

Deciphering Raw Data in Neuro-Symbolic Learning with Provable Guarantees

Lue Tao,^{1,2} Yu-Xuan Huang,^{1,2} Wang-Zhou Dai,^{1,3} Yuan Jiang^{1,2}

¹National Key Laboratory for Novel Software Technology, Nanjing University, China

²School of Artificial Intelligence, Nanjing University, China

³School of Intelligence Science and Technology, Nanjing University, China

{taol, huangyx, daiwz, jiangy}@lamda.nju.edu.cn

Abstract

Neuro-symbolic hybrid systems are promising for integrating machine learning and symbolic reasoning, where perception models are facilitated with information inferred from a symbolic knowledge base through logical reasoning. Despite empirical evidence showing the ability of hybrid systems to learn accurate perception models, the theoretical understanding of learnability is still lacking. Hence, it remains unclear why a hybrid system succeeds for a specific task and when it may fail given a different knowledge base. In this paper, we introduce a novel way of characterising supervision signals from a knowledge base, and establish a criterion for determining the knowledge’s efficacy in facilitating successful learning. This, for the first time, allows us to address the two questions above by inspecting the knowledge base under investigation. Our analysis suggests that many knowledge bases satisfy the criterion, thus enabling effective learning, while some fail to satisfy it, indicating potential failures. Comprehensive experiments confirm the utility of our criterion on benchmark tasks.

1 Introduction

Integrating machine learning and symbolic reasoning is a holy grail challenge in artificial intelligence. This pursuit has attracted much attention over the past decades (Garcez, Broda, and Gabbay 2002; Getoor and Taskar 2007; Russell 2015; De Raedt et al. 2021; Hitzler and Sarker 2022), leading to fruitful developments such as probabilistic logic programming (De Raedt and Kimmig 2015) and statistical relational artificial intelligence (De Raedt et al. 2016).

In recent years, great progress has been made in neuro-symbolic methods, equipping symbolic systems with the ability to perceive sub-symbolic data. One intriguing finding in these hybrid systems is that the perception performance of initialised classifiers can be significantly enhanced through *abduction*, a.k.a. *abductive reasoning* (Dai et al. 2019; Li et al. 2020). Moreover, it has been shown that accurate classifiers can be learned from scratch without relying on fully labelled data, given appropriate objectives and knowledge bases (Xu et al. 2018; Manhaeve et al. 2018; Tsamoura, Hospedales, and Michael 2021; Dai and Muggleton 2021).

These advances highlight the value of symbolic reasoning in many learning tasks. However, not all symbolic knowl-

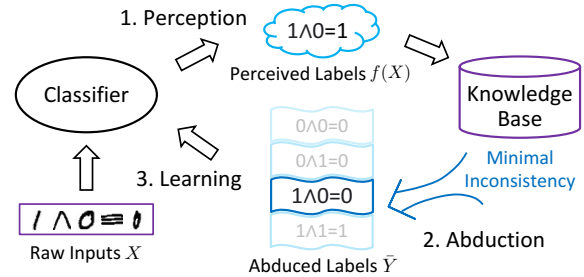


Figure 1: An illustration of the hybrid learning framework. First, raw data such as handwritten equations are perceived by a classifier. Next, the perceived labels are revised via logical abduction under the principle of minimal inconsistency. Finally, the abduced labels are used to update the classifier.

edge helps improve learning performance; there are failures in practice (Cai et al. 2021; Marconato et al. 2023a,b). More importantly, the theoretical underpinnings that drive these empirical successes or failures remain elusive, which may hinder the adoption of neuro-symbolic methods in other applications. In particular, it is unclear why such a hybrid learning system works for a specific task and when it may fail given a different knowledge base.

In this paper, we address the above questions under the framework of *abductive learning* (ABL), an expressive and representative paradigm of hybrid systems (Zhou 2019; Zhou and Huang 2022). As illustrated in Fig. 1, a hybrid learning system usually involves the perception of raw inputs using a classifier, followed by an abductive reasoning module (Kakas, Kowalski, and Toni 1992) that aims to correct the wrongly perceived labels by minimising the inconsistency with a given symbolic knowledge base.

Contributions. We present a theoretical analysis illuminating the factors underlying the success of hybrid systems. Our analysis is based on a key insight: the objective of popular hybrid methods implies a secret objective that finely characterises supervision signals from a given knowledge base. Formally, we show that the objective of minimising the inconsistency with a knowledge base, under reasonable conditions, is equivalent to optimising an upper bound of another objective denoted by L-Risk, which contains a probability matrix relating the ground-truth labels and their locations in

the ground atoms of the knowledge base. This indicates that hybrid methods can make progress by mitigating the L-Risk. Further, we show that if the matrix in the objective is full-rank, the ground-truth values of the labels perceived from raw inputs are guaranteed to be recovered. In this manner, we establish a rank criterion that indicates the knowledge’s efficacy in facilitating successful learning.

To our knowledge, this is the first study that attempts to provide a reliable diagnosis of the knowledge base in abductive learning prior to actual training. Our theoretical analysis offers practical insights: if the knowledge base meets the criterion, it can facilitate learning; otherwise, it might fail and require further refinement to ensure successful learning. Comprehensive experiments on benchmark tasks with different knowledge bases validate the utility of the criterion. We believe that our findings are instrumental in guiding the integration of machine learning and symbolic reasoning.

2 Preliminaries

Conventional Supervised Learning. Let $\mathcal{X} \subseteq \mathbb{R}^d$ be the input space and $\mathcal{Y} = [c] = \{0, \dots, c-1\}$ be the label space, where c is the number of classes. In ordinary multi-class learning, each instance-label pair $\langle x, y \rangle \in \mathcal{X} \times \mathcal{Y}$ is sampled from an underlying distribution with probability density $p(x, y)$, and the objective is to learn a mapping $h : \mathcal{X} \rightarrow \mathbb{R}^c$ that minimises the expected risk over the distribution:

$$\mathcal{R}(h) = \mathbb{E}_{p(x,y)} \ell(h(x), y), \quad (1)$$

where $\ell : \mathbb{R}^c \times \mathcal{Y} \rightarrow \mathbb{R}$ is a loss function that measures how well the classifier perceives an input. The predicted label of the classifier is represented by $f(x) = \arg \max_{i \in \mathcal{Y}} h_i(x)$, where $h_i(x)$ is the i -th element of $h(x)$.

Neuro-Symbolic Learning. In neuro-symbolic (NeSy) learning systems, it is common to assume that raw inputs $X = [x_0, x_1, \dots, x_{m-1}]$ are given, while their ground-truth labels $Y = [y_0, y_1, \dots, y_{m-1}]$ are not observable. Instead, we only know that the logical facts grounded by the labels are compatible with a given knowledge base. Specifically, we have $B \cup Y \models \tau$, where B is a knowledge base consisting of first-order logic rules, \models denotes logical entailment, and τ is a target concept in a concept space \mathcal{T} .

Example 1. Consider a binary classification task with label space $\mathcal{Y} = \{0, 1\}$ and concept space $\mathcal{T} = \{\text{conj}\}$, which contains a single target concept “conj”. Each sequence X consists of three raw inputs, whose ground-truth labels Y are unknown but satisfy the logical equation $y_0 \wedge y_1 = y_2$. Fig. 2 illustrates the logical facts abduced from the knowledge base and several sequences of raw inputs in this task.

In Example 1, when observing a sequence of raw inputs $X = [1, 0, 0]$ with the target concept “conj”, the classifier should learn to perceive the inputs so that the logical equation $f(1) \wedge f(0) = f(0)$ holds. In general, denote by $p(X, \tau)$ the underlying distribution of the input sequence $X \in \mathcal{X}^m$ and the target concept $\tau \in \mathcal{T}$, the objective is to learn a mapping $h : \mathcal{X} \rightarrow \mathbb{R}^c$ that minimises the inconsistency between the classifier and the knowledge base:

$$\mathcal{R}_{\text{NeSy}}(h) = \mathbb{E}_{p(X,\tau)} \mathcal{L}(X, \bar{Y}; h), \text{ s.t. } B \cup \bar{Y} \models \tau, \quad (2)$$




Figure 2: Illustration of the knowledge base about conjunction, the facts abduced from the knowledge base, and the raw inputs corresponding to the target concept “conj”.

where $\mathcal{L}(X, \bar{Y}; h) = \frac{1}{m} \sum_{k=0}^{m-1} \ell(h(x_k), \bar{y}_k)$, and $\bar{Y} = [\bar{y}_0, \bar{y}_1, \dots, \bar{y}_{m-1}]$ denotes the *abduced labels* that are consistent with the knowledge base B . The abduced labels \bar{Y} are inferred through *abduction*, a basic form of logical reasoning that seeks the most likely explanation for observations based on background knowledge (Peirce 1955; Simon and Newell 1971; Garcez et al. 2007). Often, there are multiple candidates for abduced labels (Dai et al. 2019), e.g., both $0 \wedge 1 = 0$ and $1 \wedge 0 = 0$ are correct equations. Hence, various heuristics have been proposed to guide the search for the most likely labels from the candidate set $\mathcal{S}(\tau) = \{\bar{Y} \in \mathcal{Y}^m \mid B \cup \bar{Y} \models \tau\}$. For example, Cai et al. (2021) constrained the Hamming distance between the abduced labels \bar{Y} and the predicted labels $f(X)$,¹ and Dai and Muggleton (2021) chose to pick the most probable labels based on the likelihood $p(\bar{Y}|X) = \prod_{k=0}^{m-1} p(\bar{y}_k|x_k)$, where $p(\bar{y}|x)$ is approximated by the output of softmax function, i.e., $\hat{p}(\bar{y}|x) = \exp(h_{\bar{y}}(x)) / \sum_{i=0}^{c-1} \exp(h_i(x))$. The overall pipeline of these algorithms is described in Appendix A.

3 Theoretical Analysis

Previous studies have showcased the practicality of neuro-symbolic learning systems—the objective of minimal inconsistency empirically yields classifiers adept at accurately predicting labels. In this section, we aim to disclose the ingredients of success from a theoretical perspective. We begin by considering a simple yet representative task. This motivates us to formulate a novel way of characterising supervision signals from a given knowledge base, and provide conditions under which the signals are sufficient for learning to succeed. Specifically, we show that the objective in Eq. (2) essentially addresses an upper bound of a location-based risk, whose minimisers are guaranteed to recover ground-truth labels when a rank criterion is satisfied.

Location Signals

Let us first consider the sequence $[1, \wedge, 0, \equiv, 0]$ in Fig. 1. It forms a correct equation only if the penultimate instance \equiv is labelled as “equal-sign”. This illustrates that the label of an instance could be determined by its location. However, challenges remain, e.g., the label of the last instance 0 is not determined by its location. To resolve this, our intuition is that the last instance is more likely to be 0 than 1, since there are four possible cases for a correct equation: $[0, \wedge, 0, =, 0]$,

¹With a slight abuse of notation, the predicted labels of raw inputs are denoted by $f(X) = [f(x_0), f(x_1), \dots, f(x_{m-1})]$.




Figure 3: Illustration of the location signals in Example 1. From the input sequences, we can observe instance-location pairs $\langle x, \iota \rangle$, while ground-truth labels y are unobservable. Intuitively, the instances of $y = 0$ are more likely to occur at the 2-th position, if candidate label probabilities are equal.

$[0, \wedge, 1, =, 0]$, $[1, \wedge, 0, =, 0]$, and $[1, \wedge, 1, =, 1]$. In 3/4 cases, the last instance is 0. This intuition is utilised as follows.

Now, consider the task in Example 1. For each input sequence, there are four possibilities of abduced labels including $[0, 0, 0]$, $[0, 1, 0]$, $[1, 0, 0]$, and $[1, 1, 1]$. This sequence-level information is crude, while our interest is still in instance-level prediction. To this end, we would like to extract instance-level supervision signals from the sequences. Specifically, in Example 1, each input sequence $[x_0, x_1, x_2]$ naturally yields three instance-location pairs $\{\langle x_i, \iota \rangle\}_{i=0}^2$, where ι denotes the location of an instance in the sequence. Fig. 3 illustrate the instance-location pairs. In general, given n sequences of unlabelled data $\{X^{(i)}\}_{i=0}^{n-1}$, each sequence $X^{(i)} = [x_0^{(i)}, x_1^{(i)}, \dots, x_{m-1}^{(i)}]$ naturally yields m instance-location pairs $\{\langle x_i^{(i)}, \iota \rangle\}_{i=0}^{m-1}$, and finally a total of mn instance-location pairs $\{\langle x_i^{(i)}, \iota^{(i)} \rangle\}_{i=0}^{mn-1}$ are obtained.

Apparently, such instance-level signals are insufficient for predicting true labels y . With a sample of the instance-location pairs, we could only learn a mapping $q: \mathcal{X} \rightarrow \mathbb{R}^m$ that estimates the underlying conditional probability $p(\iota|x)$. To address this issue, we choose to express $p(\iota|x)$ in terms of the desired conditional probability $p(y|x)$; that is, $\forall k \in [m]$,

$$p(\iota = k | x) = \sum_{j=0}^{c-1} Q_{jk} \cdot p(y = j | x), \quad (3)$$

where $Q_{jk} = p(\iota = k | y = j)$ denotes the probability of the class j occurring at the k -th position in a sequence. Here, we assume $p(\iota|y, x) = p(\iota|y)$. This means that ι is independent of x given y , analogous to the ‘‘missing completely at random’’ assumption that is often made in learning with missing values (Little and Rubin 1987; Elkan and Noto 2008).

In practice, we let $q(x) = Q^\top g(x)$ and interpret $g(x)$ as probabilities via $g_j(x) = \exp(h_j(x)) / \sum_{i=0}^{c-1} \exp(h_i(x))$, $\forall j \in [c]$. Consequently, if $q(x)$ learns to predict the probability $p(\iota|x)$, then $g(x)$ serves to estimate the probability $p(y|x)$. This can be achieved by minimising the following location-based risk (L-Risk) with appropriate loss functions:

$$\mathcal{R}_L(h) = \mathbb{E}_{p(x, \iota)} \ell(q(x), \iota). \quad (4)$$

While the above indicates that it is possible to decipher ground-truth labels, the reliance on the knowledge of label distribution is indispensable. Indeed, we find a frequently-used data generation process in previous practices (Cai et al. 2021; Dai and Muggleton 2021; Huang et al. 2021), where the labels of input sequences are implicitly assumed to be uniform over the candidate set $\mathcal{S}(\tau)$. This means that the

prior probabilities of candidate labels are equal. Under this assumption, the probability Q_{jk} in Eq. (3) can be computed as $\sum_{Y \in \mathcal{S}(\tau)} \mathbb{1}(y_k = j) / \sum_{Y \in \mathcal{S}(\tau)} \sum_{k=0}^{m-1} \mathbb{1}(y_k = j)$, where $\mathbb{1}(\cdot)$ is the indicator function. For the task in Example 1, the uniform assumption implies that the probabilities of candidate labels are $p([0, 0, 0]) = p([0, 1, 0]) = p([1, 0, 0]) = p([1, 1, 1]) = 1/4$, and that the elements in the probability matrix Q are $Q_{00} = 2/7, Q_{01} = 2/7, Q_{02} = 3/7$ and $Q_{10} = 2/5, Q_{11} = 2/5, Q_{12} = 1/5$.

In what follows, we will establish connections between the location-based risk and the objective of popular hybrid methods by utilising and relaxing the uniform assumption.

Upper Bound

We first show that under reasonable conditions, the objective of minimal inconsistency in Eq. (2) is an upper bound of the L-Risk in Eq. (4) up to an additive constant.

Assumption 1 (Uniform Assumption). $\forall Y \in \mathcal{S}(\tau)$, $p(Y) = 1/|\mathcal{S}(\tau)|$, where $\mathcal{S}(\tau) = \{Y \in \mathcal{Y}^m \mid B \cup Y \models \tau\}$.

Theorem 1. *Suppose that the uniform assumption holds and the concept space contains only one target concept τ . Let $a = \max_{i \in \mathcal{Y}} \{\sum_{Y \in \mathcal{S}(\tau)} \sum_{k=0}^{m-1} \mathbb{1}(y_k = i) / |\mathcal{S}(\tau)|\}$. For any classifier h and any knowledge base B , if the abduced labels \hat{Y} are randomly selected from $\mathcal{S}(\tau)$ and ℓ is the cross-entropy loss, then we have*

$$\mathcal{R}_L(h) \leq \mathcal{R}_{\text{NeSy}}(h) + C,$$

where $C = \log a \leq \log m$ is a constant.

The proof is given in Appendix B. Theorem 1 illustrates that minimising the inconsistency with a given knowledge base in hybrid systems is equivalent to minimising an upper bound of the location-based risk. In addition, the bound is tight: equality can be achieved, for example, when $\mathcal{S}(\tau) = \{[0, 1], [1, 0]\}$ and a uniformly random classifier h such that for any x and y , $\ell(h(x), y) = -\log \hat{p}(y|x) = \log 2$.

Remark. Firstly, it is common in practice to collect a set of input sequences belonging to the same target concept. For example, the word recognition task in Cai et al. (2021) contains only one target concept ‘‘valid_word’’, and the equation decipherment experiment in Huang et al. (2021) uses only correct equations. Secondly, although many heuristics have been proposed to select the most likely labels from the candidate set, they all behave like random guessing in the early stages of training when the classifier is randomly initialised (Dai and Muggleton 2021). Finally, the cross-entropy loss is commonly adopted in previous work (Dai et al. 2019; Cai et al. 2021; Huang et al. 2021).

Above, we have mainly focused on learning with a single target concept, which is a representative case due to its broad applications, and due to the insight it offers into the location-based signals. Next, we will generalise our discussion to utilise multiple target concepts. Evidently, these target concepts provide another source of supervision signals.

Example 2. Consider a binary classification task with label space $\mathcal{Y} = \{0, 1\}$ and concept space $\mathcal{T} = \{\text{conj0}, \text{conj1}\}$.

Facts	Knowledge Base
conj0([0,0]).	conj([Y0, Y1, Y2]) ← See Figure 2.
conj0([0,1]).	conj0([Y0, Y1]) ← conj([Y0, Y1, 0]).
conj0([1,0]).	conj1([Y0, Y1]) ← conj([Y0, Y1, 1]).
conj1([1,1]).	

Figure 4: Illustration of the knowledge base about the target concepts “conj0” and “conj1”, along with the corresponding facts abduced from the knowledge base.

Each sequence X consists of two raw inputs, whose ground-truth labels Y are unknown but satisfy the logical equations $y_0 \wedge y_1 = 0$ or $y_0 \wedge y_1 = 1$. Fig. 4 illustrates the logical facts abduced from the knowledge base in this task.

For each input sequence in Example 2, three possibilities of candidate labels can be abduced from the target concept “conj0”, including $[0, 0]$, $[0, 1]$ and $[1, 0]$, while the candidate labels become $[1, 1]$ if the observed concept is “conj1”. In order to simultaneously exploit this target-based information and the aforementioned location-based signals, we propose to construct a collection of instance-location-target triplets. Specifically, given n input sequences paired with target concepts $\{\langle X^{(i)}, \tau^{(i)} \rangle\}_{i=0}^{n-1}$, each sequence naturally yields m triplets $\{\langle x_i^{(i)}, \iota, \tau^{(i)} \rangle\}_{\iota=0}^{m-1}$, and finally a set of mn triplets $\{\langle x^{(i)}, \iota^{(i)}, \tau^{(i)} \rangle\}_{i=0}^{mn-1}$ are obtained.

In combination, the signals $\langle \iota^{(i)}, \tau^{(i)} \rangle$ in the triplets have $\tilde{c} = m \cdot |\mathcal{T}|$ distinct values. For conciseness, we use synthetic labels $\tilde{y} \in [\tilde{c}]$ to denote the signals $\langle \iota, \tau \rangle \in [m] \times \mathcal{T}$. Then, the triplet set is represented as a set of instance-label pairs $\{\langle x^{(i)}, \tilde{y}^{(i)} \rangle\}_{i=0}^{mn-1}$, from which we could learn a mapping $\tilde{q}: \mathcal{X} \rightarrow \mathbb{R}^{\tilde{c}}$ that estimates the underlying conditional probability $p(\tilde{y}|x)$, i.e., $p(\tau, \iota|x)$. Similar to before, we express $p(\tilde{y}|x)$ in terms of $p(y|x)$ for predicting ground-truth labels; that is, $\forall o \in [\tilde{c}]$,

$$p(\tilde{y} = o | x) = \sum_{j=0}^{c-1} \tilde{Q}_{jo} \cdot p(y = j | x), \quad (5)$$

where $\tilde{Q}_{jo} = p(\tilde{y} = o | y = j)$. Without loss of generality, let $t \in \{0, \dots, |\mathcal{T}| - 1\}$ and $o = tm + k$. Then, \tilde{Q}_{jo} represents $p(\iota = k, \tau = t | y = j)$, which is the probability of the class j occurring at the k -th position in a sequence of concept t . In practice, we let $\tilde{q}(x) = \tilde{Q}^\top g(x)$. If $\tilde{q}(x)$ learns to predict the probability $p(\tilde{y}|x)$, then $g(x)$ serves to estimate the probability $p(y|x)$. This can be achieved by minimising the following target-location-based risk (TL-Risk):

$$\mathcal{R}_{\text{TL}}(h) = \mathbb{E}_{p(x, \tilde{y})} \ell(\tilde{q}(x), \tilde{y}). \quad (6)$$

Now we are ready to state the upper bound in the case of multiple target concepts.

Theorem 2. *Given a data distribution $p(X, \tau)$ with any label distribution $p(Y)$. Let $b = \min_{\tau \in \mathcal{T}} p(\tau)$. For any classifier h and any knowledge base B , if the abduced labels \tilde{Y} are randomly selected from $\mathcal{S}(\tau)$ and ℓ is the cross-entropy loss, then we have*

$$\mathcal{R}_{\text{TL}}(h) \leq \mathcal{R}_{\text{NeSy}}(h) + C,$$

where $C = \log m - \log b$ is a constant.

The proof is given in Appendix B. Theorem 2 illustrates that minimising the inconsistency in hybrid systems is equivalent to minimising an upper bound of the target-location-based risk. Also, Theorem 2 implies that the upper bound holds independent of the uniform assumption. However, to render Eq. (6) practical, a premise for obtaining the value of the probability matrix \tilde{Q} is still required.

To this end, we resort to a realistic data generation process described as follows. Initially, instances in a sequence are generated independently from $p(x, y)$. This sequence is subsequently submitted to a labelling oracle, which generates a target concept by identifying a ground atom in its knowledge base that corresponds to the instances’ labels. This procedure is repeated, eventually yielding a collection of sequences paired with their target concepts. With this data generation process, it is straightforward to derive the sequence-level label density $p(Y)$ from the instance-level label density $p(y)$. More specifically, the label density for a sequence is given as the product of the individual instance-level densities., i.e., $p(Y = [\bar{y}_1, \dots, \bar{y}_m]) = \prod_{k=0}^{m-1} p(y = \bar{y}_k)$. Then, by following the derivation of generality in Muggleton (2023), we obtain $p(Y|\tau) = p(\tau|Y)p(Y)/p(\tau)$, where $p(\tau = t|Y = \bar{Y})$ equals one if $B \cup \bar{Y} \models t$, otherwise it equals zero. Finally, the probability \tilde{Q}_{jo} can be computed as $p(\tau, \iota|y) = p(y|\tau, \iota)p(\tau)p(\iota)/p(y)$, where $p(y = j|\tau = t, \iota = k) = \sum_{\bar{Y} \in \mathcal{S}(t)} \mathbb{1}(\bar{y}_k = j)p(Y = \bar{Y}|\tau = t)$.

Remark. The data generation process described above has been adopted in previous work, such as the “addition” task in Manhaeve et al. (2018) and the “member” task in Tsamoura, Hospedales, and Michael (2021). We also note that when the prior distribution on y is uniform, the label density on Y given a target concept τ derived from the above process is also uniform over the candidate set. Thus, this data generation process favourably relaxes the uniform assumption.

Rank Criterion

Equipped with the probability matrix, now we are ready to present a rank criterion that indicates the knowledge’s efficacy in recovering the ground-truth labels of raw inputs.

Theorem 3. *If the probability matrix \tilde{Q} has full row rank and the cross-entropy loss is used, then the minimiser $h_{\text{TL}}^* = \arg \min_h \mathcal{R}_{\text{TL}}(h)$ recovers the true minimiser $h^* = \arg \min_h \mathcal{R}(h)$, i.e., $h_{\text{TL}}^* = h^*$.*

The proof is given in Appendix B. Theorem 3 means that the true values of the labels of raw inputs can be reliably recovered if the probability matrix has full row rank.

On one hand, Theorems 1 and 2 demonstrate that the objective of minimal inconsistency contributes to mitigating location-based risks, in which the supervision signals implied by a knowledge base are explicitly characterised into a probability matrix. On the other hand, Theorem 3 reveals that the minimiser of the TL-Risk can recover the true labels, provided that the probability matrix has full row rank. In short, these findings suggest that the objective of minimal inconsistency can succeed by mitigating the risk of predicting locations, with the rank criterion serving as an indicator of the knowledge’s efficacy in facilitating learning.

We conclude this part by illustrating the use of the rank criterion for learning with one or more target concepts.

Corollary 1. *If the probability matrix Q has full row rank and the cross-entropy loss is used, then the minimiser $h_L^* = \arg \min_h \mathcal{R}_L(h)$ recovers the true minimiser $h^* = \arg \min_h \mathcal{R}(h)$, i.e., $h_L^* = h^*$.*

Corollary 1 applies directly to the task of learning with the single target concept “conj” in Example 1. As discussed before, the candidate labels in this task under the uniform assumption lead to the following probability matrix

$$Q = \begin{pmatrix} 2/7 & 2/7 & 3/7 \\ 2/5 & 2/5 & 1/5 \end{pmatrix},$$

which has full row rank. Therefore, according to Corollary 1, the supervision signals from the knowledge base are sufficient to produce accurate classifiers.

Similarly, let us consider the case of learning with only the target concept “conj0” in Example 2. This can lead to the following probability matrix

$$Q = \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix},$$

which has rank one. Thus, it may fail to recover ground-truth labels in this case. Fortunately, there is another target concept “conj1” in Example 2. Leveraging both target concepts, we can derive a different probability matrix

$$\tilde{Q} = \begin{pmatrix} 1/2 & 1/2 & 0 & 0 \\ 1/4 & 1/4 & 1/4 & 1/4 \end{pmatrix}.$$

This matrix has full row rank, thus facilitating the learning of accurate classifiers according to Theorem 3.

4 Experiments

In this section, we conduct comprehensive experiments to validate the utility of the proposed criterion on various tasks. The code is available for download.²

Tasks. We first examine four benchmark tasks: `ConjEq`, `HED`, `Conjunction`, and `Addition`. The `ConjEq` task is a variant of the `HED` (i.e., handwritten equation decipherment) task adopted in Dai et al. (2019); Huang et al. (2021). The knowledge base in this task has been illustrated in Fig. 2, and it accepts triplets of handwritten Boolean symbols as inputs. In contrast, the original `HED` task exploits the knowledge of the correctness of binary additive equations and accepts the handwritten equations composed of digits, the plus sign, and the equal sign as inputs (Dai et al. 2019). Similarly, the `Conjunction` task is a variant of the `Addition` task introduced in Manhaeve et al. (2018). The knowledge base in this task has been provided in Fig. 4, and it accepts pairs of handwritten Boolean symbols as inputs. In contrast, the `Addition` task works with a knowledge base that defines the sum of two summands, and accepts pairs of handwritten decimal digits as inputs. Following previous work (Huang et al. 2021; Cai et al. 2021), we collect training sequences

for the tasks by representing the handwritten symbols using instances from benchmark datasets including MNIST (LeCun et al. 1998), EMNIST (Cohen et al. 2017), USPS (Hull 1994), KUZUSHIJI (Clanuwat et al. 2018), and FASHION (Xiao, Rasul, and Vollgraf 2017). All experiments are repeated six times on GeForce RTX 3090 GPUs, with the mean accuracy and standard deviation reported. More details on experimental settings are found in Appendix C.

Methods. We consider four strategies for selecting the abducted labels from the candidate set in hybrid learning systems: `RAND` (Random), `MAXP` (Maximal Probability), `MIND` (Minimal Distance), and `AVG` (Average). Specifically, `RAND` selects a consistent label sequence randomly from the candidate set for each input sequence. `MAXP` identifies the most probable labels from the candidate set, based on the likelihood $p(\bar{Y}|X)$ estimated by the classifier. This strategy aligns with common practices in previous studies (Li et al. 2020; Dai and Muggleton 2021; Huang et al. 2021). `MIND` chooses the abducted labels that have the smallest Hamming distance to the predicted labels. This strategy also mirrors approaches found in earlier research (Dai et al. 2019; Tsamoura, Hospedales, and Michael 2021; Cai et al. 2021). `AVG` regards all label sequences in the candidate set as plausible labels. It calculates a loss for each label sequence and then averages them. Thus, in expectation, this strategy is equivalent to the `RAND` strategy. It is worth noting that the `MAXP` and `MIND` strategies behave similarly to the `RAND` strategy in the initial stages of training, as both the estimated probabilities and the predicted labels are essentially random when the classifier is initialised randomly. After the above label abduction procedure, the empirical counterpart of Eq. (2) is used as the learning objective. The above methods are compared with `TL`, i.e., the method of minimising the empirical counterpart of the risk in Eq. (6).

Experimental Results on `ConjEq` and `Conjunction`. As indicated by Fig. 2 and Fig. 4, both tasks aim to learn a binary classifier that perceives raw inputs and assigns them a prediction of either 0 or 1. While the two tasks differ in terms of the number of target concepts (one versus two) and the length of input sequences (three versus two), they both satisfy the rank criterion. Thus, it is expected that neuro-symbolic learning can produce accurate classifiers in these cases. Indeed, this is confirmed by our experimental results. Table 1 presents the test performance of multi-layer perception (MLP) produced by hybrid learning methods on various datasets for the `ConjEq` and `Conjunction` tasks. Results indicate that all four strategies for selecting abducted labels work well in facilitating successful learning: they all achieve over 98% accuracy, and none appear to have a significant edge over the others. This implies that the `MAXP` and `MIND` strategies are empirically similar to the `RAND` strategy in the `ConjEq` and `Conjunction` tasks. The phenomenon can be explained by the small candidate set size (≤ 4) for label abduction in both tasks, which leads to a considerable probability of correctly selecting the labels even when employing the `RAND` strategy. We also observe that `TL` consistently achieves competitive performance across all datasets in both tasks, which corroborates our theoretical analysis.

²<https://github.com/AbductiveLearning/ABL-TL>

TASK	METHOD	MNIST	EMNIST	USPS	KUZUSHIJI	FASHION
ConjEq	RAND	99.91 ± 0.06	99.65 ± 0.04	99.33 ± 0.16	97.82 ± 0.35	98.40 ± 0.11
	MAXP	99.94 ± 0.04	99.82 ± 0.03	99.20 ± 0.00	98.80 ± 0.16	99.39 ± 0.12
	MIND	99.91 ± 0.08	99.84 ± 0.07	99.14 ± 0.17	98.91 ± 0.17	98.84 ± 0.19
	AVG	99.85 ± 0.10	99.80 ± 0.07	99.30 ± 0.17	98.34 ± 0.16	98.62 ± 0.21
	TL	99.92 ± 0.05	99.82 ± 0.06	99.25 ± 0.08	98.53 ± 0.26	98.77 ± 0.06
Conjunction	RAND	99.91 ± 0.06	99.86 ± 0.04	99.30 ± 0.13	98.79 ± 0.13	99.00 ± 0.27
	MAXP	99.93 ± 0.04	99.81 ± 0.02	99.20 ± 0.00	98.62 ± 0.15	99.05 ± 0.09
	MIND	99.94 ± 0.02	99.79 ± 0.02	99.20 ± 0.00	98.74 ± 0.10	99.08 ± 0.10
	AVG	99.94 ± 0.02	99.85 ± 0.03	99.30 ± 0.13	98.68 ± 0.33	99.23 ± 0.13
	TL	99.94 ± 0.02	99.83 ± 0.04	99.20 ± 0.00	98.87 ± 0.18	99.30 ± 0.04
Addition	RAND	92.01 ± 0.93	92.94 ± 1.45	90.96 ± 1.04	73.18 ± 0.71	79.08 ± 2.61
	MAXP	96.40 ± 4.04	95.09 ± 5.20	94.29 ± 0.27	90.00 ± 0.27	87.34 ± 2.93
	MIND	98.32 ± 0.04	98.61 ± 0.06	94.61 ± 0.17	90.85 ± 0.26	88.40 ± 0.62
	AVG	94.90 ± 0.39	95.71 ± 0.42	93.22 ± 0.30	80.94 ± 0.62	84.43 ± 0.92
	TL	98.00 ± 0.14	98.41 ± 0.05	94.68 ± 0.20	90.04 ± 0.32	88.38 ± 0.25
HED	RAND	99.89 ± 0.02	99.71 ± 0.12	99.25 ± 0.23	97.68 ± 0.70	98.43 ± 0.55
	MAXP	99.90 ± 0.02	99.77 ± 0.02	99.23 ± 0.05	98.55 ± 0.08	99.33 ± 0.10
	MIND	99.87 ± 0.07	99.77 ± 0.02	99.21 ± 0.00	98.61 ± 0.21	99.32 ± 0.10
	AVG	99.60 ± 0.09	99.38 ± 0.21	99.32 ± 0.14	96.16 ± 1.22	98.46 ± 0.33
	TL	99.90 ± 0.02	99.77 ± 0.04	99.21 ± 0.00	98.50 ± 0.16	99.21 ± 0.06

Table 1: Test accuracy (%) of each method using MLP on benchmark datasets and tasks.

Experimental Results on Addition. This task is more challenging than Conjunction, as the size of the candidate set in Addition is larger, which complicates the abduction of correct labels. This difficulty is amplified by the presence of 10 classes in Addition, compared to just 2 classes in Conjunction. Despite these challenges, our rank criterion positively indicates that the supervision signals from the knowledge base of Addition are sufficient for learning accurate classifiers. This is confirmed by Table 1, showing that all methods can learn classifiers with significantly greater accuracy than random guessing in the Addition task. An insightful observation is that the MAXP and MIND strategies consistently outperform the RAND strategy in this task. This implies that, despite their similar behaviours in the early stages of training, these strategies may diverge in later training stages. As the accuracy of the estimated probabilities and the predicted labels improves, MAXP and MIND are increasingly likely to select the correct labels over RAND. Finally, we observe that TL outperforms RAND and AVG, while showing competitive results with MAXP and MIND.

Experimental Results on HED. Results on HED using the knowledge base of binary additive equations can be found in Table 1. Again, all methods consistently perform well across all datasets for HED. This success is also indicated by our rank criterion. The knowledge base of binary additive equations satisfies the criterion: the rank of the corresponding probability matrix is 4, which equals the number of symbols in binary additive equations (i.e., “0”, “1”, “+”, and “=”). However, when the task is extended to handle additive equations with decimal digits, the corresponding matrix rank becomes 7, which is less than 12, the number of symbols in decimal additive equations (i.e., “0”, ..., “9”, “+”, and “=”). In other words, the knowledge base of decimal ad-




Figure 5: Performance on HED using different knowledge bases of numeral systems ranging from base 2 to base 10.

ditive equations does not satisfy the criterion. This indicates that learning would fail in this case.

Our experiments validate this, showing that TL performs poorly when using the knowledge base of decimal additive equations, with test accuracy falling below 50%. Thorough experiments involving additive equations across number systems from base 2 to base 10 further confirm the utility of the proposed criterion, as shown by the performance drop when the base is larger than 5. Finally, we note that, although the rank criterion is negative in the case of base 10 for HED, Huang et al. (2021) have shown that effective learning is still possible with additional assumptions. Concretely, they assumed that there exists a similarity between raw inputs. While this assumption may not always hold in practice, it is interesting to theoretically analyse its usefulness in helping neuro-symbolic learning. We leave this as future work.

Experimental Results on Random Knowledge Bases. To further demonstrate the utility of our rank criterion, we conduct experiments on random knowledge bases. Following Cai et al. (2021), we create knowledge bases by randomly generating binary rules in disjunctive normal form




Figure 6: Comparison of knowledge bases satisfying or not satisfying the rank criterion. Knowledge bases are created by randomly generating rules in disjunctive or conjunctive norm forms, with clause lengths varying from 3 to 5. The rank criterion effectively indicates the success of learning accurate classifiers.

(DNF), i.e., disjunctions of conjunctive clauses. In DNF, we control the clause length m , and the number of clauses is randomly chosen from the range $[1, 2^{m-1}]$. Most generated knowledge bases satisfy the rank criterion, while there are some knowledge bases whose corresponding matrix rank is deficient. Examples of these knowledge bases are illustrated in Appendix E. Numerical results are summarised in Fig. 6. We observe that our rank criterion consistently indicates the learning performance of all methods in all cases. When the probability matrix corresponding to a knowledge base has full row rank, the learned classifiers mostly achieve high test accuracy; conversely, rank deficiency indicates low performance, close to random guessing. While a few outliers exist, such as the rank-deficient case in Fig. 6(b), where the MAXP and MIND strategies perform slightly better than random guessing, their results remain significantly lower than the full-rank case. We also notice that the MAXP and MIND strategies perform slightly worse in Fig. 6(c), potentially due to the size of the candidate set being large in those cases, which may complicate the abduction of correct labels. Nevertheless, we observe that RAND, AVG, and TL consistently perform well across all cases.

5 Related Work

Neuro-Symbolic Learning. Neuro-symbolic integration had been studied decades ago (Towell and Shavlik 1994; Garcez, Broda, and Gabbay 2002). In recent years, the increasing efficacy of modern machine learning techniques, such as deep learning, sparked a surge of interest in integrating machine learning and symbolic reasoning. Many efforts have been devoted to this integration (Dai and Zhou 2017; Donadello, Serafini, and Garcez 2017; Gaunt et al. 2017; Santoro et al. 2017; Grover et al. 2018; Xu et al. 2018; Trask et al. 2018; Manhaeve et al. 2018; Dai et al. 2019; Dong et al. 2019; Wang et al. 2019; Cohen, Yang, and Mazaitis 2020; Yang, Ishay, and Lee 2021; Evans et al. 2021; Li et al. 2023; Liu et al. 2023; Huang et al. 2023a,b), with abductive learning (ABL) emerging as one of the most expressive frameworks for hybrid systems (Zhou 2019; Zhou and Huang 2022). It is noteworthy that, beyond the basic settings in this paper, the ABL framework is highly general and flexible, accommodating various machine learning mechanisms and allowing the exploitation of abundant labelled data and

inaccurate knowledge bases (Zhou 2019). ABL has found applications in many areas, such as theft judicial sentencing (Huang et al. 2020), stroke evaluation (Wang et al. 2021), and optical character recognition (Cai et al. 2021).

Weakly Supervised Learning. The ABL framework has been viewed as enlarging the scope of weakly supervised learning (WSL) (Zhou 2018), where the supervision information can come from knowledge reasoning (Zhou and Huang 2022). Yet, traditional WSL settings, such as multi-instance learning (Dietterich, Lathrop, and Lozano-Pérez 1997; Zhou, Sun, and Li 2009) and partial-label learning (Jin and Ghahramani 2002; Cour, Sapp, and Taskar 2011; Liu and Dietterich 2014; Feng et al. 2020), struggle to tackle the challenges in ABL. The consistency of learning from aggregate information has been explored in Zhang et al. (2020) and a concurrent work by Wang, Tsamoura, and Roth (2023), but they can only handle limited types of aggregate functions and loss functions; in contrast, our analysis applies to any type of knowledge bases under the ABL framework. While the idea of summarising probabilities into a stochastic matrix has been widely exploited in Markov chains (Gagniuć 2017) and noisy-label learning (Natarajan et al. 2013; Patrini et al. 2017; Yu et al. 2018), we are the first to show its relation with the objective of hybrid learning systems. This work is also closely related to the concept of reasoning shortcuts in Marconato et al. (2023a,b), where perception models may learn unintended semantics. Our rank criterion is useful in indicating such failures before actual training.

6 Conclusion

In this work, we introduce a novel characterisation of the supervision signals from a given knowledge base and establish a rank criterion capable of indicating the practical efficacy of a given knowledge base in improving learning performance. Both theoretical and empirical results shed light on the success of hybrid learning systems while pinpointing potential failures when the supervision signals from a symbolic knowledge base are insufficient to ensure effective learning. Future work includes the detailed analysis of mutual promotion between learning and reasoning, the incorporation of other machine learning models, and the exploitation of abundant labelled data and inaccurate knowledge bases.

Acknowledgments

This research was supported by the NSFC (62176117, 62206124) and JiangsuSF (BK20232003).

References

- Cai, L.-W.; Dai, W.-Z.; Huang, Y.-X.; Li, Y.-F.; Muggleton, S. H.; and Jiang, Y. 2021. Abductive Learning with Ground Knowledge Base. In *IJCAI*, 1815–1821.
- Clanuwat, T.; Bober-Irizar, M.; Kitamoto, A.; Lamb, A.; Yamamoto, K.; and Ha, D. 2018. Deep learning for classical Japanese literature. *arXiv preprint arXiv:1812.01718*.
- Cohen, G.; Afshar, S.; Tapson, J.; and Van Schaik, A. 2017. EMNIST: Extending MNIST to handwritten letters. In *IJCNN*, 2921–2926.
- Cohen, W.; Yang, F.; and Mazaitis, K. R. 2020. Tensorlog: A probabilistic database implemented using deep-learning infrastructure. *Journal of Artificial Intelligence Research*, 67: 285–325.
- Cour, T.; Sapp, B.; and Taskar, B. 2011. Learning from partial labels. *Journal of Machine Learning Research*, 12: 1501–1536.
- Dai, W.-Z.; and Muggleton, S. H. 2021. Abductive knowledge induction from raw data. In *IJCAI*, 1845–1851.
- Dai, W.-Z.; Xu, Q.; Yu, Y.; and Zhou, Z.-H. 2019. Bridging machine learning and logical reasoning by abductive learning. In *NeurIPS*, 2811–2822.
- Dai, W.-Z.; and Zhou, Z.-H. 2017. Combining logical abduction and statistical induction: Discovering written primitives with human knowledge. In *AAAI*, 4392–4398.
- De Raedt, L.; Dumančić, S.; Manhaeve, R.; and Marra, G. 2021. From statistical relational to neural-symbolic artificial intelligence. In *IJCAI*, 4943–4950.
- De Raedt, L.; Kersting, K.; Natarajan, S.; and Poole, D. 2016. Statistical relational artificial intelligence: Logic, probability, and computation. *Synthesis lectures on artificial intelligence and machine learning*, 10(2): 1–189.
- De Raedt, L.; and Kimmig, A. 2015. Probabilistic (logic) programming concepts. *Machine Learning*, 100: 5–47.
- Dietterich, T. G.; Lathrop, R. H.; and Lozano-Pérez, T. 1997. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2): 31–71.
- Donadello, I.; Serafini, L.; and Garcez, A. D. 2017. Logic tensor networks for semantic image interpretation. In *IJCAI*.
- Dong, H.; Mao, J.; Lin, T.; Wang, C.; Li, L.; and Zhou, D. 2019. Neural Logic Machines. In *ICLR*.
- Elkan, C.; and Noto, K. 2008. Learning classifiers from only positive and unlabeled data. In *KDD*, 213–220.
- Evans, R.; Bošnjak, M.; Buesing, L.; Ellis, K.; Pfau, D.; Kohli, P.; and Sergot, M. 2021. Making sense of raw input. *Artificial Intelligence*, 299: 103521.
- Feng, L.; Lv, J.; Han, B.; Xu, M.; Niu, G.; Geng, X.; An, B.; and Sugiyama, M. 2020. Provably consistent partial-label learning. In *NeurIPS*, 10948–10960.
- Gagniuc, P. A. 2017. *Markov chains: from theory to implementation and experimentation*. John Wiley & Sons.
- Garcez, A. S. d.; Broda, K.; and Gabbay, D. M. 2002. *Neural-symbolic learning systems: foundations and applications*. Springer Science & Business Media.
- Garcez, A. S. d.; Gabbay, D. M.; Ray, O.; and Woods, J. 2007. Abductive reasoning in neural-symbolic systems. *Topoi*, 26: 37–49.
- Gaunt, A. L.; Brockschmidt, M.; Kushman, N.; and Tarlow, D. 2017. Differentiable programs with neural libraries. In *ICML*, 1213–1222.
- Getoor, L.; and Taskar, B. 2007. *Introduction to statistical relational learning*. MIT press.
- Grover, A.; Wang, E.; Zweig, A.; and Ermon, S. 2018. Stochastic Optimization of Sorting Networks via Continuous Relaxations. In *ICLR*.
- Hitzler, P.; and Sarker, M. K. 2022. *Neuro-Symbolic Artificial Intelligence - The State of the Art*. Frontiers in Artificial Intelligence and Applications. IOS Press.
- Huang, Y.-X.; Dai, W.-Z.; Cai, L.-W.; Muggleton, S. H.; and Jiang, Y. 2021. Fast abductive learning by similarity-based consistency optimization. In *NeurIPS*, 26574–26584.
- Huang, Y.-X.; Dai, W.-Z.; Jiang, Y.; and Zhou, Z.-H. 2023a. Enabling Knowledge Refinement upon New Concepts in Abductive Learning. In *AAAI*, 7928–7935.
- Huang, Y.-X.; Dai, W.-Z.; Yang, J.; Cai, L.-W.; Cheng, S.; Huang, R.; Li, Y.-F.; and Zhou, Z.-H. 2020. Semi-supervised abductive learning and its application to theft judicial sentencing. In *ICDM*, 1070–1075.
- Huang, Y.-X.; Sun, Z.; Li, G.; Tian, X.; Dai, W.-Z.; Hu, W.; Jiang, Y.; and Zhou, Z.-H. 2023b. Enabling abductive learning to exploit knowledge graph. In *IJCAI*, 3839–3847.
- Hull, J. J. 1994. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5): 550–554.
- Jin, R.; and Ghahramani, Z. 2002. Learning with multiple labels. In *NeurIPS*, 897–904.
- Kakas, A. C.; Kowalski, R. A.; and Toni, F. 1992. Abductive logic programming. *Journal of logic and computation*, 2(6): 719–770.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Li, Q.; Huang, S.; Hong, Y.; Chen, Y.; Wu, Y. N.; and Zhu, S.-C. 2020. Closed loop neural-symbolic learning via integrating neural perception, grammar parsing, and symbolic reasoning. In *ICML*, 5884–5894.
- Li, Z.; Yao, Y.; Chen, T.; Xu, J.; Cao, C.; Ma, X.; Jian, L.; et al. 2023. Softened Symbol Grounding for Neuro-symbolic Systems. In *ICLR*.
- Little, R.; and Rubin, D. 1987. *Statistical Analysis With Missing Data*. Wiley Series in Probability and Statistics. Wiley. ISBN 9780471802549.
- Liu, A.; Xu, H.; Van den Broeck, G.; and Liang, Y. 2023. Out-of-Distribution Generalization by Neural-Symbolic Joint Training. In *AAAI*, 12252–12259.

- Liu, L.; and Dietterich, T. 2014. Learnability of the superset label learning problem. In *ICML*, 1629–1637.
- Manhaeve, R.; Dumancic, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2018. Deepproblog: Neural probabilistic logic programming. In *NeurIPS*, 3753–3763.
- Marconato, E.; Bontempo, G.; Ficarra, E.; Calderara, S.; Passerini, A.; and Teso, S. 2023a. Neuro Symbolic Continual Learning: Knowledge, Reasoning Shortcuts and Concept Rehearsal. In *ICML*, 23915–23936.
- Marconato, E.; Teso, S.; Vergari, A.; and Passerini, A. 2023b. Not All Neuro-Symbolic Concepts Are Created Equal: Analysis and Mitigation of Reasoning Shortcuts. In *NeurIPS*.
- Muggleton, S. H. 2023. Hypothesizing an algorithm from one example: the role of specificity. *Philosophical Transactions of the Royal Society A*, 381(2251): 20220046.
- Natarajan, N.; Dhillon, I. S.; Ravikumar, P. K.; and Tewari, A. 2013. Learning with noisy labels. In *NeurIPS*, 1196–1204.
- Patrini, G.; Rozza, A.; Krishna Menon, A.; Nock, R.; and Qu, L. 2017. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, 1944–1952.
- Peirce, C. S. 1955. Abduction and induction. *Philosophical Writings of Pierce*, 150–56.
- Russell, S. 2015. Unifying logic and probability. *Communications of the ACM*, 58(7): 88–97.
- Santoro, A.; Raposo, D.; Barrett, D. G.; Malinowski, M.; Pascanu, R.; Battaglia, P.; and Lillicrap, T. 2017. A simple neural network module for relational reasoning. In *NeurIPS*, 4967–4976.
- Simon, H. A.; and Newell, A. 1971. Human problem solving: The state of the theory in 1970. *American psychologist*, 26(2): 145.
- Towell, G. G.; and Shavlik, J. W. 1994. Knowledge-based artificial neural networks. *Artificial intelligence*, 70(1-2): 119–165.
- Trask, A.; Hill, F.; Reed, S. E.; Rae, J.; Dyer, C.; and Blunsom, P. 2018. Neural arithmetic logic units. In *NeurIPS*, 8046–8055.
- Tsamoura, E.; Hospedales, T.; and Michael, L. 2021. Neural-symbolic integration: A compositional perspective. In *AAAI*, 5051–5060.
- Wang, J.; Deng, D.; Xie, X.; Shu, X.; Huang, Y.-X.; Cai, L.-W.; Zhang, H.; Zhang, M.-L.; Zhou, Z.-H.; and Wu, Y. 2021. Tac-valuer: Knowledge-based stroke evaluation in table tennis. In *KDD*, 3688–3696.
- Wang, K.; Tsamoura, E.; and Roth, D. 2023. On Learning Latent Models with Multi-Instance Weak Supervision. In *NeurIPS*.
- Wang, P.-W.; Donti, P.; Wilder, B.; and Kolter, Z. 2019. Satnet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *ICML*, 6545–6554.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Xu, J.; Zhang, Z.; Friedman, T.; Liang, Y.; and Broeck, G. 2018. A semantic loss function for deep learning with symbolic knowledge. In *ICML*, 5502–5511.
- Yang, Z.; Ishay, A.; and Lee, J. 2021. NeurASP: embracing neural networks into answer set programming. In *IJCAI*, 1755–1762.
- Yu, X.; Liu, T.; Gong, M.; and Tao, D. 2018. Learning with biased complementary labels. In *ECCV*, 68–83.
- Zhang, Y.; Charoenphakdee, N.; Wu, Z.; and Sugiyama, M. 2020. Learning from aggregate observations. In *NeurIPS*, 7993–8005.
- Zhou, Z. 2019. Abductive learning: towards bridging machine learning and logical reasoning. *Science China Information Sciences*, 62(7): 76101:1–76101:3.
- Zhou, Z.-H. 2018. A brief introduction to weakly supervised learning. *National science review*, 5(1): 44–53.
- Zhou, Z.-H.; and Huang, Y.-X. 2022. Abductive Learning. In *Neuro-Symbolic Artificial Intelligence: The State of the Art*, 353–369. IOS Press.
- Zhou, Z.-H.; Sun, Y.-Y.; and Li, Y.-F. 2009. Multi-instance learning by treating instances as non-iid samples. In *ICML*, 1249–1256.