

# Neuro-Symbolic Learning of Answer Set Programs from Raw Data

Daniel Cunningham<sup>1,2</sup>, Mark Law<sup>3</sup>, Jorge Lobo<sup>4</sup> and Alessandra Russo<sup>2</sup>

<sup>1</sup>IBM Research Europe

<sup>2</sup>Imperial College London

<sup>3</sup>ILASP Limited

<sup>4</sup>ICREA-Universitat Pompeu Fabra

dancunnington@uk.ibm.com, mark@ilasp.com, jorge.lobo@upf.edu, a.russo@imperial.ac.uk

## Abstract

One of the ultimate goals of Artificial Intelligence is to assist humans in complex decision making. A promising direction for achieving this goal is Neuro-Symbolic AI, which aims to combine the interpretability of symbolic techniques with the ability of deep learning to learn from raw data. However, most current approaches require manually engineered symbolic knowledge, and where end-to-end training is considered, such approaches are either restricted to learning definite programs, or are restricted to training binary neural networks. In this paper, we introduce Neuro-Symbolic Inductive Learner (NSIL), an approach that trains a general neural network to extract latent concepts from raw data, whilst learning symbolic knowledge that maps latent concepts to target labels. The novelty of our approach is a method for biasing the learning of symbolic knowledge, based on the in-training performance of both neural and symbolic components. We evaluate NSIL on three problem domains of different complexity, including an NP-complete problem. Our results demonstrate that NSIL learns expressive knowledge, solves computationally complex problems, and achieves state-of-the-art performance in terms of accuracy and data efficiency. Code and technical appendix: <https://github.com/DanCunnington/NSIL>

## 1 Introduction

Within Artificial Intelligence (AI), one of the ultimate goals is to assist humans in complex decision making, a key challenge in multiple industries such as healthcare, automated maintenance, and security [Lyn Paul *et al.*, 2019; Sittón *et al.*, 2019; Han *et al.*, 2021]. Neuro-Symbolic AI aims to address this challenge by combining the best features of both deep learning and symbolic reasoning techniques [d’Avila Garcez *et al.*, 2019; De Raedt *et al.*, 2020]. For example, many existing approaches improve the training of a neural network using a given symbolic knowledge base [Manhaeve *et al.*, 2018; Dai *et al.*, 2019; Riegel *et al.*, 2020; Yang *et al.*, 2020; Badreddine *et al.*, 2022], and systems that learn symbolic knowledge can utilise pre-trained neural networks to handle raw input data [Evans and Grefenstette, 2018; Cunning-

ton *et al.*, 2023]. However, the assumption that either the neural or symbolic component is given is not practical in many real-world situations, as there may be no pre-trained neural network(s) available (e.g., domain experts are time constrained and can’t label vast quantities of raw data), or the symbolic knowledge may be unknown, and therefore needs to be learned. Recent work aims to lift this assumption by training both neural and symbolic components in an end-to-end fashion [Dai and Muggleton, 2021; Evans *et al.*, 2021]. However, [Dai and Muggleton, 2021] lacks the expressivity required to efficiently represent solutions to many common-sense learning and reasoning tasks [Dantsin *et al.*, 2001], and [Evans *et al.*, 2021] is restricted to training Binary Neural Networks (BNNs) only.

In this paper, we introduce *Neuro-Symbolic Inductive Learner (NSIL)*, which trains a *general* neural network to classify latent concepts from raw data, whilst learning an expressive and interpretable knowledge base that solves computationally *complex* problems. NSIL only requires a target label for each training data point, and no latent concept labels are given. NSIL uses a state-of-the-art symbolic learner to learn a logic program in the language of Answer Set Programming (ASP) [Gelfond and Kahl, 2014], a highly expressive form of knowledge representation which can efficiently represent solutions to problems of computational complexity greater than P [Dantsin *et al.*, 2001]. A general neural network architecture is trained by *reasoning* over the learned knowledge, where the latent concept values that the network should output in order to obtain each target label are identified. We use NeurASP to perform such computation [Yang *et al.*, 2020], and seamlessly integrate symbolic learning by creating *corrective examples* for the symbolic learner. A specific mechanism for weighting these examples is proposed, which biases the symbolic learner to either refine or retain the learned knowledge, balancing *exploration* with *exploitation* of the symbolic search space.

We evaluate NSIL on three problem domains: (1) *Cumulative Arithmetic*, where the learned programs perform the cumulative addition or product over sequences of MNIST digit images. (2) *Two-Digit Arithmetic*, where the learned programs include arithmetic and logical operations over two MNIST digit images. (3) the *Hitting Set* problem, which includes the standard decision problem and a variant that remains NP-complete [Karp, 1972]. In this domain, the learned programs

not only solve the decision problem, but can also generate *all* the hitting sets of a given collection. In all domains, the input is a set of training data points, each containing a sequence of raw images together with a target label. Also, the full space of possible latent concepts is assumed to be known, together with some structural properties of the search space for the symbolic learner. The output is a trained neural network capable of classifying raw images into latent concepts, and an ASP program that maps latent concepts into target labels.

Finally, we compare NSILs performance with purely neural baselines, and where appropriate, *Meta\_Abd* [Dai and Muggleton, 2021], the only existing system that trains a general neural network and a symbolic learner. The results show that NSIL: (1) Outperforms the baselines in terms of overall accuracy and data efficiency. (2) Learns expressive knowledge to represent solutions to an NP-complete problem, which would be difficult to express with any of the other systems. (3) Requires significantly fewer data points than the neural baselines. (4) Trains the neural network to predict latent concepts with an accuracy comparable to fully supervised training, and (5) achieves comparable performance to approaches that either assume the latent concept labels are given, or the symbolic component is given and only the neural network needs to be trained.

## 2 Related Work

Many Neuro-Symbolic integrations have been proposed in the literature to enhance data efficiency, transferability and interpretability of neural networks using symbolic techniques [Besold *et al.*, 2022]. Some approaches inject symbolic knowledge directly into neural architectures [Riegel *et al.*, 2020; Badreddine *et al.*, 2022], whereas others preserve a clear distinction between neural and symbolic components [Manhaeve *et al.*, 2018; Yang *et al.*, 2020; Aspis *et al.*, 2022; Dasaratha *et al.*, 2023]. The main drawback of these approaches is that they require a complete and manually engineered symbolic knowledge base. On the other hand, symbolic learning systems [Muggleton and De Raedt, 1994; Corapi *et al.*, 2010; Muggleton *et al.*, 2015; Law, 2018; Law *et al.*, 2020] are capable of learning interpretable knowledge in a data efficient manner, but are restricted to learning from structured data, even when differentiable training methods are used [Payani and Fekri, 2019; Shindo *et al.*, 2021; Sen *et al.*, 2022]. Therefore, pre-trained neural networks are often required when learning from raw data [Evans and Grefenstette, 2018; Ferreira *et al.*, 2022; Cunningham *et al.*, 2023]. In contrast to all of the aforementioned approaches, our method performs joint Neuro-Symbolic learning, by training a neural network and learning symbolic knowledge simultaneously, which significantly reduces the amount of engineering and labelling required.

NSIL is therefore closely related to the existing approaches that also train both neural and symbolic components [Dai and Muggleton, 2021; Evans *et al.*, 2021; Daniele *et al.*, 2022]. *Meta\_Abd* [Dai and Muggleton, 2021] extends [Dai *et al.*, 2019] by using abduction and induction to jointly train a neural network and learn symbolic knowledge from raw data. Specifically, the symbolic knowledge learned by *Meta\_Abd*

is used to abduce possible latent concept values for training the neural network. These values are pruned further and a unique value is selected for each image within an input sequence, using the neural network confidence. As the authors note, the network is therefore vulnerable to becoming stuck in a local optima if the abductive space is very dense. During the early stages of training, the network predicts with a more equal distribution, which reduces the likelihood of selecting the correct latent concept values. In contrast, NSIL relies on a semantic loss function [Yang *et al.*, 2020], and instead trains the network with a set of possible values. This avoids incorrectly pruning the abductive space and the network is able to escape local optima.

Also, the symbolic learner used by *Meta\_Abd* is restricted in the expressivity of knowledge that can be learned [Law, 2018]. As NSIL learns ASP programs, we can represent common-sense knowledge involving defaults, exceptions, constraints and choice [Gelfond and Kahl, 2014], all of which are outside the scope of *Meta\_Abd*. Using ASP also enables our learned programs to have multiple answer sets (multiple models), whereas *Meta\_Abd* is restricted to learning definite programs that accept only one model. These features enable NSIL to learn knowledge that efficiently represents solutions to complex, NP-complete problems, which would be difficult to represent using *Meta\_Abd*.

The approach in [Evans *et al.*, 2021] encodes the neural network into an ASP program, which enables joint neural network training and symbolic learning within ASP. However, ASP does not support continuous arithmetic, which restricts [Evans *et al.*, 2021] to using BNNs. BNNs are an active area of research, but are currently limited to relatively small network architectures. Due to our modular approach, NSIL does not place such restrictions on the neural network architecture, and can take advantage of GPU-accelerated hardware and software. Finally, [Daniele *et al.*, 2022] learns a perception and reasoning function, but the reasoning function is a look-up table that simply maps symbolic inputs to outputs. In contrast, NSIL is able to learn complex reasoning functions in the form of ASP programs, which are more expressive than a look-up table, and can generalise beyond the specific task used during training.

## 3 Neuro-Symbolic Inductive Learner

### 3.1 Problem Setting

We decompose the task of learning complex knowledge from raw data into a perception function  $g : \mathcal{X} \rightarrow \mathcal{Z}$  and a reasoning function  $h : \mathcal{Z}^{a,b} \rightarrow \mathcal{Y}$ .  $\mathcal{X}$  is the space of raw inputs (e.g., images),  $\mathcal{Z}$  is the space of latent concept values, and  $\mathcal{Y}$  is the space of target labels.  $h$  is defined over sequences of latent concept values, with (possibly) varying length. Therefore,  $\mathcal{Z}^{a,b}$  is the set of all possible sequences of  $\mathcal{Z}$  with length  $\geq a$  and  $\leq b$ . A latent concept is a tuple  $C = \langle n, \mathcal{Z} \rangle$ , where  $n$  is the name for the latent concept. In this paper, we assume w.l.o.g. there is a single latent concept for a given task. Supporting multiple latent concepts could be achieved by partitioning  $\mathcal{Z}$  into different classes, and assigning each class a different name with a different perception function. During training, we observe a dataset of samples  $D = \{\langle X, y \rangle, \dots\}$ , where

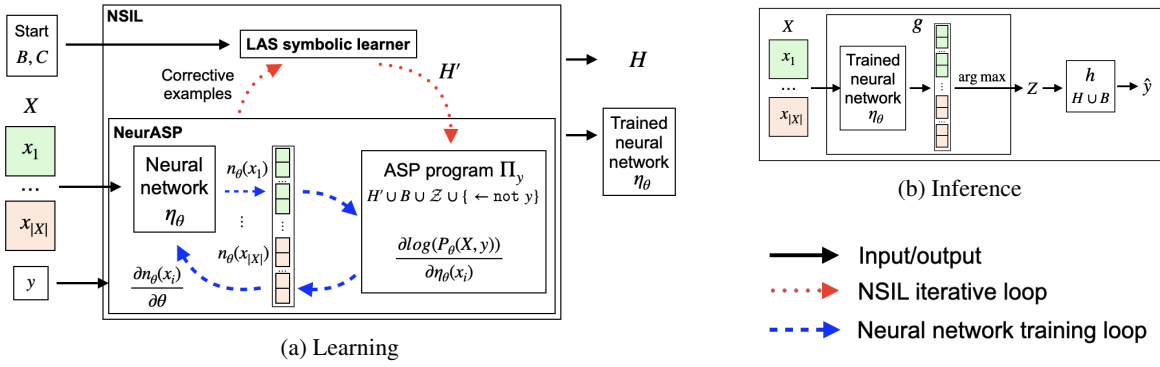


Figure 1: (a) NSIL learning with a single data point  $\langle X, y \rangle \in D$ . (b) NSIL inference over a single input  $X$ .

$X$  is a sequence of raw inputs  $x \in \mathcal{X}$ , and  $y \in \mathcal{Y}$  is a target label. We assume the latent concept  $C$  is given as background knowledge, but crucially, the samples in  $D$  are not annotated with latent concept labels.

**Example 1.** Consider the MNIST Addition task from [Manhaeve et al., 2018]. The goal is to learn a perception function  $g$  that classifies raw MNIST images into a latent concept  $C = \langle \text{digit}, \{0 \dots 9\} \rangle$ , and a reasoning function  $h$  that returns the sum of two latent concept values, where  $\mathcal{Y} = \{0 \dots 18\}$ .  $D$  contains training samples of two MNIST images, and the label indicates the sum, e.g.,  $X = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$  and  $y = 3$ .

To combine the benefits of neural perception with symbolic reasoning, we implement  $g$  using a neural network, and  $h$  using a first-order logic program in the language of ASP. Jointly learning both  $g$  and  $h$  is challenging, as  $h$  is not differentiable and cannot be learned using standard gradient-based optimisation. To learn  $h$ , we use a symbolic learner that explores a search space constructed from symbolic relations given in a domain knowledge  $B$ . Therefore, assuming  $B$  and  $C$  are given as input, the objective of NSIL is to learn  $g$  and  $h$ , s.t.  $\forall \langle X, y \rangle \in D, h([g(x_i) : x_i \in X]) = y$ . Let us now describe how  $g$  and  $h$  are learned within the NSIL architecture.

### 3.2 Neural and Symbolic Components

The NSIL architecture is presented in Figure 1. During learning, NSIL trains a neural network using knowledge learned by a symbolic learner. Corrective examples then bias the symbolic learner to refine or retain the learned knowledge based on the current system performance. The neural network is then updated again, as part of an iterative training process. During inference, the trained neural network forms the perception function  $g$ , and the learned knowledge forms the reasoning function  $h$ , to enable target predictions  $\hat{y} \in \mathcal{Y}$  from unseen raw data. Let us now describe the optimisation of the neural and symbolic components.

**Symbolic learner.** The symbolic component is a Learning from Answer Sets (LAS) symbolic learner [Law et al., 2019]. The input is a domain knowledge  $B$  that contains any (optional) background knowledge, and relations used to construct a search space  $\mathcal{H}$ . The symbolic learner also accepts a set of examples  $E$  as input, where in NSIL, each example associates

a sequence of latent concept values  $Z \in \mathcal{Z}^{a,b}$  with a target label  $y \in \mathcal{Y}$ . The output is a learned knowledge base  $H \in \mathcal{H}$  which is a first-order logic program in the language of ASP.  $H$ , together with  $B$ , forms the reasoning function  $h = H \cup B$  that maps sequences of latent concept values to target labels. We configure each example  $e \in E$  to represent the label either positively (i.e.,  $y$  should be output for  $Z$ ), or negatively ( $y$  should not be output for  $Z$ ).  $H$  is said to *cover* an example, if  $H \cup B$  outputs a desired label for the sequence of latent concept values (where a desired label depends on the positive or negative example configuration). Also,  $e$  is weighted with an integer penalty  $W(e)$ , paid by  $H$  for leaving the example uncovered. Details of how the examples are constructed are presented in Section 3.3. During learning, we add constraints to  $B$  to ensure  $H$  is guaranteed to output each target label we observe in  $D$ , i.e.,  $\forall \langle X, y \rangle \in D$ , there exists a sequence of latent concept values  $Z \in \mathcal{Z}^{a,b}$  s.t.  $h(Z) = y$ , where  $h = H \cup B$ . Also,  $H$  minimises a scoring function, based on the length of  $H$  (i.e., the number of literals), and the coverage over the examples  $E$ . Let us assume the set of examples uncovered by  $H$  is denoted by  $UNCOV(H, (B, E))$ . The score of  $H$  is given by  $score(H, (B, E)) = |H| + \sum_{e \in UNCOV(H, (B, E))} W(e)$ . A LAS symbolic learner solves the following optimisation:

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

This can be equivalently interpreted as jointly maximising the generality of  $H$  (i.e., the most compressed ASP program), and the coverage of the examples in  $E$ .

**Neural network.** For the neural component, we define a neural network  $\eta_\theta : \mathcal{X} \rightarrow [0, 1]^{|Z|}$  that maps a raw input  $x \in \mathcal{X}$  to a probability distribution over the possible latent concept values  $\mathcal{Z}$ .  $\theta$  denotes the neural network parameters which need to be learned. During inference, the perception function  $g$  returns the latent concept value with the maximum probability using the standard arg max function, i.e.,  $g(x) = \arg \max_{z \in \mathcal{Z}} (\eta_\theta(x)[z])$ . As there are no latent concept labels available, the neural network is trained using a symbolic knowledge  $H$ , alongside  $B$  and the target label  $y$ . Recall that the *space* of possible latent concept values  $\mathcal{Z}$  is given as background knowledge. As  $H$  is in the language of ASP, we use NeurASP [Yang et al., 2020] to optimise a semantic loss function [Xu et al., 2018] for neural network

training. Informally, for each training data point, the neural network is trained to predict the latent concept values that result in the target label, given  $H$  and  $B$ .

**Example 2.** Following Example 1, let us assume  $H$  is the correct addition rule. For  $X = \begin{bmatrix} 2 \\ 7 \end{bmatrix}$ ,  $y = 3$ , the neural network would be trained to output either  $[0, 3]$ ,  $[1, 2]$ ,  $[2, 1]$ , or  $[3, 0]$  for these images. Note that each of these latent concept values receives a gradient update, and the network learns to distinguish individual digits by observing multiple data points.

Formally, consider a training sample  $\langle X, y \rangle \in D$ . We define a matrix  $M_\theta(X)$  which contains all neural network outputs for each image within  $X$ .  $M_\theta(X)[i, z]$  denotes the neural network output for image  $x_i \in X$  and latent concept value  $z$ . E.g., following Example 2,  $M_\theta(X)[1, 2]$  corresponds to the probability of the first image  $\begin{bmatrix} 2 \\ 7 \end{bmatrix}$  having latent concept value 2. We also define an ASP program  $\Pi_y = H \cup B \cup Z \cup \{\leftarrow \text{not } y\}$ , with answer sets<sup>1</sup> denoted  $AS(\Pi_y)$ . Each answer set contains a sequence of latent concept values  $Z$  that output label  $y$ , given  $H$  and  $B$ . We define the probability of each  $Z \in AS(\Pi_y)$  in terms of the neural network output:  $P_{\theta, X}(Z) = \prod_{z_i \in Z} M_\theta(X)[i, z_i]$ . The probability of an  $\langle X, y \rangle$  training sample is then defined as  $P_\theta(X, y) = \sum_{Z \in AS(\Pi_y)} P_{\theta, X}(Z)$ , and the neural network is trained to optimise:

$$\theta^* = \arg \max_{\theta} \left[ \sum_{\langle X, y \rangle \in D} \log(P_\theta(X, y)) \right] \quad (2)$$

Intuitively, this can be interpreted as training the neural network to maximise the probability of the latent concept values that output each downstream label, given  $H$  and  $B$ . For further details, we refer the reader to [Yang *et al.*, 2020].<sup>2</sup> Now we have outlined the neural and symbolic components, we can introduce our novel method of integration. Recall that the symbolic learner learns a knowledge base which is used to train the neural network. An immediate question is; what happens if the learned knowledge is incorrect, as that may lead to a sub-optimal neural network? To tackle this problem, NSIL adopts an iterative training procedure where both the symbolic learner and the neural network are updated on each iteration. This integration relies upon a set of weighted *corrective examples* for the symbolic learner, that encourages the current knowledge to be refined or retained on the next iteration. Let us now outline how the corrective examples are structured and how their weights are updated on each iteration.

### 3.3 Corrective Examples

A corrective example  $e_{Z, y}$  associates a sequence of latent concept values  $Z \in \mathcal{Z}^{a, b}$  with a target label  $y \in \mathcal{Y}$ . Each example is configured to represent the label either positively (i.e.,  $y$  should be output for  $Z$ , denoted  $e_{Z, y}^{\text{pos}}$ ), or negatively ( $y$  should not be output for  $Z$ , denoted  $e_{Z, y}^{\text{neg}}$ ). Also, recall that each example is weighted with a penalty, paid if the example is left

<sup>1</sup>The answer sets correspond to the solutions of the program. See [Gelfond and Kahl, 2014] for an overview of ASP.

<sup>2</sup>In this paper, we assume each data point has a single  $y$  label. To simplify notation, we have presented a modified version of NeurASP that also assumes a single label. Generalising NSIL to multi-label tasks is left as future work.

*uncovered*. A pair of corrective examples  $\langle e_{Z, y}^{\text{pos}}, e_{Z, y}^{\text{neg}} \rangle$  represents both positive and negative cases for a given  $Z, y$  combination. In NSIL, the set of symbolic examples  $E$  contains a pair of corrective examples for all possible  $Z, y$  combinations, and initially, all of the examples have equal weight set to 0. This means there is no bias given to the symbolic learner as to which examples should be covered. We restrict the weights to the interval  $[0, 100]$ , and the goal is to maximise the weights of the positive examples that contain correct combinations of latent concept values with target labels, whilst minimising the weights of the corresponding negative examples.

On each iteration, the neural network is trained with a candidate knowledge base, denoted  $H'$ . The corrective example weights are then updated, using two sources of information: (1) The overall performance when predicting training set labels  $y$ . (2) The neural network confidence when predicting  $Z$ . In each case, we calculate a weight value, and associate this with a specific  $Z, y$  combination. To capture the overall performance, we obtain  $\hat{y}$  predictions with a forward pass over the training set using both neural and symbolic components. We then calculate the False Negative Rate (FNR) of each label  $y \in \mathcal{Y}$ , which depends on the current  $H'$  and current neural network parameters, denoted  $\theta'$ :

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

where  $TP_{H', \theta'}(y)$  and  $FN_{H', \theta'}(y)$  are the number of true positives and false negatives of label  $y$  respectively. We multiply by 100 to map the FNR into our  $[0, 100]$  weight range. Intuitively, if NSIL accurately predicts label  $y$ ,  $FNR_{H', \theta'}(y)$  will be close to 0, otherwise  $FNR_{H', \theta'}(y)$  will be close to 100. To associate the FNR with a  $Z, y$  combination, we obtain all possible  $Z$  that currently output label  $y$  given the current  $H'$ , by intercepting the answer sets of  $\Pi_y$  within NeurASP (see Section 3.2). We divide the FNR by the number of answer sets of  $\Pi_y$ , i.e.,  $\frac{FNR_{H', \theta'}(y)}{|AS(\Pi_y)|}$ , as there could be multiple  $Z$  that lead to each  $y$ . E.g., in Example 2 there are four sequences of latent concept values that lead to the label  $y = 3$ , so the FNR is shared between these possibilities.

To capture the neural network confidence, and associate it with a  $Z, y$  combination, we compute an average confidence over all training data points with a label  $y$ , where the neural network predicts  $Z$ . Formally, let  $D_{Z, y}$  be the set of  $\langle X, y \rangle \in D$  with label  $y$ , where the neural network predicts the sequence of latent concept values  $Z$  for the input sequence  $X$ , i.e.,  $[g(x_i) : x_i \in X] = Z$ . Let us also define the probability of  $X$  as the product of the maximal neural network confidence score for each input in the sequence:  $P_{\theta'}(X) = \prod_{x_i \in X} \max(\eta_{\theta'}(x_i))$ . The overall confidence for a  $Z, y$  combination is defined as:

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

Again, we multiply by 100 to map into our  $[0, 100]$  weight range. Given our two sources of information, we now outline

how the corrective example weights are updated. On each iteration, we firstly set the weights of the  $Z, y$  pairs that we have observed from the answer sets of  $\Pi_y$  as follows:

$$W(e_{Z,y}^{\text{pos}}) = 0 \quad (5)$$

$$W(e_{Z,y}^{\text{neg}}) = \frac{FNR_{H',\theta'}(y)}{|AS(\pi_y)|} \quad (6)$$

The negative example weight is set by the FNR, as if the FNR is large, this will encourage a different label for  $Z$ . The positive example weight is set to 0 because the FNR may apply to multiple  $Z, y$  pairs, and may not indicate any positive information for an individual pair. We then update the example weights using the neural network confidence to ensure positive information for an individual  $Z, y$  pair is captured:

$$W(e_{Z,y}^{\text{pos}}) = W(e_{Z,y}^{\text{pos}}) + \lambda \text{CONF}_{\theta'}(Z, y) \quad (7)$$

$$W(e_{Z,y}^{\text{neg}}) = W(e_{Z,y}^{\text{neg}}) - \lambda \text{CONF}_{\theta'}(Z, y) \quad (8)$$

where  $\lambda \in [0, 1]$  controls the effect of the  $\text{CONF}_{\theta'}$  update and is set to 1 in all of our experiments.<sup>3</sup> Example weights are allowed to persist across iterations, and in Equations 7 and 8, the  $W$  terms in the equation body could equal the weights set by Equations 5 and 6, or the weights set on a previous iteration, if a  $Z, y$  pair predicted by the neural network does not exist in the answer sets of  $\Pi_y$ . If there are no latent concept predictions of a certain  $Z$ ,  $\text{CONF}_{\theta'}(Z, y) = 0$ . Finally, the weights are clipped to the range  $[0, 100]$ .

**Example 3.** *Following Example 2, let  $H'$  be the correct addition rule,  $y = 3$  and  $FNR_{H',\theta'}(y) = 5$ , indicating 95% of training data points with label 3 have been predicted correctly. The answer sets of  $\Pi_y$  are  $[0, 3]$ ,  $[1, 2]$ ,  $[2, 1]$ , and  $[3, 0]$ . Let us also assume that for training data points with label 3, the correct  $Z$  are predicted with 98% confidence, and  $\lambda = 1$ . Therefore, the positive corrective examples associated with these latent concept values all have weight 98, and the corresponding negative examples all have weight  $(\frac{5}{4}) - 98 = -96.75$ , which is clipped to 0. This encourages the symbolic learner to cover the positive examples on the next iteration, and leave the negative examples uncovered.*

**Example 4.** *Conversely, let  $H'$  be an incorrect knowledge base,  $y = 3$ ,  $\lambda = 1$ , and  $FNR_{H',\theta'}(y) = 95$  (indicating only 5% of training data points with label 3 have been predicted correctly). Let us assume the answer sets of  $\Pi_y$  are  $[0, 4]$ ,  $[1, 2]$ ,  $[2, 3]$ , and, for all data points labelled 3, the neural network only predicts  $[1, 2]$  with 75% confidence. Therefore:*

$$W(e_{[1,2],3}^{\text{pos}}) = 75, \quad W(e_{[0,4],3}^{\text{pos}}) = W(e_{[2,3],3}^{\text{pos}}) = 0,$$

$$W(e_{[0,4],3}^{\text{neg}}) = W(e_{[2,3],3}^{\text{neg}}) = \frac{95}{3} = 31.67,$$

$$W(e_{[1,2],3}^{\text{neg}}) = \left(\frac{95}{3}\right) - 75 = -43.33 \text{ (which is clipped to 0).}$$

*This encourages the symbolic learner to cover  $e_{[1,2],3}^{\text{pos}}$ ,  $e_{[0,4],3}^{\text{neg}}$  and  $e_{[2,3],3}^{\text{neg}}$  on the next iteration.*

<sup>3</sup>Please see Appendix A.1 for results with varying values of  $\lambda$ .

For computational efficiency, example pairs with equal weights are not given to the symbolic learner. Also, weights are rounded to the nearest integer and shifted to the range  $\{1 \dots 101\}$  as the LAS systems require integer weight penalties  $> 0$ . To summarise, NSIL solves Equations 1 and 2 with the following steps:

1. The symbolic learner learns an initial knowledge base  $H'$  that satisfies  $B$  and covers each possible label in  $\mathcal{Y}$ . The neural network parameters  $\theta'$  are initialised randomly.
2. The neural network is trained for 1 epoch on the training set  $D$  to update  $\theta'$ , using the ASP program  $\Pi_y$  for each  $\langle X, y \rangle \in D$ . Note that  $\Pi_y$  is constructed using  $H'$ .
3. The weights of the corrective examples are updated (see Section 3.3), and a new  $H'$  is learned.
4. Steps 2-3 are repeated for a fixed number of iterations. NSIL outputs the learned neural network parameters  $\theta$  and the learned knowledge base  $H$ .

## 4 Experiments

Our experiments address the following questions: (1) Can NSIL train both neural and symbolic components, learning general, expressive and interpretable knowledge from raw data? (2) Is NSIL vulnerable to becoming stuck in local optima when training with limited data? (3) Can NSIL learn the correct knowledge in large search spaces? (4) How does NSIL compare to the closest existing system  $Meta_{Abd}$ ? (5) How does NSIL compare to end-to-end neural networks? Finally, (6) can NSIL learn expressive solutions to complex tasks that would be difficult to represent with other systems?

To answer these questions, we consider three problem domains; Cumulative Arithmetic, Two-Digit Arithmetic, and the Hitting Set problem. The Cumulative Arithmetic tasks follow [Dai and Muggleton, 2021] and enable a comparison to  $Meta_{Abd}$ , where NSIL achieves state-of-the-art performance, thus addressing question 4. To concretely demonstrate the contribution of NSIL, the other domains require learning more expressive knowledge that  $Meta_{Abd}$  can't learn, including negation as failure, choice rules, and programs with multiple answer sets. We demonstrate NSIL trains both neural and symbolic components successfully, outperforms fully neural baselines, and learns the correct knowledge in large search spaces. This addresses questions 1, 3, and 5. Question 2 is also addressed here, where a reduced amount of training data is used in the Two-Digit Arithmetic domain, and NSIL is able to train successfully under these conditions. Question 6 is addressed in the Hitting Set domain, where NSIL is able to solve an NP-complete problem and generate *all* the hitting sets of a given collection, despite learning from binary labels.

**Baselines.** Aside from the comparison to  $Meta_{Abd}$ , we compare to Concept Bottleneck Models (CBMs) [Koh *et al.*, 2020], which are end-to-end neural networks that model both perception and reasoning components. We ensure the perception network is the same as in NSIL, but the reasoning network is a multi-layer perceptron instead of a symbolic learner. We train two variants end-to-end (i.e., no concept labels are used), therefore using the same weak supervision as NSIL. One variant (denoted  $CBM$ ) uses a linear layer as output from the

Seq. Len.	Addition				Product			
	Acc.		MAE		Acc.		log MAE	
	1	5	10	100	1	5	10	15
<i>Meta_Abd</i>	0.953	0.510	1.299	6.587	0.977	0.334	0.495	2.374
NSIL	<b>0.984</b>	<b>0.208</b>	<b>0.550</b>	<b>3.960</b>	<b>0.983</b>	<b>0.291</b>	<b>0.469</b>	<b>2.098</b>

Table 1: Cumulative Arithmetic task results.

perception network, and the second variant (denoted *CBM-S*) uses a softmax layer, replicating more closely the perception network used in NSIL. For the Two-Digit Arithmetic tasks, we also evaluate: (1) The CNN from [Manhaeve *et al.*, 2018] that accepts a concatenated set of images as input and is trained to output the final prediction directly. (2) A CNN-LSTM with a Neural Arithmetic Logic Unit (NALU), and (3) a CNN-LSTM with a Neural Accumulator (NAC). For the Hitting Set tasks, we instead use a CNN-LSTM alongside the two CBM variants.

We also perform an ablation study where we firstly assume the symbolic knowledge is given, and secondly, assume the neural network is pre-trained with latent concept labels. NSIL has equivalent performance in these cases, whilst *learning* symbolic knowledge, and, training *without* latent concept labels. Finally, we demonstrate the effect of the novel corrective examples introduced in this paper, by presenting NSIL results *without* corrective examples. The ablation results are discussed in more detail in Section 5.

**Experiment setup.** In the Cumulative Arithmetic tasks, we use the same setup as [Dai and Muggleton, 2021], although use ILASP [Law, 2018] as our symbolic learner. In all other tasks, the neural network in NSIL is the MNIST CNN from [Manhaeve *et al.*, 2018], and the symbolic learner is FastLAS [Law *et al.*, 2020] for the Two-Digit Arithmetic tasks, and ILASP [Law, 2018] for the Hitting Set tasks. We use ILASP for the Cumulative Arithmetic task as ILASP enables the search space to be expressed using meta-rules. This ensures a fair comparison to *Meta\_Abd*, which requires a meta-rule specification. ILASP also enables us to learn choice rules in the Hitting Set tasks. We use FastLAS for the Two-Digit Arithmetic tasks to demonstrate NSIL’s flexibility to support different symbolic learners, and due to FastLAS’ increased efficiency in larger search spaces for this specific class of problems. To evaluate NSIL, we compare the final prediction to the ground truth label. Each experiment is repeated 20 times using 20 randomly generated seeds, and we measure mean classification accuracy per training epoch, as well as the standard error over the 20 repeats, where 1 epoch = 1 NSIL iteration. We perform hyper-parameter tuning for all methods using a held-out validation set and a separate random seed. Each experimental repeat is allowed to run for a maximum of 24 hours. The ASP encodings of the domain knowledge and corrective examples are presented in Appendix A.1.

### 4.1 Cumulative Arithmetic

We first consider the Cumulative Addition and Product tasks from [Dai and Muggleton, 2021]. The dataset  $D$  contains  $\langle X, y \rangle$  samples where  $X$  is a sequence of MNIST digit images, and  $y$  is the cumulative addition or product of the digits in  $X$ . The latent concept has the name *digit* with values  $\mathcal{Z} =$




Figure 2: Cumulative Arithmetic network accuracy.

$\{0 \dots 9\}$ . The goal is to train the neural network to classify the digit images in  $X$ , and learn a recursive knowledge base that defines the cumulative arithmetic operation. The training set contains 3,000 samples of sequences of length 2-5, and three test sets with sequences of lengths 5, 10, and 100 in the Addition task, and 5, 10, and 15 in the Product task, to verify the trained network and learned knowledge can extrapolate to longer inputs. Each test set contains 10,000 samples. As in [Dai and Muggleton, 2021], we evaluate the neural network accuracy (**Acc.**) on predicting single digits, and the mean average error (**MAE**), and log mean average error (**log MAE**) when extrapolating to longer inputs in the addition and product tasks respectively. For this experiment only, we report results over 5 repeats to match [Dai and Muggleton, 2021].

In both tasks, NSIL learns the correct recursive knowledge (see Figure 5), and outperforms *Meta\_Abd* (Table 1). NSIL has a higher neural network accuracy, closer to the state-of-the-art accuracy of 0.9991 when a CNN is trained in a fully supervised fashion [An *et al.*, 2020], despite NSIL being trained without latent concept labels. Also, NSIL trains the neural network to its maximal accuracy with significantly less training epochs than *Meta\_Abd* (see Figure 2). *Meta\_Abd* relies upon computing the exact psuedo-labels for each image in order to train the network accurately, whereas in NSIL, NeurASP can train the network with multiple digit choices for each image using a semantic loss function. In Figure 2, the error bars and shaded regions indicate standard error over the 5 repeats. We plot the *Meta\_Abd* Addition results from their GitHub repository, and re-run their Product task with 5 randomly generated seeds as these results are not available. We also plot the best *Meta\_Abd* result compared to the average, to demonstrate that *Meta\_Abd* can be highly sensitive to the random seed used to initialise the network in tasks with dense abductive spaces.

### 4.2 Two-Digit Arithmetic

We consider the MNIST Addition task for two single digit numbers used in [Manhaeve *et al.*, 2018], and a variation called MNIST Even9Plus (E9P) that requires learning negation as failure, which is not supported by *Meta\_Abd*. In this task, the same pair of images are used, but the label 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. The purpose of the E9P task is to demonstrate NSIL can be used to solve multiple tasks within the same domain, i.e., the search space does not have to be engineered *exactly* to solve a particular task. In this case, both sub-tasks are defined over the same set of inputs and outputs, which is more challenging than the Cumulative Arithmetic domain where the output space is

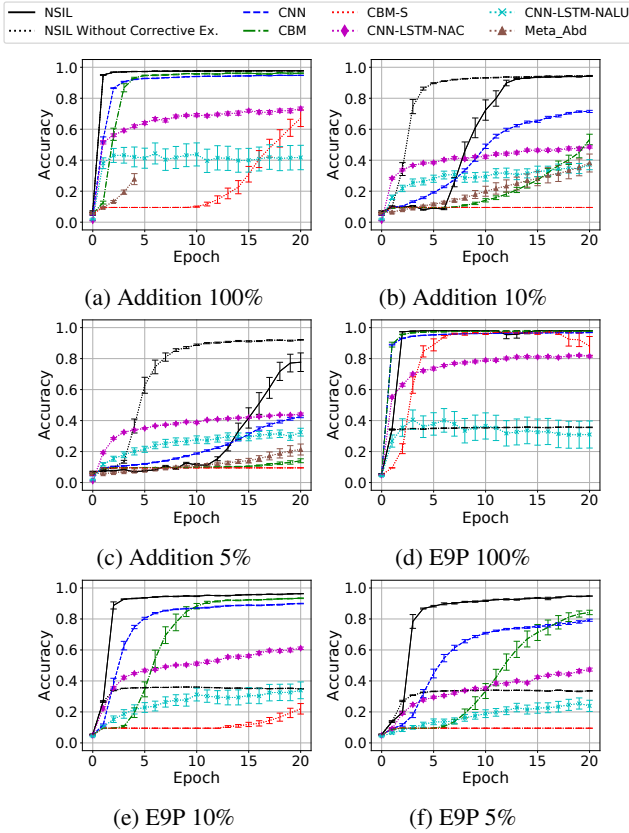


Figure 3: Two-Digit Arithmetic results with reduced training sets.

significantly different between the Addition and Product sub-tasks. We include a comparison to  $Meta_{Abd}$  on the MNIST Addition task, but note that a significantly reduced search space is used compared to NSIL, as the relations required for E9P are not included, given that  $Meta_{Abd}$  does not support learning negation as failure. Both tasks use training, validation and test datasets of size 24,000, 6,000, and 5,000 samples respectively. The relations even, not even, plus\_nine, and '=', as well as the function '+' are given in  $B$  to specify the search space for the symbolic learner. The latent concept is the same as in the Cumulative Arithmetic tasks.

The overall accuracy results are shown in Figure 3 for reducing percentages of training data (100% - 5%). NSIL outperforms the neural baselines on both tasks for all dataset percentages, and learns the correct knowledge (see Figure 5). On the Addition task, NSIL achieves superior performance in a much larger search space compared to  $Meta_{Abd}$ , which also times out after iteration 4 with 100% of the data. NSIL trains the neural network to classify MNIST images with a mean accuracy of 0.989, 0.970, and 0.891, for each dataset percentage respectively, which in the 100% case is within 1% of the performance of a fully supervised model (0.9991) [An *et al.*, 2020]. For the E9P task, the mean network accuracy is 0.986, 0.974, and 0.962, for each dataset percentage respectively. NSIL has greater performance in the E9P task than in the Addition task because the label is more informative, despite the same domain knowledge and search space. There is a reduced

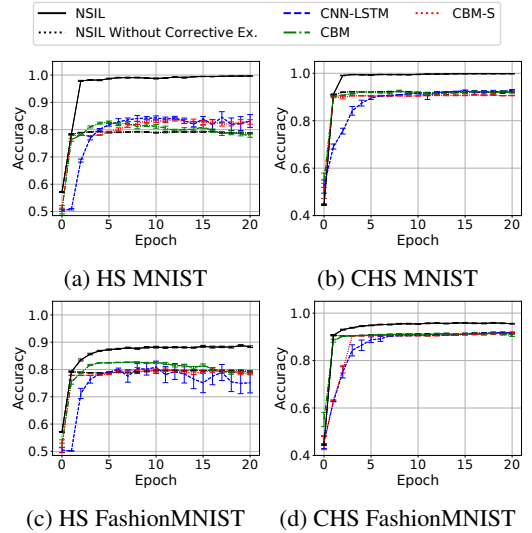


Figure 4: Hitting Set accuracy.

number of possible latent concept values for certain labels, which gives a stronger back-propagation signal for training the network. Finally, we increased the search space in the E9P task with the additional functions; '-', '×', '÷', as well as plus\_eight, ..., plus\_one relations, and NSIL converged to the expected knowledge in all of these cases (see Table 2a).

### 4.3 Hitting Set Problem

We consider two variations of the Hitting Set problem [Karp, 1972]. The first is the standard formulation: Consider a universe of elements  $U$ , and a collection of sets  $S$  which contains subsets of  $U$ . A hitting set is a special set  $\Phi \subseteq U$  that intersects with every subset in  $S$ . Given  $S$  and an integer  $k$ , the task is to decide if there is a hitting set of  $S$  with size  $\leq k$ .

**Example 5.** Let  $U = \{1 \dots 5\}$ ,  $S_1 = \{\{1\}, \{1, 3\}, \{4\}\}$ , and  $k = 3$ .  $S_1$  has a hitting set of size  $\leq 3$ :  $\Phi = \{1, 4\}$ . We say that each subset in  $S$  is "hit" if it intersects with the hitting set. Now consider  $S_2 = \{\{1\}, \{2, 3\}, \{4\}\}$ .  $S_2$  has no hitting set of size  $\leq 3$ .

In our experiments, we replace elements  $\{1 \dots 5\}$  in  $U$  with corresponding MNIST images from classes 1-5 (i.e.,  $\mathcal{Z} = \{1 \dots 5\}$ ). We assume  $k = 2$  and construct  $\langle X, y \rangle$  data points where  $X$  represents a collection of sets  $S$ , and  $y$  is a binary label indicating whether  $X$  has a hitting set of size  $\leq k$ . We also assume a maximum of 4 elements,  $X$  is well-formed, with no duplicate elements in a set, and elements within the same set are arranged in ascending order. For example;  $X = \{\{\mathbf{7}\}, \{\mathbf{1}, \mathbf{3}\}, \{\mathbf{4}\}\}$ ,  $y = 1$ . We denote this standard hitting set variant as HS. To demonstrate NSIL can solve multiple tasks within the same domain, we introduce a second variant, denoted CHS, which adds the constraint that no hitting sets occur if  $|X| \geq 3$ , and element 1 (in image form) exists in any subset in  $X$ . We also present results for both variants with a more difficult perception task using FashionMNIST images instead of MNIST. The goal is to train the neural network to classify images, whilst learning the HS or CHS rules to decide the existence of a hitting set. Training,

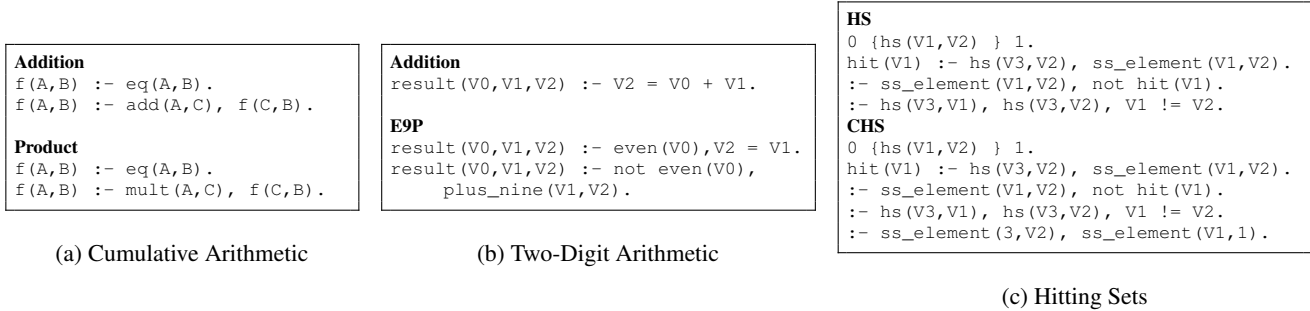


Figure 5: NSIL learned knowledge.

validation and test datasets contain 1502, 376, and 316 examples respectively. Each latent concept value is represented in ASP as `ss_element(s, z)`, where  $s \in \{1..4\}$  is a subset identifier, and  $z \in \{1..5\}$  is the associated element. The domain knowledge  $B$  contains  $k = 2$ , the relation `!=` and the relation `hs`, which defines an element of the hitting set.  $B$  also contains the relations for the CHS task: `ss_element(3, V1)` and `ss_element(V2, 1)` that define respectively subset 3 (and any element  $V_1$ ), and element 1 in any subset  $V_2$ . In these tasks, we use a CNN-LSTM baseline in addition to the CBM variants, where a CNN firstly encodes a feature vector for each image in a sequence  $X$ , and an LSTM returns a binary classification as to the existence of a hitting set. To encode the subset structure of  $X$  for the CNN-LSTM, we create an image of ‘{’ and ‘}’ characters and pass in the entire sequence as input. Note that in NSIL, the structure is passed directly into the symbolic learner using the `ss_element` relations.

Figure 4 shows the overall accuracy for both Hitting Set tasks using MNIST and FashionMNIST images. NSIL outperforms the baselines on both tasks and learns the correct knowledge (see Figure 5). When classifying images, the neural network achieves a mean accuracy of 0.993 (MNIST) and 0.897 (FashionMNIST) on the HS task, and 0.993 (MNIST) and 0.894 (FashionMNIST) on the CHS task. Note that the neural network accuracy is for concepts 1-5 only. All methods perform better on the CHS task than the HS task. This is due to a shortcut; a negative test example with 3 subsets can be correctly predicted if the neural network predicts element 1 for *any* image within the subset, as opposed to having to rely on accurate predictions for more images in the HS task.

The knowledge learned by NSIL has multiple answer sets as there could be multiple hitting sets of a given collection. Therefore, despite training labels indicating the existence of a hitting set, NSIL can generate *all* the hitting sets at inference time. We verified this by computing the hamming loss against the ground truth hitting sets for all test samples. The results are 0.003 and 0.150 on the HS task, and 0.003 and 0.126 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 knowledge. Also, it is crucial to note that the learned knowledge does not depend on the number of elements, the value of  $k$ , or the possible latent concept values, despite these restrictions during training.

This demonstrates NSIL is able to *generalise* beyond the training setup. Finally, in the more challenging HS FashionMNIST task, we also sampled 6 sets of additional `ss_element` domain relations to extend the search space for the symbolic learner. The results are presented in Table 2b, and NSIL learned the correct knowledge in all of these cases, but sometimes required more iterations and/or time to do so. The *Standard* domain knowledge contains only the predicates required to learn the expected rules of the HS task. This includes  $k = 2$ , the `!=` relation, and the relation `hs` which defines an element of the hitting set. The additional `ss_element` relations contain a combination of elements  $\{1..5\}$  (denoted ‘el’) and subset identifiers  $\{1..4\}$  (denoted ‘ssID’).

	Domain knowledge	Convergence iteration	Convergence time (s)
*	+, =, even, not even, plus_nine	2	475.44
	+, -, ×, ÷, =, even, not even, plus_nine	2	547.81
	=, even, not even, plus_nine, ..., plus_one	1	<i>402.03</i>
	+, =, even, not even, plus_nine, ..., plus_one	2	<i>881.35</i>
	+, -, ×, ÷, =, even, not even, plus_nine, ..., plus_one	2	<i>57869.12</i>

(a) E9P

	Domain knowledge	Convergence iteration	Convergence time (s)
	Standard, ssID 2, ssID 4	1	111.34
	Standard, ssID 2, el 2	1	113.21
	Standard	1	115.77
	Standard, ssID 1, el 4	1	130.71
	Standard, ssID 3, el 4	2	734.70
*	Standard, ssID 3, el 1	2	852.84
	Standard, ssID 1, ssID 3	2	888.54

(b) HS FashionMNIST

Table 2: Increasing the search space for the symbolic learner. The asterisks indicate the configuration used in the main results. Run-times in italics indicate that weight pruning is used to remove corrective examples with weight  $< 5$ , to obtain results efficiently.



Dataset %	Addition			E9P		
	100	10	5	100	10	5
FF-NSL	<b>0.9753</b> (0.0021)	0.9362 (0.0029)	<b>0.9151</b> (0.0058)	<b>0.9809</b> (0.0016)	0.9513 (0.0030)	0.9346 (0.0051)
NeurASP	<b>0.9762</b> (0.0013)	<b>0.9492</b> (0.0016)	<b>0.9149</b> (0.0051)	<b>0.9797</b> (0.0015)	<b>0.9642</b> (0.0009)	<b>0.9500</b> (0.0014)
NSIL	<b>0.9762</b> (0.0013)	0.9449 (0.0025)	0.8782 (0.0134)	<b>0.9816</b> (0.0009)	<b>0.9634</b> (0.0007)	<b>0.9510</b> (0.0016)

Table 3: Two-Digit Arithmetic ablation results. Best results highlighted in bold (we declare a tie in the case of overlapping error bars).

Images	HS		CHS	
	MNIST	FashionMNIST	MNIST	FashionMNIST
FF-NSL	0.9937 (0.0017)	0.8816 (0.0110)	0.9962 (0.0012)	<b>0.9563</b> (0.0034)
NeurASP	<b>0.9981</b> (0.0013)	<b>0.8975</b> (0.0041)	<b>0.9994</b> (0.0006)	<b>0.9538</b> (0.0070)
NSIL	<b>0.9962</b> (0.0012)	0.8747 (0.0053)	<b>0.9981</b> (0.0013)	<b>0.9544</b> (0.0021)

Table 4: Hitting Set ablation results. Best results highlighted in bold (we declare a tie in the case of overlapping error bars).

### 5 Ablations

In this section, we present three ablations: (1) We remove the requirement to learn the symbolic knowledge and assume the knowledge is given. (2) We pre-train the neural network with manually annotated latent concept labels, and (3) we evaluate NSIL without corrective examples in order to demonstrate the impact of the novel contribution of this paper. To implement the first two ablations, we respectively evaluate NeurASP [Yang *et al.*, 2020], which assumes the symbolic knowledge is given, and FF-NSL [Cunnington *et al.*, 2023] which pre-trains the neural network on a separate dataset. FF-NSL then uses the neural network predictions to train the symbolic learner. We provide a comparison to NSIL, and train each method on the experiments presented in Section 4 for 20 epochs (where 1 epoch = 1 NSIL iteration). We evaluate the mean performance and standard error over 5 repeats.

The Two-Digit Arithmetic results are presented in Table 3 for reducing percentages of training data (100% - 5%). The results indicate final task accuracy after performing a forward pass on the neural and symbolic components. In both the Addition and E9P tasks, NSIL achieves comparable accuracy to the fully supervised approach of FF-NSL, even when *no latent concept labels* are observed during training. Also, NSIL achieves comparable accuracy to NeurASP, whilst also *learning* the symbolic knowledge. The Hitting Set results are shown in Table 4, where the results also indicate final task accuracy. Again, NSIL performs comparably to FF-NSL and NeurASP in each experiment.

In the final ablation study, the corrective examples enable NSIL to converge to a more accurate solution in most cases, by biasing the symbolic learner to explore the symbolic search space. This is evidenced in the E9P task shown in Figures 3d-3f, where the correct knowledge is *not* learned on the first iteration (compare *NSIL* to *NSIL Without Corrective Ex.*). In cases where the correct knowledge *is* learned on the first iteration, the effect of the corrective examples depends on the size of the dataset. When there are sufficient data points to train the neural network accurately, the corrective examples don’t offer any improvement, because the network can be trained given the correct knowledge (see Figure 3a). For small datasets, the corrective examples together with a sub-optimal

network may lead to a delay in convergence compared to the case where there are no corrective examples (see Figures 3b and 3c). However, in real-world applications, we consider these to be very special cases, as it’s unlikely the correct knowledge will be learned on the first iteration. Therefore, corrective examples are important to NSIL’s performance.

### 6 Conclusion

This paper presents NSIL which learns expressive knowledge to solve computationally complex problems, whilst training a general neural network to classify latent concepts from raw data. The novelty is the method of biasing the symbolic learner using corrective examples weighted to balance exploration and exploitation of the symbolic search space. The results show that NSIL outperforms the closest existing system *Meta\_Abd*, and a variety of purely neural baselines, whilst learning more expressive and general knowledge.

**Limitations.** During NSIL’s iterative training cycle, using the corrective examples may require more iterations to converge given a small dataset. However, the benefit is that a high accuracy is reached in cases where the correct knowledge is not learned on the first iteration. Also, the neural baselines learn faster than NSIL, although given more time, NSIL achieves a greater accuracy (see Appendix A.2).

**Future work.** One could investigate theoretical properties of NSIL’s corrective examples, and also apply NSIL to more challenging tasks with more complex imagery. In these cases, the neural network could be trained for more epochs on each NSIL iteration, or utilise pre-trained weights on a larger dataset such as ImageNet. This would enable NSIL to be used as a “fine-tuning” approach to learn the weights of a linear layer for classifying latent concepts. Finally, future work could generalise NSIL to support multiple latent concepts by integrating multiple neural networks, and extend NSIL to take advantage of unsupervised learning techniques.

### Ethical Statement

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 as ensuring training data is unbiased, and model inputs/outputs are monitored for adversarial attacks. As NSIL learns human interpretable knowledge 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 knowledge. 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.

## Acknowledgements

The authors would like to thank Krysia Broda for helpful discussions, and also the reviewers for their insightful comments.

## References

- [An *et al.*, 2020] Sanghyeon An, Minjun Lee, Sanglee Park, Heerin Yang, and Jungmin So. An ensemble of simple convolutional neural network models for mnist digit recognition. *arXiv preprint arXiv:2008.10400*, 2020.
- [Aspis *et al.*, 2022] Yaniv Aspis, Krysia Broda, Jorge Lobo, and Alessandra Russo. Embed2sym-scalable neuro-symbolic reasoning via clustered embeddings. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 19, pages 421–431, 2022.
- [Badreddine *et al.*, 2022] Samy Badreddine, Artur d’Avila Garcez, Luciano Serafini, and Michael Spranger. Logic tensor networks. *Artificial Intelligence*, 303:103649, 2022.
- [Besold *et al.*, 2022] Tarek R Besold, Artur d’Avila Garcez, Sebastian Bader, Howard Bowman, Luis C Lamb, Leo de Penning, BV Illuminoo, Hoifung Poon, and COPPE Gerson Zaverucha. Neural-symbolic learning and reasoning: A survey and interpretation. *Neuro-Symbolic Artificial Intelligence: The State of the Art*, 342:1, 2022.
- [Corapi *et al.*, 2010] Domenico Corapi, Alessandra Russo, and Emil Lupu. Inductive logic programming as abductive search. In *Technical communications of the 26th international conference on logic programming*, pages 54–63. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2010.
- [Cunnington *et al.*, 2023] Daniel Cunnington, Mark Law, Jorge Lobo, and Alessandra Russo. Ffnsl: Feed-forward neural-symbolic learner. *Machine Learning*, pages 1–55, 2023.
- [Dai and Muggleton, 2021] Wang-Zhou Dai and Stephen Muggleton. Abductive knowledge induction from raw data. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 1845–1851. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.
- [Dai *et al.*, 2019] Wang-Zhou Dai, Qiuling Xu, Yang Yu, and Zhi-Hua Zhou. Bridging machine learning and logical reasoning by abductive learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [Daniele *et al.*, 2022] Alessandro Daniele, Tommaso Campari, Sagar Malhotra, and Luciano Serafini. Deep symbolic learning: Discovering symbols and rules from perceptions. *arXiv preprint arXiv:2208.11561*, 2022.
- [Dantsin *et al.*, 2001] Evgeny Dantsin, Thomas Eiter, Georg Gottlob, and Andrei Voronkov. Complexity and expressive power of logic programming. *ACM Computing Surveys (CSUR)*, 33(3):374–425, 2001.
- [Dasaratha *et al.*, 2023] Sridhar Dasaratha, Sai Akhil Puranam, Karmvir Singh Phogat, Sunil Reddy Tiyyagura, and Nigel P. Duffy. Deeppsl: End-to-end perception and reasoning, 2023.
- [d’Avila Garcez *et al.*, 2019] Artur d’Avila Garcez, Marco Gori, Luís C. Lamb, Luciano Serafini, Michael Spranger, and Son N. Tran. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. *FLAP*, 6(4):611–632, 2019.
- [De Raedt *et al.*, 2020] Luc De Raedt, Sebastijan Dumančić, Robin Manhaeve, and Giuseppe Marra. From statistical relational to neuro-symbolic artificial intelligence. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4943–4950. International Joint Conferences on Artificial Intelligence Organization, 7 2020. Survey track.
- [Evans and Grefenstette, 2018] Richard Evans and Edward Grefenstette. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018.
- [Evans *et al.*, 2021] Richard Evans, Matko Bošnjak, Lars Buesing, Kevin Ellis, David Pfau, Pushmeet Kohli, and Marek Sergot. Making sense of raw input. *Artificial Intelligence*, 299:103521, 2021.
- [Ferreira *et al.*, 2022] Joao Ferreira, Manuel de Sousa Ribeiro, Ricardo Gonçalves, and Joao Leite. Looking inside the black-box: Logic-based explanations for neural networks. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 19, pages 432–442, 2022.
- [Gelfond and Kahl, 2014] Michael Gelfond and Yulia Kahl. *Knowledge representation, reasoning, and the design of intelligent agents: The answer-set programming approach*. Cambridge University Press, Cambridge, UK, 2014.
- [Han *et al.*, 2021] Zhongyi Han, Benzhen Wei, Xiaoming Xi, Bo Chen, Yilong Yin, and Shuo Li. Unifying neural learning and symbolic reasoning for spinal medical report generation. *Medical Image Analysis*, 67:101872, 2021.
- [Karp, 1972] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- [Koh *et al.*, 2020] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *International Conference on Machine Learning*, pages 5338–5348, 2020.
- [Law *et al.*, 2019] Mark Law, Alessandra Russo, and Krysia 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*, pages 196–231, 2019.
- [Law *et al.*, 2020] Mark Law, Alessandra Russo, Elisa Bertino, Krysia Broda, and Jorge Lobo. Fastlas: scalable inductive logic programming incorporating domain-specific optimisation criteria. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2877–2885, 2020.
- [Law, 2018] Mark Law. *Inductive learning of answer set programs*. PhD thesis, Imperial College London, 2018.
- [Lyn Paul *et al.*, 2019] Celeste Lyn Paul, Leslie M Blaha, Corey K Fallon, Cleotilde Gonzalez, and Robert S Gutzwiller. Opportunities and challenges for human-machine teaming in cybersecurity operations. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 63, pages 442–446. SAGE Publications Sage CA: Los Angeles, CA, 2019.
- [Manhaeve *et al.*, 2018] Robin Manhaeve, Sebastijan Dumanic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Deepproblog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems*, pages 3749–3759, 2018.
- [Muggleton and De Raedt, 1994] Stephen Muggleton and Luc De Raedt. Inductive logic programming: Theory and methods. *The Journal of Logic Programming*, 19:629–679, 1994.
- [Muggleton *et al.*, 2015] Stephen H Muggleton, Dianhuan Lin, and Alireza Tamaddoni-Nezhad. Meta-interpretive learning of higher-order dyadic datalog: Predicate invention revisited. *Machine Learning*, 100(1):49–73, 2015.
- [Payani and Fekri, 2019] Ali Payani and Faramarz Fekri. Inductive logic programming via differentiable deep neural networks. *CoRR*, abs/1906.03523, 2019.
- [Riegel *et al.*, 2020] Ryan Riegel, Alexander G. Gray, Francois P. S. Luus, Naweed Khan, Ndivhuwo Makondo, Ismail Yunus Akhalwaya, Haifeng Qian, Ronald Fagin, Francisco Barahona, Udit Sharma, Shajith Iqbal, Hima Karanam, Sumit Neelam, Ankita Likhyan, and Santosh K. Srivastava. Logical neural networks. *CoRR*, abs/2006.13155, 2020.
- [Sen *et al.*, 2022] Prithviraj Sen, Breno WSR de Carvalho, Ryan Riegel, and Alexander Gray. Neuro-symbolic inductive logic programming with logical neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8212–8219, 2022.
- [Shindo *et al.*, 2021] Hikaru Shindo, Masaaki Nishino, and Akihiro Yamamoto. Differentiable inductive logic programming for structured examples. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5034–5041, May 2021.
- [Sittón *et al.*, 2019] Inés Sittón, Ricardo S Alonso, Elena Hernández-Nieves, Sara Rodríguez-Gonzalez, and Alberto Rivas. Neuro-symbolic hybrid systems for industry 4.0: A systematic mapping study. In *International Conference on Knowledge Management in Organizations*, pages 455–465. Springer, 2019.
- [Xu *et al.*, 2018] Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Broeck. A semantic loss function for deep learning with symbolic knowledge. In *International conference on machine learning*, pages 5502–5511. PMLR, 2018.
- [Yang *et al.*, 2020] Zhun Yang, Adam Ishay, and Joohyung Lee. Neurasp: Embracing neural networks into answer set programming. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 1755–1762. International Joint Conferences on Artificial Intelligence Organization, 7 2020.