

Efficient Rectification of Neuro-Symbolic Reasoning Inconsistencies by Abductive Reflection

Wen-Chao Hu^{1,2}, Wang-Zhou Dai^{1,3}, Yuan Jiang^{1,2}, Zhi-Hua Zhou^{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

{huwc, daiwz, jiangy, zhouzh}@lamda.nju.edu.cn

Abstract

Neuro-Symbolic (NeSy) AI could be regarded as an analogy to human dual-process cognition, modeling the intuitive System 1 with neural networks and the algorithmic System 2 with symbolic reasoning. However, for complex learning targets, NeSy systems often generate outputs inconsistent with domain knowledge and it is challenging to rectify them. Inspired by the human Cognitive Reflection, which promptly detects errors in our intuitive response and revises them by invoking the System 2 reasoning, we propose to improve NeSy systems by introducing Abductive Reflection (ABL-Refl) based on the Abductive Learning (ABL) framework. ABL-Refl leverages domain knowledge to abduce a reflection vector during training, which can then flag potential errors in the neural network outputs and invoke abduction to rectify them and generate consistent outputs during inference. ABL-Refl is highly efficient in contrast to previous ABL implementations. Experiments show that ABL-Refl outperforms state-of-the-art NeSy methods, achieving excellent accuracy with fewer training resources and enhanced efficiency.

1 Introduction

Human decision-making is generally recognized as an interaction between two systems: System 1 quickly generates an intuitive response, and System 2 engages in further algorithmic and slow reasoning (Frederick 2005; Kahneman 2011). In Neuro-Symbolic (NeSy) Artificial Intelligence (AI), neural networks often resemble System 1 for rapid pattern recognition, and symbolic reasoning mirrors System 2 to leverage domain knowledge and handle complex problems thoughtfully, yet in a slower and more controlled way (Bengio 2019). Like human System 1 reasoning, when facing complicated tasks, neural networks often produce unreliable outputs which cause inconsistencies with domain knowledge. These inconsistencies can then be reconciled with the help of the symbolic reasoning counterpart (Hitzler 2022).

To achieve the above process, some methods relax symbolic domain knowledge as neural network constraints (Xu et al. 2018; Yang, Lee, and Park 2022), some attempt to approximate logical calculus using distributed representations within neural networks (Wang et al. 2019). However, a loss

of full symbolic reasoning ability often occurs during these relaxation or approximation, hampering the ability of generating reliable output.

Abductive Learning (ABL) (Zhou 2019; Zhou and Huang 2022) is a framework for bridging machine learning and logical reasoning while preserving full expressive power in each side. In ABL, the machine learning component first converts raw data into primitive symbolic outputs. These outputs can be utilized by the symbolic reasoning component, which leverages domain knowledge and performs abduction to generate a revised, more reliable output. However, previous implementations of ABL require a highly discrete combinatorial consistency optimization before applying abduction, and this optimization has high complexity which encumbers, thereby severely limiting the efficiency and applicability to large-scale scenarios.

Human reasoning naturally exploits both sides efficiently, a hypothetical model for this process is called Cognitive Reflection, where the fast System 1 thinking is called to quickly generate an approximate over-all solution, and then seamlessly hands the complicated parts to System 2 (Frederick 2005). The key to this process is the reflection mechanism, which promptly detects which part in the intuitive response may contain inconsistencies with domain knowledge and invokes System 2 to rectify them. This reflection typically positively associates with System 2 capabilities, as both are closely linked to an individual’s mastery of domain knowledge (Sinayev and Peters 2015). Following the reflection, the process of the step-by-step formal reasoning becomes less complex: With a largely reduced search space, deriving the correct solution for System 2 becomes straightforward.

Inspired by this phenomenon, we propose a general enhancement, *Abductive Reflection (ABL-Refl)*. Based on ABL framework, ABL-Refl preserves full expressive power of neural networks and symbolic reasoning, while replacing the time-consuming consistency optimization with the reflection mechanism, thereby significantly improves efficiency and applicability. Specifically, in ABL-Refl, a reflection vector is concurrently generated with the neural network intuitive output, which flags potential errors in the output and invokes symbolic reasoning to perform abduction, thereby rectifying these errors and generating a new output that is more consistent with domain knowledge. During model training, the training information for the reflection derives from domain

knowledge. In essence, the reflection vector is abducted from domain knowledge and serves as an attention mechanism for narrowing the problem space of symbolic reasoning. The reflection can be trained unsupervisedly, requiring only the same amount of domain knowledge as state-of-the-art NeSy systems without generating extra training data.

We validate the effectiveness of ABL-Refl in solving Sudoku NeSy benchmarks in both symbolic and visual forms. Compared to previous NeSy methods, ABL-Refl performs significantly better, achieving higher reasoning accuracy efficiently with fewer training resources. We also compare our method to symbolic solvers, and show that the reduced search space in ABL-Refl improves the reasoning efficiency. Further experiments on solving combinatorial optimization on graphs validate that ABL-Refl can handle diverse types of data in varied dimensions, and exploit knowledge base in different forms.

2 Related Work

Recently, there has been notable progress in enhancing neural networks with reliable symbolic reasoning. Some methods use differentiable fuzzy logic (Serafini and Garcez 2016; Marra et al. 2020) or relax symbolic domain knowledge as constraints for neural network training (Xu et al. 2018; Yang, Lee, and Park 2022; Hoernle et al. 2022; Ahmed et al. 2022), while others learn constraints within neural networks by approximating logic reasoning with distributed representations (Amos and Kolter 2017; Selsam et al. 2018; Wang et al. 2019). These models tend to soften the requirements in symbolic reasoning, impacting the reliability of output generation. Models like DeepProbLog (Manhaeve et al. 2018) and NeurASP (Yang, Ishay, and Lee 2020) interpret the neural network output as a distribution over symbols and then apply a symbolic solver, incurring substantial computational costs. Abductive Learning (ABL) (Zhou 2019; Zhou and Huang 2022) attempts to integrate machine learning and logical reasoning in a balanced and mutually supporting way. It features an easy-to-use open-source toolkit (Huang et al. 2024) with many practical applications (Huang et al. 2020; Cai et al. 2021; Wang et al. 2021; Gao et al. 2024). However, the consistency optimization is with high complexity.

Another category of work related to our study also follows a similar process of prediction, error identification, and reasoning (Nair et al. 2020; Nye et al. 2021; Han et al. 2023). These methods are usually constrained in a narrow scope of domain knowledge, confined to specific mathematical problems or are bounded within a minimal world model.

Cornelio *et al.* (2023) generates a selection module to identify errors requiring symbolic reasoning rectification. In contrast to their approach which requires the preparation of a large synthetic dataset in advance, our approach automatically abduces the reflection vector during model training.

3 Abductive Reflection

This section presents problem setting and the *Abductive Reflection (ABL-Refl)* method.

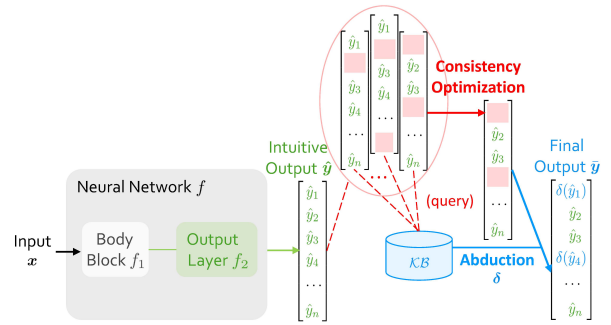


Figure 1: Abductive Learning (ABL) framework.

3.1 Problem Setting

The main task of this paper is as follows: The input is raw data x , which can be in either symbolic or sub-symbolic form, and the target output is $y = [y_1, y_2, \dots, y_n]$, with each y_i being a symbol from a set \mathcal{Y} that contains all possible output symbols. We assume two key components at our disposal: neural network f and domain knowledge base \mathcal{KB} . f can directly map x to y , and \mathcal{KB} holds constraints between the symbols in y . \mathcal{KB} can assume various forms, including propositional logic, first-order logic, mathematical or physical equations, etc., and can perform symbolic reasoning operations by exploiting the corresponding symbolic solver. The output y should adhere to the constraints in \mathcal{KB} , otherwise it will inevitably contain errors that lead to inconsistencies with the domain knowledge and incorrect reasoning results.

This problem type has broad applications. For example, it can be used to solve Sudoku puzzles, where the output y consists of $n = 81$ symbols from the set $\mathcal{Y} = \{1, 2, \dots, 9\}$, and the constraints in \mathcal{KB} are the rules of Sudoku. It can also be applied in deploying generative models for text generation, gene prediction, mathematical problem-solving, etc., producing outputs that adhere to intricate commonsense, biological, or mathematical logics in \mathcal{KB} .

3.2 Brief Introduction to Abductive Learning

When Abductive Learning (ABL) receives an input x , it initially employs f to map x into an intuitive output $\hat{y} = [\hat{y}_1, \hat{y}_2, \dots, \hat{y}_n]$. When f is under-trained, \hat{y} might contain errors leading to inconsistencies with \mathcal{KB} . ABL then tries to rectify them, and obtains a revised \bar{y} . As shown in Figure 1, the final output, \bar{y} , consists of two parts: the green part retains the results from neural network, and the blue part is the modified result obtained by abduction, a basic form of symbolic reasoning that seeks plausible explanations for observations based on \mathcal{KB} .

Specifically, the process of obtaining \bar{y} can be divided into two sequential steps. The first step, consistency optimization, determines which positions in \hat{y} include elements that contain errors causing inconsistencies, so that performing abduction at these positions will yield a \bar{y} consistent with \mathcal{KB} . Essentially, this process is pinpointing propositions (or ground atoms, etc.) which have incorrect truth assignments,

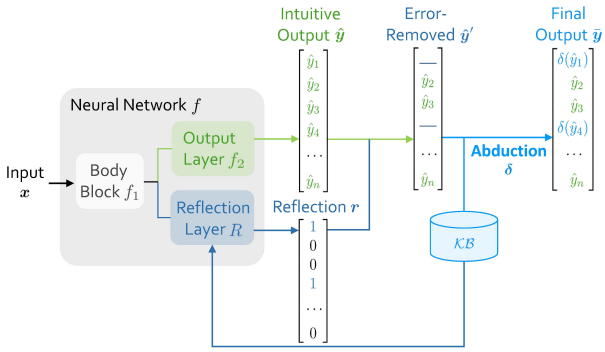


Figure 2: Architecture of Abductive Reflection (ABL-Refl). It replaces the external consistency optimization module with an efficient reflection mechanism, which is abduced directly from \mathcal{KB} .

and most neuro-symbolic tasks can be formalized into this form. Once these positions are determined, the second step is rectifying by abduction, which then becomes easy for \mathcal{KB} and its corresponding symbolic solver.

Challenge. In previous ABL, consistency optimization has always been a computational bottleneck. It operates as an external module using zeroth-order optimization methods, independent from both f and \mathcal{KB} (Dai et al. 2019; Zhou and Huang 2022). For each time of inference, it involves repetitively selecting various possible positions and querying the \mathcal{KB} to see if a consistent result can be inferred. Each query involves an invocation of \mathcal{KB} for slow symbolic reasoning. Also, since it is a complex combinatorial problem with a highly discrete nature, the number of such queries required escalates exponentially as data scale increases. This large number leads to a marked increase in time consumption, hence confines the applicability of ABL to only small datasets, usually those with output dimension n less than 10.

3.3 Architecture

To address the challenges above, we propose Abductive Reflection (ABL-Refl). In this section, we will provide a detailed description of its architecture.

Let’s first revisit the role of the neural network f when we map the input to symbols from the set \mathcal{Y} . Typically, the raw data is first passed through the body block of the network, denoted by f_1 , resulting in a high-dimensional embedding which encapsulates a wealth of feature information of the raw data. The form of f_1 varies, including structures like recurrent layers, graph convolution layers, or Transformers, etc. The result of f_1 is subsequently passed into several layers, usually linear layers, denoted by f_2 , to obtain the intuitive output: $\hat{y} = \text{argmax}(f_2(f_1(x))) \in \mathcal{Y}^n$.

Besides the structure described above, as shown in Figure 2, our architecture further incorporates a reflection layer R after the body block f_1 , generating a reflection vector: $r = \text{argmax}(R(f_1(x))) \in \{0, 1\}^n$. The reflection layer R and reflection vector r together constitute the reflection mechanism. This vector r has the same dimensionality n as the intuitive output \hat{y} , and each element, r_i , acts as a binary

classifier to indicate whether the corresponding element \hat{y}_i is an error leading to inconsistencies with \mathcal{KB} (flagged as 1 for an error, and 0 otherwise). The reflection vector r is generated concurrently with the intuitive response during inference, resonating with human cognition where cognitive reflection typically forms right upon generation of an intuitive response (Frederick 2005).

With the initial intuitive output \hat{y} and the corresponding reflection vector r , we seamlessly obtain the error-removed output \hat{y}' : In \hat{y}' , elements flagged as error by r are removed and left as blanks, while the rest are retained. Subsequently, \mathcal{KB} applies abduction to fill in these blanks, thereby generating an output \bar{y} that is consistent with \mathcal{KB} . That is:

$$\bar{y}_i = \begin{cases} \hat{y}_i, & r_i = 0 \\ \delta(\hat{y}_i), & r_i = 1 \end{cases} \quad i = 1, 2, \dots, n$$

where δ denotes abduction. We treat $\bar{y} = [\bar{y}_1, \bar{y}_2, \dots, \bar{y}_n]$ as the final output.

During model training, the reflection is abduced from \mathcal{KB} by directly leveraging information from domain knowledge (discussed later in Section 3.4). It can be seen as an attention mechanism generated from neural networks, which can help quickly focus symbolic reasoning specifically on areas it identifies as errors, hence largely narrowing the problem space of deliberate symbolic reasoning (Zhang et al. 2020).

Benefits. Compared to previous ABL implementations, ABL-Refl replaces the zeroth-order consistency optimization module with the reflection mechanism to address the computational bottleneck. In this way, the need for a substantial number of querying \mathcal{KB} is mitigated: After promptly pinpointing inconsistencies in System 1 output, regardless of the data scale, only a single invocation of \mathcal{KB} is required to obtain a rectified and more consistent output.

Another thing worth noticing is that, in the architecture, the reflection layer directly connects to the body block, which helps leveraging information from the embeddings and linking more closely with the raw data. Therefore, the reflection vector r establishes a more direct and tighter bridge between raw data and domain knowledge.

3.4 Training Paradigm

In this section, we will discuss how to train the ABL-Refl method, especially the reflection in it.

In ABL-Refl, when each input x is processed by the neural network, we obtain the intuitive output \hat{y} and the reflection vector r , and subsequently obtain the error-removed (by r) output \hat{y}' . With \hat{y} and \hat{y}' , we can measure their consistency with \mathcal{KB} , respectively. We denote these consistency measurements as $\text{Con}(\hat{y}, \mathcal{KB})$ and $\text{Con}(\hat{y}', \mathcal{KB})$, as shown in Figure 3. For a simplest example, if all elements in \hat{y} (or \hat{y}') adhere to constraints in \mathcal{KB} , the consistency measurement is 1; otherwise, it is 0.

Consequently, the improvement in consistency measurement after reflection, as denoted by

$$\Delta \text{Con}_r(\hat{y}) = \text{Con}(\hat{y}', \mathcal{KB}) - \text{Con}(\hat{y}, \mathcal{KB})$$

naturally indicates the effectiveness of the reflection vector: A higher value of it signifies that reflection r can more effectively detect inconsistencies within \hat{y} . Our training goal

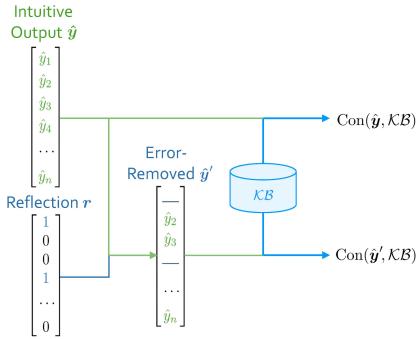


Figure 3: Consistency measurements.

is to guide the neural network’s parameters towards generating reflections that can maximize this value. Given that $\Delta\text{Con}_r(\hat{\mathbf{y}})$ is usually a discrete value, we employ the REINFORCE algorithm to achieve this goal (Williams 1992), which optimizes the policy (implicitly defined by neural network f) through maximizing a specified reward — in this case, $\Delta\text{Con}_r(\hat{\mathbf{y}})$. This process leads to the following consistency loss:

$$L_{con}(\mathbf{x}) = -\Delta\text{Con}_r(\hat{\mathbf{y}}) \cdot \nabla_{\theta} \log f_{\theta}(\hat{\mathbf{y}}, \mathbf{r} \mid \mathbf{x}) \quad (1)$$

where θ are parameters of neural network f .

Additionally, given that the time abduction required often escalates with problem size, we want to invoke it judiciously during inference, applying it only when it is truly necessary. Therefore, we aim to avoid the reflection vector from flagging too many elements in $\hat{\mathbf{y}}$ as error. To achieve this, we then introduce a reflection size loss:

$$L_{size}(\mathbf{x}) = \Phi\left(C - \frac{1}{n} \sum_{i=1}^n (1 - R(f_1(\mathbf{x}))_i)\right) \quad (2)$$

where $\Phi(a) \triangleq \max(0, a)^2$ and C is a hyperparameter ranging between 0 and 1. When C is set at a higher value, the reflection vector tends to retain a greater number of intuitive output elements instead of flagging them as error and delegating to abduction.

In addition to the above-mentioned training methods, using labeled data, we employ data-driven supervised training methods similar to common neural network training paradigm. The loss function in this process, e.g., cross-entropy loss, is denoted by $L_{labeled}(\mathbf{x}, \mathbf{y})$.

Therefore, combining all the training loss, the total loss for ABL-Refl is represented as follows:

$$\begin{aligned} \mathcal{L} &= \frac{1}{|D_l|} \sum_{(\mathbf{x}, \mathbf{y}) \in D_l} L_{labeled}(\mathbf{x}, \mathbf{y}) \\ &+ \frac{1}{|D_l \cup D_u|} \sum_{\mathbf{x} \in D_l \cup D_u} (\alpha L_{con}(\mathbf{x}) + \beta L_{size}(\mathbf{x})) \end{aligned} \quad (3)$$

where α and β are hyperparameters, $D_l = \{(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2), \dots\}$ are the labeled datasets and $D_u = \{\mathbf{x}_1, \mathbf{x}_2, \dots\}$ are the unlabeled datasets.

Note that neither L_{con} nor L_{size} , which are loss functions specifically related to the reflection, incorporate information from the data label. Instead, we leverage training information directly from \mathcal{KB} to train the reflection. Also, despite sharing the prior feature layers, the output layer f_2 and reflection layer R utilize different training information, thereby decoupling the objectives of intuitive problem-solving and inconsistency reflection.

4 Experiments

In this section, we will conduct several experiments. First, we will test our method on the NeSy benchmark task of solving Sudoku to comprehensively verify its effectiveness. Next, we will change the Sudoku input from symbols to images, which requires integrating and simultaneous reasoning with both sub-symbolic and symbolic elements, representing one of the most challenging tasks in this field. Finally, we will tackle NP-hard combinatorial optimization problems on graphs, using a knowledge base of only mathematical definitions, to demonstrate our method’s versatility. Through these experiments, we aim to answer the following questions:

- Q1** Compared to existing neuro-symbolic learning methods, can ABL-Refl achieve better performance in tasks requiring complex reasoning?
- Q2** Can ABL-Refl reduce the training resources required?
- Q3** Can ABL-Refl narrow the problem space for symbolic reasoning to achieve acceleration?
- Q4** Does ABL-Refl possess the capability for broad application, such as handling diverse data scenarios or various forms of domain knowledge?

All experiments are performed on a server with Intel Xeon Gold 6226R CPU and Tesla A100 GPU. In our experiments, we simply set hyperparameters α and β in Eq. (3) to 1, since adjusting them does not have a noticeable impact on the results. For the hyperparameter C in (2), we set it to 0.8, and have provided discussions in Appendix C, demonstrating that setting it to a value within a broad moderate range (e.g., 0.6-0.9) would always be a recommended choice. All experiments are repeated 5 times.

4.1 Solving Sudoku

Dataset and Setting. This task aims to solve a 9×9 Sudoku: Given 81 digits of 0-9 (where 0 represents a blank space) in a 9×9 board, we aim to find a solution $\mathbf{y} \in \{1, 2, \dots, 9\}^{81}$ that adhere to the Sudoku rules: no duplicate numbers are allowed in any row, column, or 3×3 subgrid. In this section, we first consider inputs in symbolic form, $\mathbf{x} \in \{0, 1, \dots, 9\}^{81}$, and use datasets from a publicly available Kaggle site (Vopani 2019).

For the neural network f , we use a simple graph neural network (GNN): the body block f_1 consists of one embedding layer and eight iterations of message-passing layers, resulting in a 128-dimensional embedding for each number, and then connects to both a linear output layer f_2 to obtain the intuitive output $\hat{\mathbf{y}}$ and a linear reflection layer R to obtain the reflection vector \mathbf{r} . We use the cross-entropy

Method	Training Time (min)	Inference Time (s)	Inference Accuracy
RRN	114.8±7.8	0.19±0.01	73.1±1.2
CL-STE	173.6±9.9	0.19±0.02	76.5±1.8
SATNet	140.3±6.8	0.11±0.01	74.1±0.4
ABL-Refl	109.8±10.8	0.22±0.02	97.4±0.3
Simple GNN	29.7±2.6	0.02±0.00	55.6±0.3

Table 1: Training time (for a total of 100 epochs using 20K training data), inference time and accuracy (on 1K test data) on solving Sudoku.

loss as $L_{labeled}$. For the domain knowledge base \mathcal{KB} , it contains the Sudoku rules mentioned above. We express \mathcal{KB} in the form of propositional logic and utilize the MiniSAT solver (Sörensson 2010), an open-source SAT solver, as the symbolic solver to leverage \mathcal{KB} and perform abduction.

For the consistency measurement, we define it as follows: one point is awarded for each row, each column and each 3×3 subgrid with no duplicate numbers, additionally, ten points are awarded if the entire board has no inconsistencies with \mathcal{KB} . In this way, it is entirely based on \mathcal{KB} . Notice that we deviated from the 1 or 0 measurement example setup mentioned in Section 3.4 to avoid a predominance of zero values in $\Delta\text{Con}_r(\hat{y})$ of Eq. (1), facilitating effective training with the REINFORCE algorithm. Similar considerations are applied in subsequent experiments.

Compared Methods and Results. We compare ABL-Refl with the following baseline methods: 1) Recurrent Relational Network (RRN) (Palm, Paquet, and Winther 2018), a pure neural network method, 2) CL-STE (Yang, Lee, and Park 2022), a representative method of logic-based regularized loss, and 3) SATNet (Wang et al. 2019). A detailed description of these methods is provided in Appendix A. We also report the result for Simple GNN, which is the very same neural network used in our setting, yet directly treats the intuitive output \hat{y} as the final output.

We report the training time (for a total of 100 epochs using 20K training data), inference time (on 1K test data) and accuracy (the percentage of completely accurate Sudoku solution boards on test data) in Table 1. We may see that our method outperforms the baselines significantly, improving by over 20% while maintaining a comparable inference time. This suggests an answer to **Q1**: ABL-Refl can achieve better reasoning performance. This improvement is primarily due to the use of abduction to rectify the neural network’s output during inference.

Furthermore, our method reaches high accuracy in only a few epochs (training curve is shown in Appendix B), significantly reducing training time. Even considering under identical training epochs, our total training time is less than baseline methods, despite involving a time-consuming symbolic solver. This partly stems from the neural network in our approach being less complex than those in baseline methods while achieving high accuracy. Overall, this suggests an answer to **Q2**: ABL-Refl can reduce the training time required.

Labeled Data	Unlabeled Data	Inference Accuracy
20K	0	97.4±0.3
10K	10K	96.3±0.3
5K	15K	95.8±0.6
2K	18K	94.7±0.8

Table 2: Inference accuracy on solving Sudoku after reducing the amount of labeled data.

We also attempt to reduce the amount of labeled data, removing labels from 50%, 75%, and 90% of the training data. We record the inference accuracy in Table 2. It can be observed that even with only 2K labeled training data, our method still achieves far better accuracy than the baseline methods with 20K labeled training data. This suggests an answer to **Q2** from another aspect: ABL-Refl can reduce the labeled training data required.

Comparing to Symbolic Solvers. We next compare our method with merely employing symbolic solvers from scratch, to demonstrate its capability in accelerating symbolic reasoning. We perform inference on 1K test data and record the accuracy and time in Table 3. The inference time for our method includes the combined duration for data processing through both the neural network (NN time) and symbolic reasoning (abduction time).

As observed in the former two lines, our method achieves a notable acceleration in the abduction process, consequently decreasing the overall inference time, with only a minor compromise in accuracy. This efficiency gain is due to the fact that in ABL-Refl, after quickly generating an intuition through the neural network, abduction only needs to focus on areas identified as necessary by the reflection vector, whereas using only symbolic solvers requires abduction to reason through all blanks in a Sudoku puzzle. Overall, this suggests an answer to **Q3**: ABL-Refl can quickly generate the reflection, thereby reducing the symbolic reasoning search space and enhancing reasoning efficiency.

We also compared with Prolog with CLP(FD) (Triska 2012) solver, by expressing the same \mathcal{KB} with a first-order constraint logic program. As shown in the table, we observe a significant reduction in abduction time and overall inference time, which puts another evidence to our previous answer to **Q3**, and also suggests an answer to **Q4**: ABL-Refl can effectively utilize the two most commonly used forms in symbolic knowledge representation, propositional logic and first-order logic.

4.2 Solving Visual Sudoku

Dataset and Setting. In this section, we modify the input from 81 symbolic digits to 81 MNIST images (handwritten digits of 0-9). We use the dataset provided in SATNet (Wang et al. 2019) and use 9K Sudoku boards for training and 1K for testing.

In order to process image data, we first pass each image through a LeNet convolutional neural network (CNN) (LeCun et al. 1998) to obtain the probability of each digit. The rest of our setting follows from that described in Section 4.1.

\mathcal{KB} Form	Solver	Method	Inference Accuracy	Inference Time (s)		
				NN Time	Abduction Time	Overall Time
Propositional logic	MiniSAT	Solver only	100±0	-	0.227±0.024	0.227±0.024
		ABL-Refl	97.4±0.3	0.021±0.004	0.196±0.015	0.217±0.019
First-order logic	Prolog with CLP(FD)	Solver only	100±0	-	105.81±5.62	105.81±5.62
		ABL-Refl	97.4±0.3	0.021±0.004	31.86±1.88	31.88±1.89

Table 3: Inference accuracy and time (on 1K test data) on solving Sudoku. For \mathcal{KB} expressed in two different forms, ABL-Refl shows notable acceleration compared to symbolic solvers in both cases.

Method	Inference Time (s)	Inference Accuracy
SATNet	0.12±0.01	63.5±2.2
CNN+Solver	0.23±0.02	67.8±4.2
ABL-Refl	0.22±0.02	77.8±5.8
ABL-Refl (with pretrained CNN)	0.22±0.02	93.5±3.2

Table 4: Inference time (on 1K test data) and accuracy on solving visual Sudoku.

Compared Methods and Results. We compare ABL-Refl with SATNet, as both methods allow for end-to-end training from visual inputs. We report the results in Table 4 and the training curve in Appendix B. Compared to SATNet, ABL-Refl shows notable improvement in reasoning accuracy within only a few training epochs. We then consider pretraining the CNN in advance using self-supervised learning methods (Chen et al. 2020) and find that this can further improve accuracy. Overall, the results further suggest positive answers to **Q1** and **Q2**.

We also compare with CNN+Solver: each image is first mapped to symbolic form by a fully trained CNN (with 99.6% accuracy on the MNIST dataset) and then directly fed into the symbolic solver to fill in the blanks and derive the final output. In such scenarios, the problem space for the symbolic solver includes all the Sudoku blanks, and additionally, since the symbolic solver cannot revise errors from CNN, any inaccuracies in CNN’s output could lead the symbolic solver to crash (i.e., output no solution). Consequently, inference accuracy and time are adversely affected. This confirms the positive answer to **Q3**.

Finally, an overview of Sections 4.1 and 4.2 also suggests an answer to **Q4**: ABL-Refl is capable of handling both symbolic and sub-symbolic forms of input data.

4.3 Solving Combinatorial Optimization Problems on Graphs

In this section, we will further expand the application domain of our method. We apply ABL-Refl to solving combinatorial optimization problems on graphs. We conduct the experiment on finding the maximum clique in this section, and provide an additional experiment in Appendix E.

Dataset and Setting. In this task, we are given a graph $G = (V, E)$ with $|V| = n$ nodes, and aim to output $\mathbf{y} \in \{0, 1\}^n$, where each index corresponds to a node, and the set of indices assigned the value of 1 collectively constitute the maximum clique. Note that this problem is a challenging NP-hard problem with extensive applications in real-life scenarios, and is generally considered challenging for neural networks (Zhang et al. 2023).

We use several datasets from the TUDatasets (Morris et al. 2020), with their basic information shown in Table 5. We use 80% of the data for training and 20% for testing.

In our method, the body layer f_1 consists of a single GAT layer (Veličković et al. 2017) and 16 gated graph convolution layers (Li et al. 2015), and the output layer f_2 and reflection layer R are both linear layers. We use binary cross-entropy loss as $L_{labeled}$. The domain knowledge base \mathcal{KB} expresses the mathematical definition of maximum clique, i.e., every pair of vertices in the output set should be connected by an edge. We use Gurobi solver, an efficient mixed-integer program solver, to perform abduction. We define the consistency measurement as follows: one point is awarded for each pair of vertices if they are not connected by an edge; additionally, the size of the output set multiplied by 10 is added if the output set is indeed a clique.

Compared Methods and Results. We compare our methods with the following baselines: 1) Erdos (Karalias and Loukas 2020), 2) Neural SFE (Karalias et al. 2022), both leading methods for solving graph combinatorial problems. Their detailed descriptions are provided in Appendix A.

We report the approximation ratios in Table 5. The approximation ratio, indicating the result set size relative to the actual maximum set size, is better when closer to 1. We may observe that our method outperforms the baseline methods, achieving near-perfect results on all datasets. This confirms the positive answer to **Q1**. Also, as the scale of the data increases, our method maintains a high level of accuracy, showing a more pronounced improvement compared to baseline methods. This suggests an answer to **Q4**: ABL-Refl is capable of handling scalable data scenarios, even in high-dimensional settings that are challenging for previous methods. Finally, an overview of this section provides another aspect to **Q4**: ABL-Refl can utilize a wide range of \mathcal{KB} , not limited to logical expressions but can also operate effectively with just the basic mathematical formulations.

Method	Dataset (Graph nums./Avg. nodes per graph/Avg. edges per graph)			
	ENZYMES (600/33/62)	PROTEINS (1113/39/73)	IMDB-Binary (1000/19/97)	COLLAB (5000/74/2457)
Erdos	0.883±0.156	0.905±0.133	0.936±0.175	0.852±0.212
Neural SFE	0.933±0.148	0.926±0.165	0.961±0.143	0.781±0.316
ABL-Refl	0.991±0.017	0.985±0.020	0.979±0.029	0.982±0.015

Table 5: Approximation ratios on finding maximum clique on different datasets.

5 Effects of Reflection Mechanism

This section provides a further analysis on the reflection mechanism. In ABL-Refl, the reflection is abduced from domain knowledge, and acts as an efficient attention mechanism to direct the focus for symbolic search. This reflection is the key in our method to accomplish the NeSy reasoning rectification pipeline, i.e., a pipeline that detects errors in neural networks and then invokes symbolic reasoning to rectify these positions. To corroborate the effectiveness of the reflection, we conduct direct comparison with other methods that achieve the same pipeline:

- 1) **ABL**, minimizing the inconsistency of intuitive output and knowledge base with an external zeroth-order consistency optimization module, as detailed in Section 3.2;
- 2) **NN Confidence**, retaining intuitive output with the top 80% confidence from the neural network result (other retain thresholds are explored in Appendix D) and passing the remaining into symbolic reasoning;
- 3) **NASR** (Cornelio et al. 2023), using a Transformer-based external selection module to detect error, and the module is trained on a large synthetic dataset in advance.

We compare them on the solving visual Sudoku task in Section 4.2. For a fair comparison, all methods employ the same neural network, \mathcal{KB} and MiniSAT solver setup. We report the recall (the percentage of errors from neural networks that can be identified), inference time and accuracy (on 1K test data) in Table 6. Note that “recall” directly evaluates the effectiveness of the detection module itself. The following analysis examines the results:

- The consistency optimization in ABL faces significant efficiency challenges due to the large data scale (output dimension $n = 81$). In such scenarios, the potential rectifications can reach up to 2^{81} , resulting in an overwhelmingly large search space for consistency optimization. Also, as an external module, its only way of interacting with \mathcal{KB} is to treat it as a black box and repetitively submit queries for consistency evaluation. As a result, it may require more than 10^9 queries to identify errors for each Sudoku example, resulting in several hours to complete inference on 1K test data.
- NN Confidence performs poorly in identifying outputs with errors. Since the pure data-driven neural network training does not explicitly incorporate \mathcal{KB} information, a low confidence from it does not necessarily indicate an inconsistency with the domain knowledge. This subsequently results in the frequent crashing in symbolic solver, therefore hampering the overall inference time

Method	Recall	Inference Time (s)	Inference Accuracy
ABL	Timeout	Timeout	Timeout
NN Confidence	82.64±2.78	0.24±0.03	64.3±6.2
NASR	95.86±0.96	0.26±0.02	82.7±4.4
ABL-Refl	99.04±0.85	0.22±0.02	93.5±3.2

Table 6: Recall, inference time and accuracy. “Timeout” indicates that inference takes more than 1 hour.

and accuracy. This result parallels human cognitive reflection abilities, which do not show much positive correlation with System 1 intuition (Pennycook et al. 2016). To further illustrate this point, we provide additional analysis, including a case study, in Appendix D.

- Our method also outperforms NASR, and notably, without the need of a synthetic dataset. This could be due to the fact that NASR’s error-selection module is trained independently from other components, and operates sequentially and separately during inference. Therefore, it can only rely on information from the output label, in contrast to our method, which can leverage information directly from the body block of neural network, establishing a deeper connection with the raw data. Additionally, in NASR, traversing the separate selection module takes additional time, whereas in ABL-Refl, the reflection is generated concurrently with the neural network output, avoiding efficiency loss.

6 Conclusion

In this paper, we present *Abductive Reflection (ABL-Refl)*. It leverages domain knowledge to abduce a reflection vector, which flags potential errors in neural network outputs and then invokes abduction, serving as an attention mechanism for symbolic reasoning to focus on a much smaller problem space. Experiments show that ABL-Refl significantly outperforms other NeSy methods, achieving excellent reasoning accuracy with fewer training resources, and has successfully enhanced reasoning efficiency.

ABL-Refl preserves the integrity of both machine learning and logical reasoning with superior inference speed and high versatility. Therefore, it has the potential for broad application. In the future, it can be applied to large language models (Mialon et al. 2023) to help identify errors within their outputs, and subsequently exploit symbolic reasoning to enhance their trustworthiness and reliability.

Acknowledgments

This research was supported by the NSFC (62176117, 62206124) and Jiangsu Science Foundation Leading-edge Technology Program (BK20232003).

References

- Ahmed, K.; Wang, E.; Chang, K.-W.; and Van den Broeck, G. 2022. Neuro-symbolic entropy regularization. In *Uncertainty in Artificial Intelligence*, 43–53. PMLR.
- Amos, B.; and Kolter, J. Z. 2017. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, 136–145. PMLR.
- Bengio, Y. 2019. From system 1 deep learning to system 2 deep learning. In *Neural Information Processing Systems*.
- 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 *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI'21)*, 1815–1821.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Cornelio, C.; Stuehmer, J.; Hu, S. X.; and Hospedales, T. 2023. Learning where and when to reason in neuro-symbolic inference. In *The Eleventh International Conference on Learning Representations*.
- Dai, W.-Z.; Xu, Q.; Yu, Y.; and Zhou, Z.-H. 2019. Bridging machine learning and logical reasoning by abductive learning. *Advances in Neural Information Processing Systems*, 32.
- Frederick, S. 2005. Cognitive reflection and decision making. *Journal of Economic perspectives*, 19(4): 25–42.
- Gao, E.-H.; Huang, Y.-X.; Hu, W.-C.; Zhu, X.-H.; and Dai, W.-Z. 2024. Knowledge-Enhanced Historical Document Segmentation and Recognition. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI'24)*, 8409–8416.
- Han, Q.; Yang, L.; Chen, Q.; Zhou, X.; Zhang, D.; Wang, A.; Sun, R.; and Luo, X. 2023. A GNN-Guided Predict-and-Search Framework for Mixed-Integer Linear Programming. *arXiv preprint arXiv:2302.05636*.
- Hitzler, P. 2022. *Neuro-symbolic artificial intelligence: The state of the art*. IOS Press.
- Hoernle, N.; Karampatsis, R. M.; Belle, V.; and Gal, K. 2022. Multiplexnet: Towards fully satisfied logical constraints in neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5700–5709.
- 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 *Proceedings of the 20th IEEE International Conference on Data Mining (ICDM'20)*, 1070–1075.
- Huang, Y.-X.; Hu, W.-C.; Gao, E.-H.; and Jiang, Y. 2024. ABLkit: A Python Toolkit for Abductive Learning. *Frontiers of Computer Science*, pp. to appear.
- Kahneman, D. 2011. *Thinking, fast and slow*. macmillan.
- Karalias, N.; and Loukas, A. 2020. Erdos goes neural: an unsupervised learning framework for combinatorial optimization on graphs. *Advances in Neural Information Processing Systems*, 33: 6659–6672.
- Karalias, N.; Robinson, J.; Loukas, A.; and Jegelka, S. 2022. Neural Set Function Extensions: Learning with Discrete Functions in High Dimensions. *arXiv preprint arXiv:2208.04055*.
- 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, Y.; Tarlow, D.; Brockschmidt, M.; and Zemel, R. 2015. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*.
- Manhaeve, R.; Dumancic, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2018. Deepproblog: Neural probabilistic logic programming. *advances in neural information processing systems*, 31.
- Marra, G.; Giannini, F.; Diligenti, M.; and Gori, M. 2020. Integrating learning and reasoning with deep logic models. In *Machine Learning and Knowledge Discovery in Databases*, 517–532. Springer.
- Mialon, G.; Fourrier, C.; Swift, C.; Wolf, T.; LeCun, Y.; and Scialom, T. 2023. GAIA: a benchmark for General AI Assistants. *arXiv preprint arXiv:2311.12983*.
- Morris, C.; Kriege, N. M.; Bause, F.; Kersting, K.; Mutzel, P.; and Neumann, M. 2020. Tudataset: A collection of benchmark datasets for learning with graphs. *arXiv preprint arXiv:2007.08663*.
- Nair, V.; Bartunov, S.; Gimeno, F.; Von Glehn, I.; Lichocki, P.; Lobov, I.; O'Donoghue, B.; Sonnerat, N.; Tjandraatmadja, C.; Wang, P.; et al. 2020. Solving mixed integer programs using neural networks. *arXiv preprint arXiv:2012.13349*.
- Nye, M.; Tessler, M.; Tenenbaum, J.; and Lake, B. M. 2021. Improving coherence and consistency in neural sequence models with dual-system, neuro-symbolic reasoning. *Advances in Neural Information Processing Systems*, 34: 25192–25204.
- Palm, R.; Paquet, U.; and Winther, O. 2018. Recurrent relational networks. *Advances in neural information processing systems*, 31.
- Pennycook, G.; Cheyne, J. A.; Koehler, D. J.; and Fugelsang, J. A. 2016. Is the cognitive reflection test a measure of both reflection and intuition? *Behavior research methods*, 48: 341–348.
- Selsam, D.; Lamm, M.; Bünz, B.; Liang, P.; de Moura, L.; and Dill, D. L. 2018. Learning a SAT solver from single-bit supervision. *arXiv preprint arXiv:1802.03685*.
- Serafini, L.; and Garcez, A. d. 2016. Logic tensor networks: Deep learning and logical reasoning from data and knowledge. *arXiv preprint arXiv:1606.04422*.

- Sinayev, A.; and Peters, E. 2015. Cognitive reflection vs. calculation in decision making. *Frontiers in psychology*, 6: 532.
- Sörensson, N. 2010. Minisat 2.2 and minisat++ 1.1. *A short description in SAT Race*.
- Triska, M. 2012. The finite domain constraint solver of SWI-Prolog. In *Functional and Logic Programming: 11th International Symposium*, 307–316. Springer.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Vopani. 2019. 9 Million Sudoku Puzzles and Solutions. <https://www.kaggle.com/datasets/rohanrao/sudoku>. Accessed: 2024-08-01.
- 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 *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'21)*, 3688–3696.
- Wang, P.-W.; Donti, P.; Wilder, B.; and Kolter, Z. 2019. Satnet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *International Conference on Machine Learning*, 6545–6554. PMLR.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Reinforcement learning*, 5–32.
- Xu, J.; Zhang, Z.; Friedman, T.; Liang, Y.; and Broeck, G. 2018. A semantic loss function for deep learning with symbolic knowledge. In *International conference on machine learning*, 5502–5511. PMLR.
- Yang, Z.; Ishay, A.; and Lee, J. 2020. Neurasp: Embracing neural networks into answer set programming. In *29th International Joint Conference on Artificial Intelligence*.
- Yang, Z.; Lee, J.; and Park, C. 2022. Injecting logical constraints into neural networks via straight-through estimators. In *International Conference on Machine Learning*, 25096–25122. PMLR.
- Zhang, B.; Luo, S.; Wang, L.; and He, D. 2023. Rethinking the expressive power of gnns via graph biconnectivity. *arXiv preprint arXiv:2301.09505*.
- Zhang, W.; Sun, Z.; Zhu, Q.; Li, G.; Cai, S.; Xiong, Y.; and Zhang, L. 2020. NLocalSAT: Boosting local search with solution prediction. *arXiv preprint arXiv:2001.09398*.
- Zhou, Z.-H. 2019. Abductive learning: towards bridging machine learning and logical reasoning. *Science China Information Sciences*, 62: 1–3.
- Zhou, Z.-H.; and Huang, Y.-X. 2022. Abductive Learning. In Hitzler, P.; and Sarker, M. K., eds., *Neuro-Symbolic Artificial Intelligence: The State of the Art*, 353–369. Amsterdam: IOS Press.