

Inductive Learning of Complex Knowledge from Raw Data

Daniel Cunningham^{1,2}, Mark Law^{2,3}, Jorge Lobo⁴ and Alessandra Russo²

¹IBM Research Europe, Winchester, UK

²Imperial College London, London, UK

³ILASP Limited, Grantham, UK

⁴Universitat Pompeu Fabra, Barcelona, Spain

Abstract

One of the ultimate goals of Artificial Intelligence is to learn generalised and human-interpretable knowledge from raw data. Neuro-symbolic reasoning approaches partly tackle this problem by improving the training of a neural network using a manually engineered symbolic knowledge base. In the case where symbolic knowledge is *learned* from raw data, this knowledge lacks the expressivity required to solve complex problems. In this paper, we introduce Neuro-Symbolic Inductive Learner (NSIL), an approach that trains a neural network to extract latent concepts from raw data, whilst learning symbolic knowledge that solves complex problems, defined in terms of these latent concepts. The novelty of our approach is a method for biasing a symbolic learner to learn improved knowledge, based on the in-training performance of both neural and symbolic components. We evaluate NSIL on two problem domains that require learning knowledge with different levels of complexity, and demonstrate that NSIL learns knowledge that is not possible to learn with other neuro-symbolic systems, whilst outperforming baseline models in terms of accuracy and data efficiency.

Keywords

neuro-symbolic learning, inductive logic programming, raw data

1. Introduction

Within Artificial Intelligence, one of the ultimate goals is to learn generalised and human interpretable knowledge from raw data. Neuro-symbolic approaches tackle this problem by combining the best features of both neural networks and symbolic learning and reasoning techniques [1, 2, 3]. For example, many existing approaches improve the training of a neural network using a symbolic knowledge base that is manually engineered [4, 5, 6, 7]. In contrast, systems that learn symbolic knowledge are generally only applied to structured data, and use pre-trained neural networks when applied to raw data. To address this limitation, *Meta_Abd* jointly trains a neural network whilst learning symbolic knowledge [8], using a meta-interpreter learner that only learns first-order definite logic programs [9]. Consequently, the learned symbolic knowledge lacks the expressivity required for many common-sense learning and reasoning tasks [10].

AAAI 2022 FALL SYMPOSIUM SERIES, *Thinking Fast and Slow and Other Cognitive Theories in AI*, November 17-19, 2022, Westin Arlington Gateway in Arlington, Virginia, USA

 dancunnington@uk.ibm.com (D. Cunningham)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

In this paper, we present Neuro-Symbolic Inductive Learner (NSIL), that jointly trains a neural network whilst learning a complex, first-order symbolic knowledge, called an *inductive hypothesis*. Building upon state-of-the-art symbolic learners [11, 12], NSIL generalises *Meta_Abd* by learning more expressive knowledge from raw data in the language of Answer Set Programming (ASP) [13], solving computationally harder problems that *Meta_Abd* cannot solve. In NSIL, the learned hypothesis maps latent concepts to downstream labels, and the neural network is trained to classify these latent concepts from raw data, by reasoning over the learned hypothesis. This is a form of neuro-symbolic reasoning, where the neural network is trained in a weakly supervised fashion, as no labels are given for the latent concepts. In NSIL, we use the NeurASP system [7] to perform such computation, and seamlessly integrate symbolic learning by creating *corrective examples* for the symbolic learner. These examples use the current predictions of the neural network and the space of possible latent concept values relevant to the downstream label. A specific mechanism for weighting these examples is proposed, which biases the symbolic learner towards *exploration* or *exploitation* of the inductive hypothesis search space. NSIL is the first system that jointly trains a general¹ neural network from scratch, whilst inducing a complex first-order hypothesis in the form of an ASP program.

Our experiments apply NSIL to two tasks: *MNIST Arithmetic* and *MNIST Hitting Sets*, where each training sample contains a collection of raw MNIST inputs [15, 16], together with a downstream label. We include specific relations as domain knowledge for each task, and compare NSILs performance with purely differentiable baseline models. The results show that NSIL; (1) outperforms the baselines in terms of its overall accuracy and data efficiency, (2) trains the neural network to predict latent concepts with comparable accuracy to fully supervised differentiable models, and (3) learns interpretable knowledge. In comparison with *Meta_Abd*, the MNIST Hitting Sets task [17] demonstrates that our NSIL approach is capable of learning inductive hypotheses of increased expressivity, including constraints and choice rules. The learned hypotheses indeed solve such computationally harder problems, and, go a step further by generating all the hitting sets of a given collection. This can only be achieved when the learned hypothesis is expressed using ASP, which is clearly not possible with *Meta_Abd*.

2. Related work

Many neuro-symbolic learning and reasoning techniques have been proposed in the literature [18]. Their primary objective is to leverage properties of symbolic methods to enhance data efficiency, transferability and interpretability of neural models. Some techniques inject symbolic knowledge directly into the neural architectures [4, 6], whereas others preserve a clear distinction between a neural and symbolic reasoning component [5, 7]. The key drawback of these neuro-symbolic reasoning approaches is that they require complete and manually engineered symbolic background knowledge. In contrast, our approach learns symbolic knowledge, expressed as a first-order rule-based program while training a neural network.

Pure symbolic learning systems are capable of learning interpretable knowledge in a data efficient manner [19, 20, 21]. However, as these systems are logic-based, they can only learn from structured data. Even recent differentiable methods [22, 23, 24] are only applied to structured

¹[14] is limited to end-to-end training of a *binary* neural network only, in conjunction with rule learning.

data, and use pre-trained neural networks when applied to raw data [25, 26]. To address this limitation, neuro-symbolic *learning* systems have been proposed [8, 27, 14].

Meta_Abd [8] extends [27] by using abduction and induction to jointly train a neural network and induce definite logic programs from raw data. However, the Metagol symbolic learner [9], used by *Meta_Abd*, can only learn symbolic knowledge expressed as definite logic programs without function symbols, which can compute only polynomial functions [10]. Hence, *Meta_Abd* cannot learn more expressive knowledge involving defaults, exceptions, constraints and choice, which are essential aspects of common-sense learning and reasoning. Similarly to *Meta_Abd*, we also train a neural network whilst learning symbolic knowledge, through weak supervision from downstream labels. However, we learn first-order complex knowledge expressed in the language of ASP, which is more general than symbolic learning of definite clauses [28, 11, 29], and can solve computationally harder problems, such as the hitting sets NP decision problem [17]. In this task, the training set consists of binary True/False labels that indicate the existence of a hitting set over a collection of up to four sets, but our learned programs are general as they can decide if there is a hitting set over collections of *any* number of sets. This is not possible with *Meta_Abd*, or indeed any neural network architecture, since at inference time their input/output format must be exactly the same as that observed during training. Furthermore, our learned programs not only solve the hitting set decision problem, but can also generate *all* the hitting sets of a given collection. Finally, *Meta_Abd* may be vulnerable to becoming stuck in a local optima, which justifies why the authors present in [8] an evaluation using a CNN pre-trained in a one-shot fashion. In contrast, we demonstrate convergence to the optimal hypothesis over very large hypothesis spaces without using ground-truth concept labels. This is also the case when a reduced training set is used, for which the neural network remains under-trained.

Finally, the approach in [14] uses Binary Neural Networks (BNN) encoded into an ASP program, together with an apperception theory as background knowledge. The task of jointly learning the weights of the BNN and the additional knowledge required to solve a given problem, is formalised as solving a single ASP problem, and as such is fully delegated to an ASP solver. BNNs are an active area of research but they are still far from solving the kind of problems that regular neural networks can solve. Hence, since ASP doesn't support floating point arithmetic, the approach in [14] is limited by the BNN and is unlikely to handle tasks that require complex neural network architectures.

3. Neuro-Symbolic Inductive Learner

3.1. Problem setting

We consider a dataset of samples $D = \{ \langle X, y \rangle, \dots \}$ and a domain knowledge B , where X is a collection of size m containing multiple raw data $x_i \in \mathcal{X}$, $y \in \mathcal{Y}$ is a target label, and B is a first-order logic program. For a given task, we assume a single *latent concept* $C = \langle n_C, \mathcal{Z}_C \rangle$, where n_C is the name of the latent concept, and \mathcal{Z}_C is its set of possible values, which we assume to be finite. Each $x_i \in X$ has an associated latent concept value $z_i \in \mathcal{Z}_C$, and Z is a collection of size m denoting the combination of latent concept values associated with X . Let \mathcal{L} be the language defined using the relations in B , C , and \mathcal{Y} . Informally, we wish to learn a complex symbolic knowledge H , called an *inductive hypothesis*, expressed as a first-order logic

program in the language \mathcal{L} , that defines the label y for a sample X in terms of its latent concept values Z , and domain knowledge B . Formally, the objective of NSIL is to learn a composite function of the form $h^* \circ f^* : \mathcal{X}^m \rightarrow \mathcal{Y}$ where $f^* : \mathcal{X}^m \rightarrow \mathcal{Z}_C^m$ and $h^* : \mathcal{Z}_C^m \cup B \rightarrow \mathcal{Y}$. We learn approximations f of f^* and h of h^* , and refer to them as the *neural* and *symbolic learning* components of NSIL respectively. Under our single latent concept assumption, f is effectively a classifier for the latent concept C in each input. Therefore, to simplify notation, we refer to the set of possible latent concept values \mathcal{Z}_C as \mathcal{Z} .

During training, our approach learns from a dataset D and a domain knowledge B . It is crucial to note that D does not contain the correct latent labels Z for any of the $\langle X, y \rangle$ samples. To solve the task, NSIL has to; (1) learn how to correctly map each raw input x_i to its latent concept value z_i , and (2) learn how to symbolically relate the combination of latent concept values Z predicted from X to the label y . These mappings correspond to our f and h functions respectively. As h is purely symbolic, available domain knowledge B can be leveraged during learning. However, h is not differentiable, which makes the joint optimisation of $h \circ f$ through standard gradient-based learning algorithms infeasible. The challenge is to train f and h via a weakly supervised learning task using D and B only, such that for every $\langle X, y \rangle \in D$, $h(f(X) \cup B) = y = h^*(f^*(X) \cup B)$.

NSIL addresses this challenge iteratively, using two steps. f is trained using the current inductive hypothesis, denoted H' , and the symbolic learner is biased to learn a new H' using the current f and h . Let us now describe the NSIL architecture, defining the neural and symbolic learning components.

3.2. NSIL architecture

In this section, we instantiate the learning problem tackled by NSIL and present a novel neuro-symbolic learning architecture. The neural component f uses a neural network $g : \mathcal{X} \rightarrow [0, 1]^{|\mathcal{Z}|}$, parameterised by θ , which computes a probability distribution over the latent concept values for each $x_i \in X$. f then uses the standard arg max function to aggregate the output of g for each input, into a collection of latent concept values Z . The symbolic learning component h learns a complex inductive hypothesis H expressed in the first-order logic programming language \mathcal{L} such that for every $\langle X, y \rangle \in D$, $B \cup H \cup Z$ satisfies y .

The *Meta_{Abd}* approach [8], addresses a similar problem of training a neural network whilst learning an inductive hypothesis. Informally, *Meta_{Abd}* induces an inductive hypothesis H , and uses H to “abduce” relevant facts (i.e., pseudo-labels), to prune the search space for Z while g is trained with respect to these computed pseudo-labels. The *Meta_{Abd}* approach relies upon a definition of $B \cup H \cup Z$ covering y in terms of logical entailment (i.e., $P(y|B, H, Z) = 1$ if $B \cup H \cup Z \models y$ and 0 otherwise), and the strong (implicit) assumption that $B \cup H \cup Z$ defines a unique *logical* model. Such assumption is related to the confined expressivity of the symbolic learner of *Meta_{Abd}*, which can only learn inductive hypotheses expressed as first-order definite clauses. Consequently, hypotheses learned by *Meta_{Abd}* can only solve tasks requiring polynomial functions [10]. In contrast, we are able to learn more complex functions, such as in the following example:

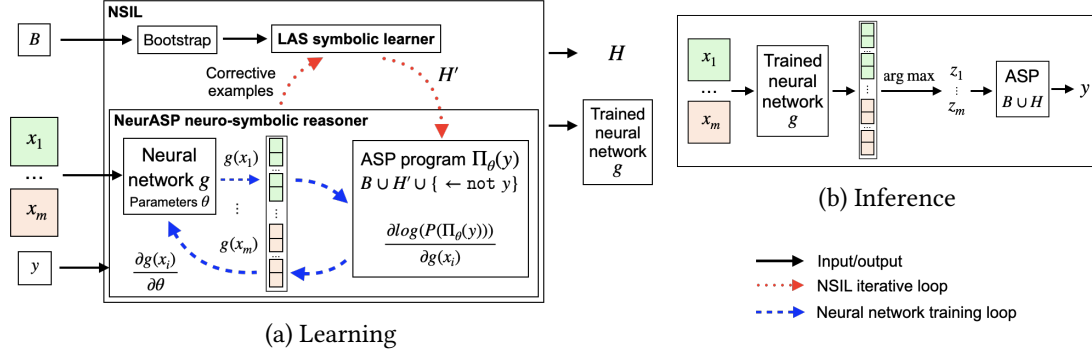


Figure 1: NSIL learning with a single data point $\langle X, y \rangle$, and inference with a single input X .

Example 1. Consider a set \mathcal{X} of MNIST images from classes 1-5. The task is to decide whether a collection of sets X , containing images from \mathcal{X} , has a hitting set of size ≤ 2 . For example, if $X = \{\{\mathbf{2}, \mathbf{7}\}, \{\mathbf{1}\}, \{\mathbf{3}\}\}$, then the answer is yes, because $\{1, 3\}$ is a hitting set. If $X = \{\{\mathbf{1}, \mathbf{2}\}, \{\mathbf{3}\}, \{\mathbf{4}\}\}$ then the answer is no, because there are no hitting sets. The training data D contains pairs, $\langle X, y \rangle$, where $y = 1$ if there is a hitting set of X , otherwise $y = 0$. Also, the domain knowledge B defines the notion of an element in a set.

Our system will not only learn an inductive hypothesis H that solves the hitting sets NP decision problem, but also the same hypothesis will be able to generate all the possible hitting sets of a given collection X . These kind of problems cannot be solved by *Meta-Abd*. NSIL provides a general neuro-symbolic learning architecture, where the learned inductive hypothesis is a type of first-order logic program called an *answer set program*, based on the ASP paradigm [13]. ASP is a rule-based declarative formalism for representing and solving combinatorial problems. Its highly expressive language allows for rules representing non-determinism, choice and constraints, among other constructs. Knowledge about a real-world problem can be represented as an ASP program, whose *answer sets*, or logical models, correspond to the solutions of the original problem. In the context of the hitting set decision problem in ASP, stating that $B \cup H \cup Z$ covers y means that there is at least one answer set of $B \cup H \cup Z$ if $y = 1$, or no answer sets if $y = 0$ [13].

ASP programs can be learned from noisy, labelled, symbolic examples [30], and can also be used to train neural networks by applying symbolic reasoning over semantic constraints [7]. As illustrated in Figure 1, the NSIL architecture leverages these two aspects of ASP, by seamlessly integrating two components: (1) a Learning from Answer Sets (LAS) symbolic learner [30], for learning an inductive hypothesis H , that maps latent concept values Z to outputs y ; (2) NeurASP [7] to train the neural network to predict Z from X , by applying neuro-symbolic reasoning over the learned H . Let us now define the optimisation criteria for these components.

Informally, let $E_{Z, Y}$ be a set of examples depending on Z and Y (see Subsection 3.3 for details on how $E_{Z, Y}$ is dynamically generated during training). Each example $e \in E_{Z, Y}$ has an associated penalty e_{pen} that indicates the cost that a candidate inductive hypothesis H' pays for not covering it. Let B be an ASP domain knowledge, defining also the set \mathcal{H} of possible learnable ASP programs. An inductive *hypothesis* $H \in \mathcal{H}$ is an ASP program with $size(H) = |H|$ given

by the number of relations in H , and a *cost* given by the sum of the penalties of uncovered examples. We denote the set of examples uncovered by H with $UNCOV(H, E_{\mathcal{Z}, \mathcal{Y}})$, and define the cost of H as $pen(H, (B, E_{\mathcal{Z}, \mathcal{Y}})) = \sum_{e \in UNCOV(H, E_{\mathcal{Z}, \mathcal{Y}})} e_{pen}$. The *score* of H , denoted as $score(H, (B, E_{\mathcal{Z}, \mathcal{Y}}))$, is the sum $|H| + pen(H, (B, E_{\mathcal{Z}, \mathcal{Y}}))$. A LAS symbolic learner computes the function $h : \mathcal{Z} \cup B \rightarrow \mathcal{Y}$ by solving the following optimisation problem:

$$H^* = \arg \min_{H \in \mathcal{H}} [score(H, (B, E_{\mathcal{Z}, \mathcal{Y}}))] \quad (1)$$

The above minimisation problem can be equivalently interpreted as jointly maximising the generality of H (i.e., the most compressed ASP program) and its coverage of the examples in $E_{\mathcal{Z}, \mathcal{Y}}$. Given a candidate inductive hypothesis H' , NSIL uses NeurASP to train the neural network. Therefore, the following is a reformulation of the NeurASP learning task [7]. Let Π_θ be the ASP program given by $\Pi_\theta = B \cup H$, where B includes a choice rule over the possible latent concept values in \mathcal{Z} , that the neural network can predict from a collection X of raw data, in a given sample $\langle X, y \rangle$. Choice in ASP means that Π_θ may have multiple answer sets which include relations of the form $g(x_i) = z_i$. We indicate with $AS_{\Pi_\theta}(g(x_i) = z_i)$ the set of answer sets of Π_θ that include the facts $g(x_i) = z_i$. For any $A \in AS_{\Pi_\theta}(g(x_i) = z_i)$, we define the probability of an answer set A as: $P(A) = [\prod_{g(x_i)=z_i \in A} P(g(x_i) = z_i)] / |AS_{\Pi_\theta}(g(x_i) = z_i)|$, where $P(g(x_i) = z_i)$ is given by the neural network. Let $\Pi_\theta(y) = \Pi_\theta \cup \{\leftarrow \text{not } y\}$, where the constraint $\leftarrow \text{not } y$ indicates that the label y must be satisfied in the answer sets of $\Pi_\theta(y)$. Also, let $AS_{\Pi_\theta(y)}(g(x_i) = z_i)$ denote the answer sets of $\Pi_\theta(y)$ that include the facts $g(x_i) = z_i$. The probability of satisfying a particular label y can be computed as: $P(\Pi_\theta(y)) = \sum_{A \in AS_{\Pi_\theta(y)}(g(x_i)=z_i)} P(A)$. This second step of NSIL computes the function $g : \mathcal{X} \rightarrow \mathcal{Z}$ by solving the following optimisation that maximises the log-likelihood of labels y under the ASP semantics of $\Pi_\theta(y)$:

$$\theta^* = \arg \max_{\theta} \sum_{\langle X, y \rangle \in D} \log(P(\Pi_\theta(y))). \quad (2)$$

Hence, NSIL integrates neural and symbolic components by iteratively solving Equations 1 and 2:

1. NSIL is initialised by a bootstrap stage that learns an initial inductive hypothesis H' that satisfies the domain knowledge B and covers each possible label in \mathcal{Y} , independently of specific neural predictions. θ is initialised randomly.
2. The parameters of the neural network θ are updated for 1 epoch on dataset D , using the ASP program $\Pi_\theta(y) = B \cup H' \cup \{\leftarrow \text{not } y\}$, for each $\langle X, y \rangle \in D$.
3. Corrective examples are constructed using neural network predictions (described in detail in Section 3.3), and a new H' is learned with a LAS symbolic learner.
4. Steps 2-3 are repeated for a fixed number of iterations. NSIL outputs the trained neural network g and the learned inductive hypothesis H .

3.3. Symbolic learning optimisation

At each iteration of NSIL, and as shown in Equation 1, the LAS symbolic learner searches for a candidate inductive hypothesis $H' \in \mathcal{H}$ that minimises the score in terms of its size and the total penalty of examples $E_{Z,Y}$ that are *not* covered by H' . The examples, and their associated weight penalties, are therefore crucial to the optimisation, as the symbolic learner is encouraged to learn a H' that covers examples with large weight penalties. The examples are symbolic and define logical relations that are expected to be true or false w.r.t. some contextual information, such that these relations are included or excluded from the answer sets of $B \cup H'$ that cover the example. Formally, an example $e \in E_{Z,Y}$ is a tuple $e = \langle e_{\text{pen}}, e_{\text{inc}}, e_{\text{exc}}, e_{\text{ctx}} \rangle$ where e_{pen} is the weight penalty, and e_{inc} and e_{exc} are respectively, sets of relations to be included and excluded w.r.t. the contextual information e_{ctx} . Examples can be either positive or negative, depending on the task; positive examples are used when the search space \mathcal{H} contains clauses without constraints, whereas positive *and* negative examples are used when \mathcal{H} includes more expressive first-order rules, including constraints. In this paper, the key intuition is that each example associates a possible combination of latent concepts Z with a label y , expressed using e_{ctx} , and $e_{\text{inc}}, e_{\text{exc}}$ respectively. For positive examples, the weight penalty therefore influences whether y should be in the answer sets of $B \cup H' \cup Z$, and the higher the penalty, the higher the cost a candidate hypothesis pays if y is not in the answer sets of $B \cup H' \cup Z$. The opposite holds for negative examples. $E_{Z,Y}$ is constructed initially during a bootstrap stage, and dynamically modified throughout training to optimise towards the final H . The example structure for each of the learning tasks we have used to evaluate NSIL is given in Section 4, and the reader is referred to [30] for further details regarding examples for the LAS systems.

Bootstrap examples Initially, the neural network is not trained and therefore cannot be used to optimise H' . We bootstrap the learning of H' using a set of examples $E_{Z,Y}$ defined to cover all the target labels $y \in \mathcal{Y}$, with respect to possible combinations of latent concept values Z . The symbolic learner ensures the initial H' is the shortest inductive hypothesis that satisfies the label coverage. Let us now define how $E_{Z,Y}$ is dynamically modified throughout training using corrective examples.

Computing corrective examples On each iteration of NSIL training, we use two sources of information to optimise H . Once the neural network has been trained with a candidate H' , we can analyse the overall performance of both neural and symbolic components in predicting training set labels y from X . Also, the neural network returns confidence scores when predicting Z . The challenge is how to create corrective examples from these information sources, that appropriately weight possible combinations of Z and y in the form of examples for the symbolic learner, such that a new H' is learned. To achieve this, we create two types of corrective examples called *explore* and *exploit*, that respectively encourage exploration and exploitation of the hypothesis search space. The explore examples relate possible Z to a different label than what H' currently predicts, whilst the exploit examples reinforce the current prediction of y for a particular Z . When a combination of Z and y is obtained, we create a pair of linked explore and exploit examples, adjusting their weights simultaneously, in order to give a consistent optimisation signal to the symbolic learner. The goal is to maximise (c.f. minimise) the weights of the exploitation (c.f. exploration) examples that contain the correct Z for each label y , such that the correct H is learned. At this stage, two questions remain: (1) How are Z and y obtained?

(2) How are the weights calculated and updated?

To obtain Z and y , for exploration we compute the answer sets of the ASP program $B \cup H' \cup \{\leftarrow \text{not } y\}$, for each $y \in \mathcal{Y}$, where B contains a choice rule for possible values of Z . The answer sets therefore contain Z that lead to label y , given the current H' . For exploitation, for each $\langle X, y \rangle \in D$, we run a forward pass of the neural network on each $x_i \in X$ to obtain Z , and y is obtained directly from the training set. When a combination of Z and y are first obtained, the weights of the corresponding explore and exploit examples are initialised as $e_{\text{pen}} = 1$ (i.e., equal penalty). Let us now define how these example weights are updated throughout NSIL training.

Updating the weights of corrective examples Weight updates are computed at the end of each NSIL iteration. Firstly, the weight of the explore (c.f. exploit) example is increased (c.f. decreased) by the False Negative Rate (FNR) that NSIL has for y , w.r.t. the current H' and θ :

$$FNR_{H',\theta}(y) = 100 \times \left(1 - \frac{TP_{H',\theta}(y)}{TP_{H',\theta}(y) + FN_{H',\theta}(y)} \right) \quad (3)$$

where $TP_{H',\theta}(y)$ and $FN_{H',\theta}(y)$ are respectively the number of true positive and false negative training samples with label y that NSIL predicts on the training set. The FNR is multiplied by 100 to ensure the weight adjustment is informative when converted to an integer, as ASP does not support real numbers. For each explore example related to y , the weight is updated as $e_{\text{pen}} + \lambda FNR_{H',\theta}(y)$, and the weight of the corresponding exploit example is updated as $e_{\text{pen}} - \lambda FNR_{H',\theta}(y)$, where in both cases, e_{pen} is the current example weight and $\lambda \in [0, 1]$ is a learning rate that controls the effect of the proposed update.² The weight is then clipped to a minimum of 1 and maximum of 101 to ensure the example weights are > 0 as required by the LAS symbolic learner. Secondly, the weight of the exploit (c.f. explore) example is increased (c.f. decreased) by the aggregated confidence score of the neural network when predicting Z from X . Let $P_\theta(X) = \prod_{x_i \in X} \max(g(x_i))$, and let $D_{Z,y}$ be the set of $\langle X, y \rangle \in D$ with label y , where the neural network predicts Z . The aggregated confidence score $CONF_\theta$ is defined as:

$$CONF_\theta(Z, y) = \frac{100}{|D_{Z,y}|} \times \sum_{\langle X, y \rangle \in D_{Z,y}} P_\theta(X) \quad (4)$$

The example weight updates are then calculated as $e_{\text{pen}} + \lambda CONF_\theta(Z, y)$ and $e_{\text{pen}} - \lambda CONF_\theta(Z, y)$ for the exploit and explore examples respectively. Again, the weights are clipped to the range $\{1..101\}$. To summarise, Equations 3 and 4 generate exploration and exploitation signals to encourage the symbolic learner to explore the hypothesis space when NSIL performs poorly, and to retain the current H' when NSIL performs well. Let us now present the results of our experiments.

4. Experiments

The experiments answer the following questions: (1) Can NSIL solve tasks involving the joint training of both neural and symbolic components? (2) Given the presence of symbolic learning,

²In this paper, we set $\lambda = 1$ as this yields the best performance. Please see Appendix A.4.1 for more details.

is NSIL more data efficient than pure differentiable models, whilst achieving the same if not better overall accuracy? (3) Given the weak supervision, does NSIL train the neural network effectively, and learn accurate, general and interpretable knowledge? (4) Can NSIL learn symbolic knowledge from raw data that solves computationally harder problems compared to computing polynomial functions? To answer these questions, we evaluate NSIL using two task domains: *MNIST Arithmetic*, and *MNIST Hitting Sets*. Questions 1 and 3 are addressed in both domains, where we evaluate the final accuracy achieved by the neural network, and also present the learned inductive hypotheses. Question 2 is answered in the *MNIST Arithmetic* tasks by using reducing percentages of training data. Finally, question 4 is answered in the *MNIST Hitting Sets* tasks that require more complex and expressive hypotheses to be learned.

Baselines To the best of our knowledge, *Meta_{Abd}* is the only method that jointly trains both a conventional neural network and a symbolic learner. However, *Meta_{Abd}* is limited to learning definite rules. We demonstrate the advantage of NSIL w.r.t. *Meta_{Abd}* in the *MNIST Hitting Sets* tasks, by showing that NSIL learns more general and expressive programs that can be applied to any collection of sets, and can generate all the hitting sets of a given collection. We compare NSIL to the following differentiable models. In both tasks, we use two variants of a Concept Bottleneck Model (CBM) [31]. The perception component is the same as the one used in NSIL, but the reasoning component is differentiable, and implemented as a neural network with 3 linear layers. We use the *joint* CBM and ensure no concept loss is used during training, therefore achieving the same weak supervision as NSIL. The second CBM variant, denoted CBM-S, adds a softmax layer to the CNN, replicating more closely the neural network used in NSIL. As a third baseline, in the *MNIST Arithmetic* tasks, we use the CNN from [5, 7], and in the *MNIST Hitting Sets* tasks, we use a CNN-LSTM. Finally, in the *MNIST Arithmetic* tasks, we attach a Neural Arithmetic Logic Unit (NALU) and a Neural Accumulator (NAC) to a CNN-LSTM [32]. NALU and NAC are not used in the *Hitting Sets* tasks as they are designed for regression and not classification. Full architectural details are presented in Appendix A.4.3.

Experiment setup The neural network in NSIL is the *MNIST CNN* model from [5, 7], and we use the FastLAS [11] symbolic learner in the *MNIST Arithmetic* tasks, and ILASP [28] in the *MNIST Hitting Sets* tasks. To evaluate NSIL, we use the inference architecture shown in Figure 1b. The neural component maps each raw input x_i to its latent concept value z_i , and the symbolic component maps Z to a final prediction, using the learned hypothesis and B . Each experiment is repeated 20 times using 20 randomly generated seeds. The performance is measured by mean classification accuracy per epoch, where 1 epoch = 1 NSIL iteration, for comparison to the baseline models. We perform hyper-parameter tuning for all models using a held-out validation set and a separate random seed.³ All experiments are performed on the following infrastructure: RHEL 8.5 x86_64 with Intel Xeon E5-2667 CPUs (20 cores total), and an NVIDIA A100 GPU, 150GB RAM. The ASP encodings of the domain knowledge are given in Appendix A.4.2.

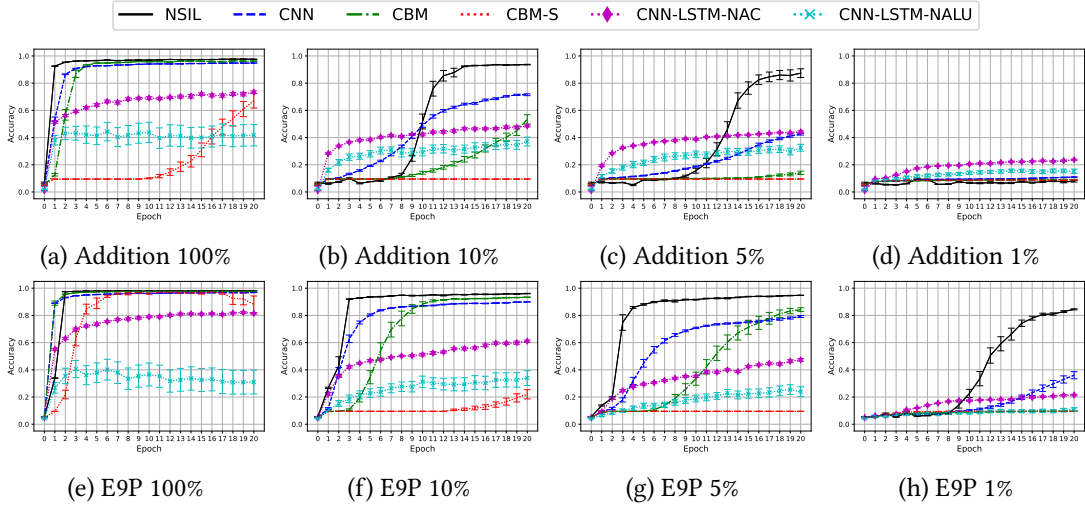


Figure 2: MNIST Arithmetic accuracy with reducing training sets. Error bars indicate standard error.

4.1. MNIST Arithmetic

We consider the MNIST Addition task for single digit numbers used in [5, 7] and a variation called MNIST Even9Plus. They both use a set \mathcal{X} composed of raw MNIST images of single digits (0..9), and the dataset D contains $\langle X, y \rangle$ samples where X is a pair of digit images. The space of target labels $\mathcal{Y} = \{0..18\}$, and the latent concept $\langle n_C, \mathcal{Z}_C \rangle$ has $n_C = \text{digit}$ and $\mathcal{Z}_C = \{0..9\}$. In the MNIST Addition task, the label y indicates the sum of the digit values in X . The goal is to train the neural network to classify the digit images in X , and jointly learn an inductive hypothesis that defines the sum of the two digits. In the MNIST Even9Plus task, X is the same pair of images, but the label y is equal to the digit value of the second image, if the digit value of the first image is even, or 9 plus the digit value in the second image otherwise. Both tasks use training, validation and test datasets of size 24,000, 6,000, and 5,000 samples respectively. The latent concept of a raw image $x_i \in X$ is represented as $\text{digit}(i, z_i)$, where $i \in \{1, 2\}$ and z_i is the associated latent concept value. B uses the relation `result` to express the final label, relations `even`, `not even`, `plus_nine`, and `=`, as well as the function `+` to specify the search space of inductive hypotheses. The positive explore examples have $e_{\text{inc}} = \{\text{label}\}$, $e_{\text{exc}} = \{\}$, and e_{ctx} given by the set of rules $\{Z. \text{label} \leftarrow \text{label}(Y1), Y1 \neq y. \leftarrow \text{label}(Y1), \text{label}(Y2), Y2 \neq Y1.\}$, which states that `label` can be true if a different y is chosen for Z . Positive exploit examples have instead $e_{\text{inc}} = \{y\}$, $e_{\text{exc}} = \{\mathcal{Y} \setminus \{y\}\}$, and $e_{\text{ctx}} = Z$.

Results and analysis The overall accuracy results are shown in Figure 2 for different percentages of the training set D (100% - 1%). NSIL significantly outperforms the differentiable models on both tasks for all dataset percentages, except 1% on the Addition task. NSIL learns the expected hypotheses, shown in Figure 4a in Appendix A.1. The hypotheses define the result as variable `V2`, and capture the correct solutions. For the Addition task, the neural network is

³Please refer to <https://github.com/dancunnington/NSIL> for details of all the hyper-parameters used in our experiments.

trained to classify MNIST images with a mean accuracy of 0.988, 0.968, 0.931, and 0.124 for each dataset percentage respectively. The state-of-the-art for MNIST image classification is 0.991, when the CNN is trained in a fully supervised fashion [33]. We achieve a similar accuracy of 0.988 with only weak-supervision during training. Also, when training with only 1200 samples (5%), NSIL trains the neural network to 0.931 accuracy. For the E9P task, NSIL learns a more expressive inductive hypothesis with negation as failure. The mean accuracy of the neural network is 0.987, 0.972, 0.962, and 0.887 for each dataset percentage respectively. Again, this is a negligible drop in performance compared to the state-of-the-art in fully supervised training [33]. Interestingly, the performance of NSIL trained with just 1% of the data is very poor in the Addition task, although this is not the case in the E9P task, as the neural network achieves 0.887 accuracy with just 240 training samples (1%). Despite the same domain knowledge and hypothesis search space, the label is more informative: There is a reduced number of possible latent concept values for certain labels which gives a stronger back-propagation signal for training the neural network. E.g., for all samples with label $y = 10$, the first digit must be odd and the second digit must be 1, as it's only possible to obtain 10 if the first digit is odd, and $9 +$ the second digit sums to 10. In the addition task, it's possible to obtain 10 with multiple combinations of the first and second digits. Finally, we increased the hypothesis search space in the more challenging E9P task with the additional functions; '-', '×', '÷', as well as `plus_eight`, ..., `plus_one` relations, and NSIL converges to the expected hypothesis in all of these extended search spaces (see Appendix A.2).

4.2. MNIST Hitting Sets

We consider two variations of the hitting sets problem [17]. The first, denoted HS, is defined in Example 1, the second, denoted CHS, adds the constraint that no hitting sets occur if $|X| \geq 3$, and digit 1 exists in any subset in X . Each $\langle X, y \rangle$ sample contains 4 MNIST or FashionMNIST images from classes 1-5, arranged into subsets in X . The label y is 1 or 0 indicating the existence of a hitting set, and the goal is to train the neural network to classify MNIST images, whilst learning the HS or CHS rules to decide the existence of and generate hitting sets. Training, validation and test datasets contain 1502, 376, and 316 examples respectively. The latent concept has $n_C = \text{MNIST_class}$ and $Z_C = \{1..5\}$, and for a raw image x_i , it is represented as `ss_element(s_i, z_i)`, where $i \in \{1..4\}$, s_i is a subset identifier, and z_i is the associated latent concept value. We consider well-formed Z , with no duplicate elements in a set, and elements within the same set are arranged in ascending order. The domain knowledge B contains $k = 2$, the '=' relations, the relation `hs` stating an element is in a hitting set, and `ss_element(3, V)` and `ss_element(V, 1)` relations that define respectively subset 3 (and any element V) and digit 1 in any subset V. Both explore and exploit corrective examples share a common $e_{\text{inc}} = \{\}$, $e_{\text{exc}} = \{\}$, and $e_{\text{ctx}} = Z$. Exploration and exploitation is achieved by representing the example positively or negatively w.r.t. the given label. In these tasks, we use a CNN-LSTM baseline, where a CNN firstly encodes a feature vector for each image in a collection X , and an LSTM followed by a linear and a sigmoid layer returns a binary classification as to the existence of a hitting set. To encode the structure of a collection, we create an image of '{' and '}' characters and pass in the entire collection as input, e.g., $X = \{\{, 2, 7, \}, \{, 1, \}, \{, 3, \}$.

Results and analysis Figure 3 shows the overall accuracy for both hitting sets tasks using

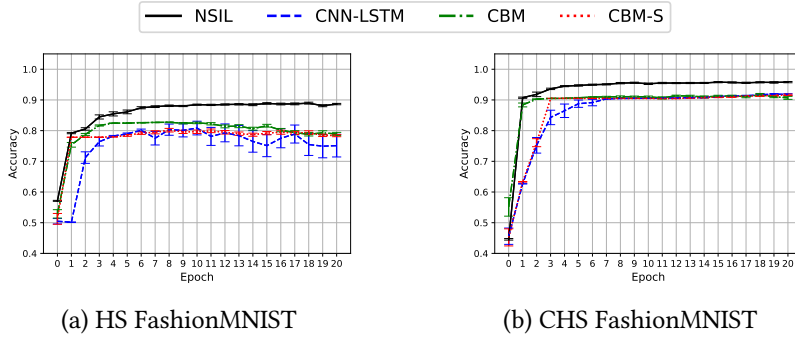


Figure 3: Hitting Sets accuracy. Error bars indicate standard error.

only FashionMNIST images, as this dataset is more challenging than MNIST. NSIL outperforms the baselines on both tasks and learns the correct hypotheses (see Figure 4b in Appendix A.1). The MNIST training curves are shown in Appendix A.3. In Figure 4b, the first rule in the HS and CHS inductive hypotheses is a choice rule that generates hitting sets containing element V_2 at index V_1 . The second rule defines whether a subset in the collection is hit (i.e., contains an element in the hitting set). The constraints ensure each subset in the collection is hit, and elements in the hitting set are different. In the CHS task, the final constraint ensures collections with 3 subsets and element 1 contain no hitting sets. When classifying images, the neural network achieves a mean accuracy of 0.992 (MNIST) and 0.906 (FashionMNIST) on the HS task, and 0.992 (MNIST) and 0.901 (FashionMNIST) on the CHS task. The neural network achieves near state-of-the-art performance on MNIST images. For FashionMNIST, the state-of-the-art is 0.969 when a CNN is trained in a fully supervised fashion [34], and NSIL achieves 0.906. Despite the hitting sets tasks having more complex hypotheses, the neural networks are trained to a similar accuracy as in the MNIST Arithmetic tasks. All models perform better on the CHS task. This is due to a shortcut; a negative test example with 3 subsets can be correctly predicted if the neural network predicts class 1 for *any* image within the subset, as opposed to having to rely on accurate predictions for more images when the constraint is not present in the HS task. Also, as a result of the constraint, 40 test examples that contain 3 subsets become negative in the CHS task, and the models can take advantage of the shortcut.

The inductive hypotheses learned by NSIL have multiple answer sets due to the choice rule $0\{\text{hs}(V_1, V_2)\}1$. Therefore, despite training labels indicating the existence of a hitting set, the programs can generate all the hitting sets. We verified this by computing the hamming loss against the ground truth hitting sets for all test samples. The hamming losses are 0.003 and 0.031 on the HS task, and 0.003 and 0.025 on the CHS task, for the MNIST and FashionMNIST images respectively. This indicates almost perfect hitting set generation, with the minor errors due to mis-classifications by the neural network, as opposed to errors in the learned hypotheses. The larger hamming loss for FashionMNIST images reflects the weaker neural network performance with this dataset. Finally, in the more challenging HS FashionMNIST task we sampled 10 sets of additional predicates containing a combination of elements 1..5 and subset IDs 1..4. NSIL converged to the expected hypothesis in all of these extended search spaces (see Appendix A.2).

5. Conclusion

This paper introduces a neuro-symbolic learner (NSIL) that learns complex knowledge whilst training a neural network to classify latent concepts from raw data. The framework uses a symbolic learner to learn an inductive hypothesis capable of solving computationally hard problems. The novel aspect of the architecture is the use of corrective examples to bias the symbolic learner. The results show that NSIL outperforms differentiable models in terms of accuracy and data efficiency, and solves the hitting sets problem that other approaches can't solve. In terms of limitations, the differentiable baselines learn faster than NSIL, although given more time, NSIL achieves a greater accuracy (see the Appendix for learning time vs. accuracy plots). Future work includes generalising NSIL to support multiple latent concepts by integrating multiple neural networks, and extending NSIL to take advantage of unsupervised learning techniques.

References

- [1] A. d'Avila Garcez, M. Gori, L. C. Lamb, L. Serafini, M. Spranger, S. N. Tran, Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning, *FLAP* 6 (2019) 611–632. URL: <https://collegepublications.co.uk/ifcolog/?00033>.
- [2] A. d'Avila Garcez, L. C. Lamb, Neurosymbolic AI: the 3rd wave, *CoRR* abs/2012.05876 (2020). URL: <https://arxiv.org/abs/2012.05876>. arXiv:2012.05876.
- [3] L. De Raedt, S. Dumančić, R. Manhaeve, G. Marra, From statistical relational to neuro-symbolic artificial intelligence, in: C. Bessiere (Ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20, International Joint Conferences on Artificial Intelligence Organization, 2020*, pp. 4943–4950. URL: <https://doi.org/10.24963/ijcai.2020/688>. doi:10.24963/ijcai.2020/688, survey track.
- [4] S. Badreddine, A. d'Avila Garcez, L. Serafini, M. Spranger, Logic tensor networks, *Artificial Intelligence* 303 (2022) 103649. URL: <https://www.sciencedirect.com/science/article/pii/S0004370221002009>. doi:<https://doi.org/10.1016/j.artint.2021.103649>.
- [5] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, L. De Raedt, Deepproblog: Neural probabilistic logic programming, in: *Advances in Neural Information Processing Systems*, 2018, pp. 3749–3759.
- [6] R. Riegel, A. G. Gray, F. P. S. Luus, N. Khan, N. Makondo, I. Y. Akhalwaya, H. Qian, R. Fagin, F. Barahona, U. Sharma, S. Ikbal, H. Karanam, S. Neelam, A. Likhyan, S. K. Srivastava, Logical neural networks, *CoRR* abs/2006.13155 (2020). URL: <https://arxiv.org/abs/2006.13155>. arXiv:2006.13155.
- [7] Z. Yang, A. Ishay, J. Lee, Neurasp: Embracing neural networks into answer set programming, in: C. Bessiere (Ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20, International Joint Conferences on Artificial Intelligence Organization, 2020*, pp. 1755–1762. URL: <https://doi.org/10.24963/ijcai.2020/243>. doi:10.24963/ijcai.2020/243.
- [8] W.-Z. Dai, S. Muggleton, Abductive knowledge induction from raw data, in: Z.-H. Zhou (Ed.), *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*,

- IJCAI-21, International Joint Conferences on Artificial Intelligence Organization, 2021, pp. 1845–1851. URL: <https://doi.org/10.24963/ijcai.2021/254>. doi:10.24963/ijcai.2021/254, main Track.
- [9] A. Cropper, S. Muggleton, Metagol system, <https://github.com/metagol/metagol>, 2016. URL: <https://github.com/metagol/metagol>.
- [10] E. Dantsin, T. Eiter, G. Gottlob, A. Voronkov, Complexity and expressive power of logic programming, *ACM Computing Surveys (CSUR)* 33 (2001) 374–425.
- [11] M. Law, A. Russo, E. Bertino, K. Broda, J. Lobo, Fastlas: scalable inductive logic programming incorporating domain-specific optimisation criteria, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2020, pp. 2877–2885.
- [12] M. Law, A. Russo, K. Broda, Inductive learning of answer set programs from noisy examples, *Advances in Cognitive Systems* (2018) 57–76.
- [13] M. Gelfond, Y. Kahl, Knowledge representation, reasoning, and the design of intelligent agents: The answer-set programming approach, Cambridge University Press, Cambridge, UK, 2014.
- [14] R. Evans, M. Bošnjak, L. Buesing, K. Ellis, D. Pfau, P. Kohli, M. Sergot, Making sense of raw input, *Artificial Intelligence* 299 (2021) 103521. URL: <https://www.sciencedirect.com/science/article/pii/S0004370221000722>. doi:<https://doi.org/10.1016/j.artint.2021.103521>.
- [15] Y. LeCun, C. Cortes, MNIST handwritten digit database, <http://yann.lecun.com/exdb/mnist/>, 2010. URL: <http://yann.lecun.com/exdb/mnist/>.
- [16] H. Xiao, K. Rasul, R. Vollgraf, Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, *CoRR abs/1708.07747* (2017). URL: <http://arxiv.org/abs/1708.07747>. arXiv:1708.07747.
- [17] R. M. Karp, Reducibility among combinatorial problems, in: *Complexity of computer computations*, Springer, 1972, pp. 85–103.
- [18] T. R. Besold, A. d. Garcez, S. Bader, H. Bowman, L. C. Lamb, L. de Penning, B. Illumino, H. Poon, C. Gerson Zaverucha, Neural-symbolic learning and reasoning: A survey and interpretation, *Neuro-Symbolic Artificial Intelligence: The State of the Art* 342 (2022) 1.
- [19] D. Corapi, A. Russo, E. Lupu, Inductive logic programming as abductive search, in: *Technical communications of the 26th international conference on logic programming*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2010, pp. 54–63.
- [20] S. Muggleton, L. De Raedt, Inductive logic programming: Theory and methods, *The Journal of Logic Programming* 19 (1994) 629–679.
- [21] S. H. Muggleton, D. Lin, A. Tamaddoni-Nezhad, Meta-interpretive learning of higher-order dyadic datalog: Predicate invention revisited, *Machine Learning* 100 (2015) 49–73.
- [22] A. Payani, F. Fekri, Inductive logic programming via differentiable deep neural logic networks, *CoRR abs/1906.03523* (2019). URL: <http://arxiv.org/abs/1906.03523>. arXiv:1906.03523.
- [23] P. Sen, B. W. S. R. de Carvalho, R. Riegel, A. Gray, Neuro-symbolic inductive logic programming with logical neural networks, 2021. arXiv:2112.03324.
- [24] H. Shindo, M. Nishino, A. Yamamoto, Differentiable inductive logic programming for structured examples, *Proceedings of the AAAI Conference on Artificial Intelligence* 35 (2021) 5034–5041. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/16637>.

- [25] D. Cunnington, M. Law, A. Russo, J. Lobo, FF-NSL: feed-forward neural-symbolic learner, CoRR abs/2106.13103 (2021). URL: <https://arxiv.org/abs/2106.13103>. arXiv: 2106.13103.
- [26] R. Evans, E. Grefenstette, Learning explanatory rules from noisy data, *Journal of Artificial Intelligence Research* 61 (2018) 1–64.
- [27] W.-Z. Dai, Q. Xu, Y. Yu, Z.-H. Zhou, Bridging machine learning and logical reasoning by abductive learning, *Advances in Neural Information Processing Systems* 32 (2019).
- [28] M. Law, Inductive learning of answer set programs, Ph.D. thesis, Imperial College London, 2018.
- [29] M. Law, A. Russo, K. Broda, The complexity and generality of learning answer set programs, *Artif. Intell.* 259 (2018) 110–146.
- [30] M. Law, A. Russo, K. Broda, Logic-based learning of answer set programs, in: *Reasoning Web. Explainable Artificial Intelligence - 15th International Summer School 2019, Bolzano, Italy, September 20-24, 2019, Tutorial Lectures, 2019*, pp. 196–231.
- [31] P. W. Koh, T. Nguyen, Y. S. Tang, S. Mussmann, E. Pierson, B. Kim, P. Liang, Concept bottleneck models, in: *International Conference on Machine Learning, 2020*, pp. 5338–5348.
- [32] A. Trask, F. Hill, S. E. Reed, J. Rae, C. Dyer, P. Blunsom, Neural arithmetic logic units, *Advances in neural information processing systems* 31 (2018).
- [33] S. An, M. Lee, S. Park, H. Yang, J. So, An ensemble of simple convolutional neural network models for mnist digit recognition, arXiv preprint arXiv:2008.10400 (2020).
- [34] M. S. Tanveer, M. U. K. Khan, C.-M. Kyung, Fine-tuning darts for image classification, in: *2020 25th International Conference on Pattern Recognition (ICPR), IEEE, 2021*, pp. 4789–4796.
- [35] S. Ruder, An overview of gradient descent optimization algorithms, arXiv preprint arXiv:1609.04747 (2016).

A. Appendix

A.1. NSIL learned hypotheses

```
MNIST Addition  
result(V0,V1,V2) :- V2 = V0 + V1.  
MNIST E9P  
result(V0,V1,V2) :- even(V0), V2 = V1.  
result(V0,V1,V2) :- not even(V0),  
    plus_nine(V1,V2).
```

(a) MNIST Arithmetic

```
HS  
0 {hs(V1,V2) } 1.  
hit(V1) :- hs(V3,V2), ss_element(V1,V2).  
:- ss_element(V1,V2), not hit(V1).  
:- hs(V3,V1), hs(V3,V2), V1 != V2.  
CHS  
0 {hs(V1,V2) } 1.  
hit(V1) :- hs(V3,V2), ss_element(V1,V2).  
:- ss_element(V1,V2), not hit(V1).  
:- hs(V3,V1), hs(V3,V2), V1 != V2.  
:- ss_element(3,V2), ss_element(V1,1).
```

(b) MNIST Hitting Sets

Figure 4: NSIL learned hypotheses

A.2. Increasing the hypothesis space

To demonstrate NSIL can converge to the correct hypothesis in a variety of hypothesis spaces, we select the more challenging E9P and HS FashionMNIST tasks, and increase the hypothesis space using additional predicates, functions and relations in B . We record the time and number of iterations NSIL requires until convergence, and the results demonstrate NSIL converges to the correct hypothesis in all of these cases, but may require more iterations and/or time to do so.

E9P We add arithmetic functions ‘-’, ‘×’, ‘÷’, as well as `plus_eight`, ..., `plus_one` relations, and increase the grounding of possible labels, such that the results of multiplication and subtraction support all combinations of digits 0..9. Given the strong performance by NSIL with reducing dataset percentages, we use 40% of the data to reduce the neural network training time (this does not affect the symbolic learning time, nor simplify the symbolic learning task). The results are shown in Table 1a.

HS FashionMNIST Firstly, we define a set of domain knowledge called *Standard*, that contains only the predicates required to learn the expected rules of the HS task. The standard domain knowledge includes $k = 2$, the ‘!=’ relations, and the relation `hs` stating an element is in a hitting set. We extend these by sampling 10 sets of additional predicates containing a combination of elements 1..5 (denoted ‘el’) and subset IDs 1..4 (denoted ‘ssID’). The results are shown in Table 1b.

A.3. Hitting Sets test accuracy: MNIST images

The results are presented in Figure 5.

Table 1

Increasing the hypothesis space for NSIL. The asterisks indicate the predicates and relations used in the results presented in Section 4. NSIL converges to the correct hypothesis in all of these cases, but sometimes requires more iterations and/or time to do so.

(a) E9P				(b) HS FashionMNIST		
Domain knowledge	Label range	Convergence iteration	Convergence time (s)	Domain knowledge	Convergence iteration	Convergence time (s)
* +, =, even, not even, plus_nine	0..18	2	237.88	Standard	1	73.12
+, -, ×, ÷, =, even, not even, plus_nine	0..18	2	258.31	Standard, sslD 4, el 3	1	87.88
"	-9..81	2	763.06	Standard, sslD 4, el 4	1	87.90
=, even, not even, plus_nine, ..., plus_one	0..18	4	2034.62	Standard, sslD 2, sslD 4	1	89.88
"	-9..81	3	23614.5	Standard, el 3, el 4	1	100.31
				Standard, sslD 2, el 2	1	101.97
				Standard, sslD 1, el 4	1	106.83
				Standard, sslD 1, sslD 3	2	867.00
				Standard, sslD 3, el 4	3	1014.50
				Standard, sslD 3, el 3	3	1242.44
				* Standard, sslD 3, el 1	3	1736.42

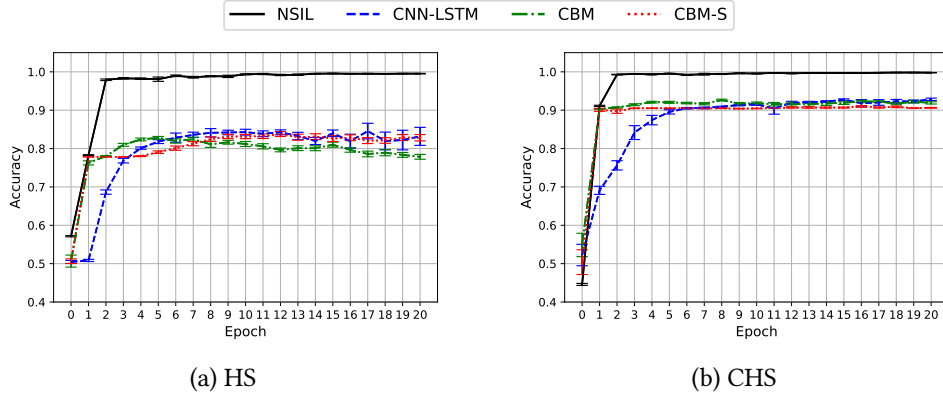


Figure 5: Hitting Sets accuracy when MNIST images are used to represent elements. Error bars indicate standard error. NSIL significantly outperforms the baseline differentiable models on the HS task, and all models perform better on the CHS task, due to the shortcut explained in Section 4.2.

A.4. Experiment setup

A.4.1. Relying on neural network confidence for corrective example weights

In this section, we investigate the λ hyper-parameter for NSIL introduced in Section 3.3. When NSIL calculates the weights of the corrective examples for the symbolic learner, two updates occur; *exploration* based on the FNR, and *exploitation*, based on the neural network confidence (see Section 3.3). λ can be adjusted to vary the effect of each update. During exploration, we set $\lambda = 1$ always, to give priority to the FNR, as the FNR is reliable and takes into account both neural and symbolic components. However, during exploitation, we may not want to fully rely on the neural network confidence when the neural network is under-trained. Therefore, we run

experiments with varying $\lambda \in \{1, 0.8, 0.6, 0.4, 0.2, 0\}$ during exploitation only. We again select the more challenging E9P and HS tasks, and use 40% of the training data in the E9P task. We run NSIL for 10 iterations, with 5 repeats, using different randomly generated seeds. Figure 6 presents the results, where the error bars indicate standard error.




Figure 6: Accuracy with varying λ , E9P and HS tasks.

In each task, $\lambda = 1$ and $\lambda = 0$ yields the best and worst performance respectively, which indicates the neural network confidence can be relied upon for these tasks. This is our justification for using $\lambda = 1$ in Section 3.3. Analysing the $\lambda = 0$ cases, the accuracy on the E9P task improves throughout training and the symbolic learner converges to more accurate hypotheses. However, in the HS tasks, after a few iterations, the symbolic learner learns the same hypothesis and performance plateaus. When corrective examples are generated, there are many examples with the same weight, because there are many choices of digits for each combination of collection structure and label, given the labels are binary. Therefore, the optimisation signal to the symbolic learner is weaker in the HS tasks compared to the E9P task when $\lambda = 0$, because in the E9P task, labels 0..18 generate corrective examples with many different weights. When $\lambda = 1$ in the HS tasks, the neural network confidence scores help the symbolic learner to converge the correct hypothesis by providing a strong exploitation signal for specific examples.

A.4.2. Domain knowledge

Domain knowledge used for all tasks is presented in Figure 7.

A.4.3. Model details

NSIL CNN The neural component in NSIL is the CNN architecture from [5, 7]. It consists of an encoder with 2 convolutional layers with kernel size 5, and output sizes 6 and 16 respectively. Each convolutional layer is followed by a max pooling layer of size 2, stride 2. The encoder is followed by 3 linear layers of size 120, 84, and 10 (in the hitting sets tasks, the last layer is of size 5), and all layers are followed by a ReLU activation function. Finally, a softmax layer returns the output probability distribution.

Baselines The baseline CNN in the MNIST Arithmetic tasks follows the baseline CNN architecture from [5, 7], which is largely the same as the NSIL CNN, except the size of the last

```

r(0..18).
d(0..9).
even(X) :- d(X), X \ 2 = 0.
plus_nine(X1,X2) :- d(X1), X2=9+X1.
res(X1,X2,Y) :- dig(1,X1), dig(2,X2).
:- dig(1,X1),dig(2,X2),res(X1,X2,Y1),res(X1
,X2,Y2), Y1 != Y2.

#modeh(res(var(d),var(d),var(r))).
#modeb(var(n) = var(d)).
#modeb(var(n) = var(d) + var(d)).
#modeb(plus_nine(var(d),var(r))).
#modeb(even(var(d))).
#modeb(not even(var(d))).

```

(a) MNIST Arithmetic tasks

```

s(1..4).
h(1..2).
e(1..5).
#modeha(hs(var(h), var(e))).
#modeh(hit(var(s))).
#modeb(hs(var(h), var(e)),(positive)).
#modeb(var(e) != var(e)).
#modeb(ss_element(var(s),var(e)),(positive)).
#modeb(ss_element(3,var(e)),(positive)).
#modeb(ss_element(var(s),1),(positive)).
#modeb(hit(var(s))).

```

(b) MNIST Hitting Sets tasks

Figure 7: ASP encodings of the domain knowledge used within NSIL.

linear layer is 19, the network accepts a pair of concatenated images as input, and a log softmax layer is used to provide a classification. The CNN-LSTM in the MNIST Hitting Sets tasks uses the same CNN encoder as NSIL, applied to each image in the input sequence, followed by an LSTM layer of tunable size, a linear layer of size 1, and a sigmoid layer. The CBM and CBM-S baselines have similar architectures, consisting of the same CNN as NSIL, applied to each image separately. The only difference is that the CBM variant doesn't contain a softmax layer for the CNN. In both variants, the CNN is followed by 3 linear layers where the size of the first two layers are tunable and the last layer is of size 19 in the MNIST arithmetic tasks, and 1 in the MNIST Hitting Sets tasks. In the MNIST Addition task, ReLU activation is used after the first two linear layers for the CBM and CBM-S baselines, as this resulted in improved performance. In the other tasks, no non-linearity was used. The CNN-LSTM-NALU and CNN-LSTM-NAC baselines contain the same CNN as NSIL and the other baselines, followed by an LSTM layer of tunable size, and then either an NALU or NAC layer. Finally, we use the stochastic gradient descent optimiser for all neural networks [35], and tune the learning rate and momentum.³

A.5. Learning time comparison

Figures 8 and 9 show the learning time vs. accuracy comparison between NSIL and the differentiable models for the MNIST Arithmetic and MNIST Hitting Sets tasks respectively. Each point shows the accuracy after an epoch of training (1 epoch = 1 NSIL iteration), and error bars indicate standard error. Also, the x -axis is adjusted to show a meaningful comparison instead of displaying the total running time for 20 epochs. As observed in both task domains, NSIL requires more time but achieves a greater accuracy than the differentiable models.

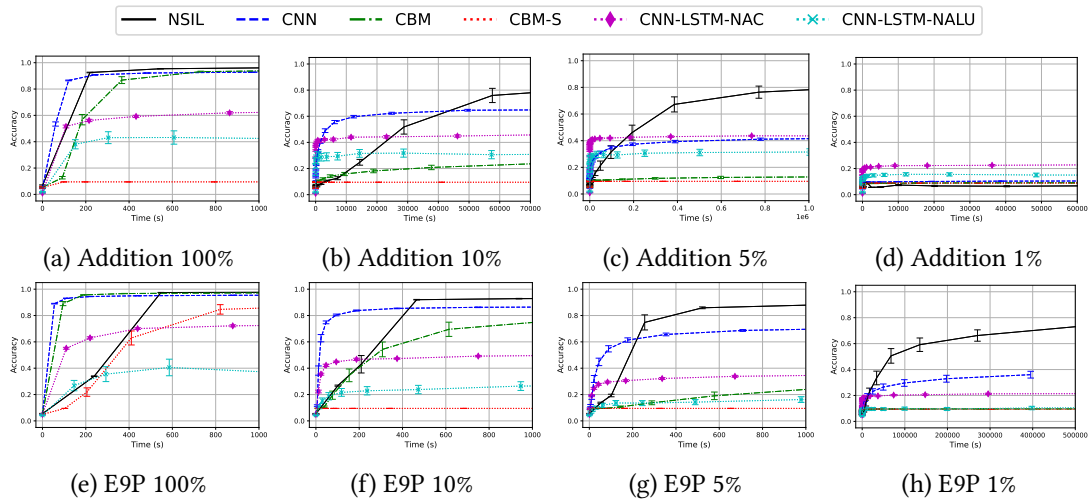


Figure 8: MNIST Arithmetic learning time vs. accuracy with reducing training set percentages.

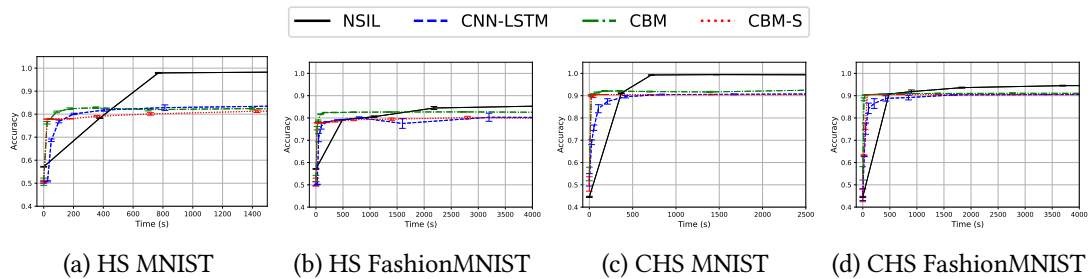


Figure 9: Hitting Sets learning time vs. accuracy.

A.6. Asset licenses

The ILASP system is free to use for research,⁴ FastLAS⁵ and the FashionMNIST dataset⁶ are both open-source with an MIT license, the MNIST dataset is licensed with Creative Commons Attribution-Share Alike 3.0,⁷ and the CNN models used from DeepProbLog are open-source and licensed with Apache 2.0.⁸

A.7. Broader impact

There are no direct negative impacts that we can envisage for our work, given we introduce a general machine learning approach. However, NSIL inherits general concerns regarding the deployment of machine learning systems, and appropriate precautions should be taken, such

⁴<https://ilasp.com/terms>

⁵<https://github.com/spike-imperial/FastLAS/blob/master/LICENSE>

⁶<https://github.com/zalandoresearch/fashion-mnist/blob/master/LICENSE>

⁷See bottom paragraph: <https://keras.io/api/datasets/mnist/>

⁸<https://github.com/ML-KULeuven/deepproblog/blob/master/LICENSE>

as ensuring training data is unbiased, and model inputs/outputs are monitored for adversarial attacks. As NSIL learns human interpretable hypotheses using symbolic learning, this may help to mitigate these issues in some applications, by revealing potential bias, and providing a level of assurance regarding possible downstream predictions based on the learned hypothesis. The usual performance monitoring will still be required if NSIL is deployed into production, to prevent adversarial attacks, and to detect when re-training may be required if the input data is subject to distributional shifts.