In the context of the NeSy tasks examined in this study, we observed a consistent pattern where implication rules are typically comprised of straightforward forward inference rules, and constraints are formed from a fewer number of rules. Each of these constraints contains a label neural atom alongside a series of negated observation atoms in their body.

Based on the definition of the combined loss function, it is clear that only one is necessary for accomplishing the MNIST Addition task. The first BCE is essentially the same as the one for a 19-label multiclass classification task (labels spanning from 0 to 18), while the second BCE corresponds to a multiclass classification with an all-0 label.

The evaluation of implication rules ensures that if the correct label's corresponding atom is derived as the head, $\mathcal{L}_{impl.}$ becomes 0. Similarly, if all constraints are satisfied, then $\mathcal{L}_{cons.}$ becomes 0. Since both the evaluation of implication rules and constraints are defined in a (almost-everywhere) differentiable manner, it is possible to



train the neural network using this loss function through backpropagation.

5 Experiments


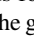


5.1 Task Description

We studied the learning performance on the following NeSy tasks.

MNIST Addition [16]

The input consists of two MNIST images of digits (i_1, i_2), and the output is their sum (e.g.,  +  = 9). The goal is to learn image classification from the sum of digits, rather than from images with individually labeled digits, as in the usual MNIST classification. This paper deals with two types: single-digit additions (two images), and two-digit additions (four images).




ADD2x2 [12]

The input is four MNIST images of digits ($i_{11}, i_{12}, i_{21}, i_{22}$) arranged in a 2x2 grid, and the output is four sums (s_1, s_2, s_3, s_4) calculated from each row and column of the grid (e.g.,     = 1, 8, 3, 6). The goal is to learn the classification problem of MNIST images from the four sums provided as labels.

APPLY2x2 [12]

The input consists of three numbers (d_1, d_2, d_3) and four handwritten operator images (o_1, o_2, o_3, o_4) arranged in a 2x2 grid, with the output being the results (r_1, r_2, r_3, r_4) of operations performed along each row and column of the grid. The operators are one of $\{+, -, \times\}$. For example, the result for the first row is calculated as $r_1 = (d_1 \text{ op}_{11} d_2) \text{ op}_{12} d_3$ (e.g., $1, 2, 4, \times, +, -, \times, 6, -4, -2, 12$). The goal is to learn the classification of handwritten operators from the four results and three numbers given as labels.

MEMBER(n) [25]

For $n=3$, the input consists of three images (i_1, i_2, i_3) and one number (d_1), with the output being a boolean indicating whether d_1 is included in the three images (e.g.,    = 4, 0). For $n=5$, the input includes five images (i_1, \dots, i_5). The goal is to learn the classification problem of MNIST images from the numbers provided as labels. This paper deals with two types: $n=3$ and $n=5$.

5.2 Implementation and Experimental Setup

The methods introduced in the previous section were implemented in PyTorch. The convolutional neural network (CNN) used in the experiments is the same as those in [16] and [28]. The experimental environment is an AMD Ryzen 7950X (16c/32t), 128GB RAM, and an NVIDIA A4000 (16GB), with settings to utilize the GPU as much as possible. The number of training data was 30,000 and 15,000 for Addition 1 and 2, respectively, and 10,000 for other tasks. Unless otherwise noted, the number of epochs and batch size for all tasks were set to 1, and Adam was used with a learning rate of 0.001. Each experiment consisted of 5 repeated trials, and the average is reported. The timeout was set to 30 minutes per trial.

5.3 Results

Table 1 shows the dimensions of the program matrix **P**, head matrix **D**, and constraint matrix **C**. In terms of the number of rules in respective programs, the constraint matrix is usually smaller than the implication matrix, because each constraint rule (row) is often an enumeration of conditions for the target atom, whereas in the implication matrix, multiple rules (rows) can share the same head. The rate of growth of the program matrix is highly task dependent; for example, adding two digits to the MNIST Addition task result in adding 10,000 elements to the matrix, whereas adding two digits to the set in the Member task results in adding a few thousand elements.

Table 1. Dimensions of program matrices

	P (implication)	D (head)	C (constraint)
Addition 1	100 × 200	19 × 100	19 × 238
Addition 2	10000 × 20000	199 × 10000	199 × 20398
Add2x2	400 × 800	76 × 400	76 × 952
Apply2x2	11979 × 2680	10597 × 11979	10597 × 23874
Member 3	40 × 60	20 × 40	40 × 100
Member 5	60 × 100	20 × 60	60 × 140

The accuracy and training time are reported in tables 2 and 3, respectively. In the tables, I indicates training using only implication rules, C indicates training using only constraints, and I+C indicates training using both implication rules and constraints. DPL, DSL and NASP are abbreviations for DeepProbLog [16], DeepStochLog [26] and NeurASP [28], respectively.

Table 2 shows the average accuracy of MNIST digit classification and math operator classification. It can be seen that while the proposed method achieved comparable accuracy to the comparison methods in MNIST Addition and Add2x2, it significantly outperformed the comparison methods in Apply2x2. However, in Member 3, the accuracy was lower. Comparing the use of implication rules, constraints, or both, there was no significant difference except for Member(n), where using either implication rules or constraints alone tended to result in higher accuracy than using both. In particular, for Member 5 tasks, combining implication rules and constraints led to significantly worse performance compared to individual applications. This suggests that we might need to modify the combination loss or consider introducing a scheduler for improving training performance.

Table 3 shows the average training times of each method. Except for Apply2x2, it is evident that the proposed method can learn faster than existing methods. Especially in Member 5, where comparison methods timed out, the proposed method processed 10,000 training instances in about 13 seconds, showing a significant speed difference. The long training time for Apply2x2 can be attributed to the rather

Table 2. Accuracy on digit and operator classification. The numbers in parentheses indicate timeouts (30 min).

Accuracy (%)	Comparisons			Ours		
	DPL	DSL	NASP	I	C	I+C
Addition 1	97.8	95.8	97.7	97.7	97.5	97.4
Addition 2	97.7	97.8	97.8	97.5	97.9	97.8
Add2x2	T/O(5)	98.0	97.5	97.7	97.6	97.9
Apply2x2	87.8	87.8	80.9	99.5	99.4	99.4
Member 3	92.3(3)	92.9	91.7	87.8	87.0	84.6
Member 5	T/O(5)	T/O(5)	T/O(5)	86.3	86.4	69.5

naive implementation used in this experiment. In our implementation, Apply2x2 required a more fine-grained control over indexing, necessitating the use of less efficient 'for-loops' to avoid neural network outputs being mapped to incorrect neural atoms. In contrast, in other tasks, the more efficient vectorized operations were used to compute the Cartesian product.

Table 3. Learning times of each method. The numbers in parentheses indicate timeouts (30 min).

Time (sec)	Comparisons			Ours		
	DPL	DSL	NASP	I	C	I+C
Addition 1	470.7	20.8	84.7	31.4	31.4	35.9
Addition 2	1120	33.6	283.5	83.5	19.9	87.4
Add2x2	T/O(5)	35.8	131.1	16.9	16.2	18.0
Apply2x2	323.3	127.7	25.5	154.4	228.5	359.5
Member 3	1782(3)	398.4	191.9	10.3	10.3	11.9
Member 5	T/O(5)	T/O(5)	T/O(5)	11.6	11.7	13.4




Figure 2. Test accuracy (%) during training.

Figure 2 shows the test accuracy during training.¹ We see that

¹ This is from a different set of experiments, as such, the results presented in this figure do not necessarily match those in Table 2 and Table 3. The timeout was set to 2 hours per trial.

the addition tasks (MNIST Addition and Add2x2) can be handled equally well by our method, NeurASP and DeepStochLog. The difference is more pronounced for Apply2x2 and Member3, and in Apply2x2, we observe that DeepStochLog plateaus quickly while NeurASP fluctuates around 80%. In Member 3, NeurASP leads in terms of accuracy up to 5,000 iterations, other methods catch up after 9,000 iterations. Finally, it is interesting that the implication rule only method performs well in Member 5 compared against constraint only or the combination of both. This suggests that there might be tasks that can be learned better by certain types or combinations of programs, although it is difficult to know beforehand which one would perform the best.

6 Related Work

Xu et al. [27] introduced the semantic loss that leverages sentential decision diagrams for efficient loss calculation, enabling effective learning in both supervised and semi-supervised settings. LTN [4] and LNN [18] embed first-order logic formulas using fuzzy logic within neural network architectures. Our approach is similar to the semantic loss, in the sense that we evaluate the neural output using a differentiable loss function to train the neural network by backpropagation. On the other hand, our approach does not require weighted model counting, nor direct embedding of logical operations, and the logic programming semantics is kept separate from the perception neural network.

Parallel to direct integration strategies, significant work has been conducted on coupling neural outputs with symbolic solvers. For instance, DeepProbLog [16], DeepStochLog [26], NeurASP [28], and NeuroLog [25] facilitate inference by integrating the neural outputs as probabilities or weighted models within a symbolic solver. While these approaches can utilize logical constraints and background knowledge represented by logical programs, the computational cost of symbolic reasoning can become a bottleneck during learning. Distinct from the aforementioned coupling approaches, the focus shifts more towards reasoning rather than simultaneous learning in neural-symbolic systems proposed by Eiter et al. [8] and Embed2Sym [3]. In the pipeline proposed by Eiter et al. [8], the neural network is trained separately, and predictions that pass a predefined confidence threshold are then translated into logic programs for reasoning with ASP.

Various neural-symbolic approaches have been developed for symbolic rule learning, with varying degree of integration between logical reasoning and neural computation. Neural Theorem Provers [19] and ∂ ILP [9] allow end-to-end differentiable learning of explanatory rules from data. Similarly, frameworks like NeuroLog [25], NSIL [6], and the Apperception Engine [10] integrate symbolic solvers to enhance the reasoning processes, utilizing the output from neural networks represented as logical constructs. Additionally, α ILP [22] and DFORL [11] extend the capabilities of inductive logic programming (ILP) by making the learning process differentiable. In particular, α ILP is designed for visual scene understanding, while DFORL focuses on relational data, demonstrating the versatility of differentiable ILP in handling diverse data types. In Neural Logic Machines [7], the architecture of the neural network itself is designed to mimic logical reasoning processes, thereby learning to approximate symbolic rules. In contrast to these methods, which predominantly focus on enhancing the logic programming aspects within neural frameworks, our approach focuses more on enhancing the training of neural networks themselves, and we leave the inductive learning of logic programs for future work.

Recent advancements have explored various methods for representing logic programming semantics in vector spaces, each with distinct embedding strategies and computational techniques. Notably, Sakama et al. [20] proposed a method to embed logic programs into vector spaces, and compute logic programming semantics using linear algebra. More specifically, their method allows for the computation of stable models as fixpoints through repeated tensor multiplications. However, a significant limitation from the standpoint of neural-symbolic integration is the non-differentiability of their method, which complicates direct integration with neural networks. Other matrix-based approaches by Sato et al. [21] and Takemura and Inoue [23] use binary program matrices to represent logic programs. While Sato et al. [21]'s method is non-differentiable, the one proposed by Takemura and Inoue [23] is based on a differentiable loss function, where the interpretation vector is treated as the input to this loss function, and the loss is minimized by updating the interpretation vector using the gradient of the loss function. Similarly, Aspis et al. [2] proposed a method based on a root-finding algorithm, which presents yet another computational strategy in this domain. In contrast to these existing methods, our approach computes supported models in vector spaces without imposing restrictive conditions on program structures that require the rewriting of same head rules, for example. The method presented in this paper does not minimize the loss function by updating the interpretation vector using the gradient information, and the loss is reduced by updating the parameters of the neural network to predict the correct intermediate labels in a distant supervision setting. To the best of our knowledge, no other implementation exists that utilizes the differentiable logic programming semantics for training neural networks in neural-symbolic settings.

7 Conclusion

We proposed a method to assist the learning of neural networks with logic programs and verified its effectiveness in NeSy tasks. This method is based on a differentiable logic programming semantics, where continuous-valued interpretation vectors contain outputs of neural networks, and evaluation of implication rules and constraints are incorporated into a loss function, enabling the learning of neural networks under the distant supervision settings. The experimental results showed that it is possible to achieve comparable accuracy to those based on symbolic solvers, and with an exception, the proposed method completed the neural network training much faster than the existing methods.

The findings of this study demonstrate the effectiveness of the approach based on the differentiable logic programming semantics for enabling high-accuracy and fast learning in NeSy. Future work includes applying the proposed method to more complex NeSy tasks. Additionally, a more detailed analysis of the balance between constraints and implication rules is necessary. In the tasks addressed in this study, there was little difference between using constraints and implication rules, but this might vary by task. Another limitation is that the semantics as presented currently is only valid for crisp 0-1 interpretations, so other continuous-valued interpretations do not necessarily have natural meanings associated with them. To this end, it might be interesting to seek connections between the differentiable logic programming semantics and fuzzy logic.

Acknowledgements

This work has been supported by JST CREST JPMJCR22D3 and JSPS KAKENHI JP21H04905.

References

- [1] K. R. Apt, H. A. Blair, and A. Walker. Towards a theory of declarative knowledge. In *Foundations of Deductive Databases and Logic Programming*, pages 89–148. Elsevier, 1988.
- [2] Y. Aspis, K. Broda, A. Russo, and J. Lobo. Stable and Supported Semantics in Continuous Vector Spaces. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning*, pages 59–68, 2020. doi: 10.24963/kr.2020/7.
- [3] Y. Aspis, K. Broda, J. Lobo, and A. Russo. Embed2Sym - Scalable Neuro-Symbolic Reasoning via Clustered Embeddings. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*, 2022.
- [4] S. Badreddine, A. d’Avila Garcez, L. Serafini, and M. Spranger. Logic Tensor Networks. *Artificial Intelligence*, 303:103649, Feb. 2022. ISSN 0004-3702. doi: 10.1016/j.artint.2021.103649.
- [5] K. L. Clark. Negation as Failure. In *Logic and Data Bases*, pages 293–322. Springer US, Boston, MA, 1978. ISBN 978-1-4684-3384-5. doi: 10.1007/978-1-4684-3384-5_11.
- [6] D. Cunningham, M. Law, J. Lobo, and A. Russo. Neuro-Symbolic Learning of Answer Set Programs from Raw Data. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 3586–3596. ijcai.org, 2023. doi: 10.24963/IJCAI.2023/399.
- [7] H. Dong, J. Mao, T. Lin, C. Wang, L. Li, and D. Zhou. Neural Logic Machines. In *International Conference on Learning Representations (ICLR)*, page 22, 2019.
- [8] T. Eiter, N. Higuera, J. Oetsch, and M. Pritz. A Neuro-Symbolic ASP Pipeline for Visual Question Answering. *Theory Pract. Log. Program.*, 22(5):739–754, 2022. doi: 10.1017/S1471068422000229.
- [9] R. Evans and E. Grefenstette. Learning Explanatory Rules from Noisy Data. *Journal of Artificial Intelligence Research*, 61:1–64, Jan. 2018. ISSN 1076-9757. doi: 10.1613/jair.5714.
- [10] R. Evans, M. Bosnjak, L. Buesing, K. Ellis, D. P. Reichert, P. Kohli, and M. J. Sergot. Making sense of raw input. *Artificial Intelligence*, 299:103521, 2021. doi: 10.1016/j.artint.2021.103521.
- [11] K. Gao, K. Inoue, Y. Cao, and H. Wang. A differentiable first-order rule learner for inductive logic programming. *Artificial Intelligence*, 331:104108, June 2024. ISSN 0004-3702. doi: 10.1016/j.artint.2024.104108.
- [12] A. L. Gaunt, M. Brockschmidt, N. Kushman, and D. Tarlow. Differentiable Programs with Neural Libraries. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1213–1222, July 2017.
- [13] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *ICLP/SLP*, volume 88, pages 1070–1080, 1988.
- [14] P. Hitzler and M. K. Sarker, editors. *Neuro-Symbolic Artificial Intelligence: The State of the Art*, volume 342 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2022. ISBN 978-1-64368-244-0. doi: 10.3233/FAIA342.
- [15] P. Hitzler, M. K. Sarker, and A. Eberhart, editors. *Compendium of Neurosymbolic Artificial Intelligence*, volume 369 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2023. ISBN 978-1-64368-406-2. doi: 10.3233/FAIA369.
- [16] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, and L. De Raedt. DeepProbLog: Neural Probabilistic Logic Programming. In *Advances in Neural Information Processing Systems 31*, pages 3749–3759, 2018.
- [17] M. Mintz, S. Bills, R. Snow, and D. Jurafsky. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011, Aug. 2009.
- [18] R. Riegel, A. Gray, F. Luus, N. Khan, N. Makondo, I. Y. Akhalwaya, H. Qian, R. Fagin, F. Barahona, U. Sharma, S. Ikbali, H. Karanam, S. Neelam, A. Likhyan, and S. Srivastava. Logical Neural Networks. *arXiv preprint: arXiv:2006.13155*, June 2020. doi: 10.48550/arXiv.2006.13155.
- [19] T. Rocktäschel and S. Riedel. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems*, pages 3788–3800, 2017.
- [20] C. Sakama, K. Inoue, and T. Sato. Logic programming in tensor spaces. *Annals of Mathematics and Artificial Intelligence*, 89:1133–1153, Aug. 2021. ISSN 1573-7470. doi: 10.1007/s10472-021-09767-x.
- [21] T. Sato, C. Sakama, and K. Inoue. From 3-valued Semantics to Supported Model Computation for Logic Programs in Vector Spaces. In *12th International Conference on Agents and Artificial Intelligence*, pages 758–765, Sept. 2020. ISBN 978-989-758-395-7.
- [22] H. Shindo, V. Pfanschilling, D. S. Dhami, and K. Kersting. α ILP: Thinking visual scenes as differentiable logic programs. *Mach. Learn.*, 112(5):1465–1497, 2023. doi: 10.1007/S10994-023-06320-1.
- [23] A. Takemura and K. Inoue. Gradient-Based Supported Model Computation in Vector Spaces. In *Logic Programming and Nonmonotonic Reasoning*, pages 336–349, 2022. ISBN 978-3-031-15707-3. doi: 10.1007/978-3-031-15707-3_26.
- [24] A. Takemura and K. Inoue. Differentiable Logic Programming for Distant Supervision. *arXiv preprint: arXiv:2408.12591*, Aug. 2024. doi: 10.48550/arXiv.2408.12591. Full version of this paper.
- [25] E. Tsamoura, T. Hospedales, and L. Michael. Neural-Symbolic Integration: A Compositional Perspective. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5051–5060, May 2021. ISSN 2374-3468.
- [26] T. Winters, G. Marra, R. Manhaeve, and L. D. Raedt. DeepStochLog: Neural Stochastic Logic Programming. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 10090–10100. AAAI Press, 2022. doi: 10.1609/AAAI.V36I9.21248.
- [27] J. Xu, Z. Zhang, T. Friedman, Y. Liang, and G. V. den Broeck. A semantic loss function for deep learning with symbolic knowledge. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, volume 80, pages 5498–5507, 2018.
- [28] Z. Yang, A. Ishay, and J. Lee. NeurASP: Embracing Neural Networks into Answer Set Programming. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 1755–1762, July 2020. ISBN 978-0-9992411-6-5. doi: 10.24963/ijcai.2020/243.