BREAKING MAJORITY DOMINANCE: PROGRESSIVE ABDUCTIVE LEARNING UNDER CONCEPT SKEWNESS

Anonymous authors
Paper under double-blind review

ABSTRACT

Emulating the human interplay between perception and reasoning, abductive learning (ABL) is proposed to improve the perception model's ability in predicting concept labels without explicit annotations by drawing knowledge-consistent supervision signals from logical abduction, thereby benefiting reasoning. However, in practice, concept distributions are often skewed, leading to a **biased feedback loop** in ABL: early predictions favor majority concepts, making majority-dominated abduction candidates more likely to be reused as pseudo-labels to supervise the update of the perception model, further amplifying the bias against minority concepts. To solve this problem, we propose distribution-prior-rectified progressive abductive learning (DPRP-ABL). DPRP-ABL limits ambiguous candidate competition early in training by progressively expanding nested knowledge sub-bases. Within each active candidate set, it uses estimated concept distribution priors to alleviate majority-favoring scoring bias. Our analysis characterizes how these components affect candidate selection at a fixed model state and establishes conditions for preserving ground-truth feasibility across progressive phases. Experiments on two typical tasks indicate that DPRP-ABL improves reasoning-output performance and concept recognition under concept skewness.

1 INTRODUCTION

Humans can perceive raw sensory signals and reason over structured knowledge to reach a conclusion, motivating neuro-symbolic learning (NeSy), which pursues replicating such synergy (Solso et al., 2008; Hitzler & Sarker, 2022). Neural models specialize in extracting patterns from unstructured data but struggle to exploit structured domain knowledge, whereas symbolic systems provide explicit and verifiable reasoning but typically require symbolic inputs (Raedt et al., 2020).

To bridge this gap, abductive learning (ABL) (Dai et al., 2019) offers a framework in which the system perceives raw observations for symbols and reasons over them with domain knowledge within a trainable and iterative loop. Specifically, given a raw input sequence, the perception model first predicts concept labels, which may initially violate domain knowledge as the model is undertrained. Abduction then generates a set of candidates compatible with domain knowledge, from which the best candidate is selected as revised pseudo-labels of concepts via consistency optimization and is subsequently fed back to supervise the update of the perception model. Once trained, the perception model maps new observations to symbolic concepts, based on which system derives reasoning outputs. Nowadays, ABL has demonstrated potential in tasks where knowledge can compensate for the absence of explicit concept annotations (Huang et al., 2020; Yang et al., 2024b; Wu et al., 2025).

However, in practice, widespread concept skewness, i.e., concepts having sharply different marginals with some occurring much more frequently than others, severely impedes the deployment of ABL (Yang et al., 2024a; Tsamoura et al., 2025). Unlike traditional class imbalance in supervised learning, skewness-induced bias in ABL can be amplified by the abduction process itself. Specifically, since concept labels are not available as direct training supervision, the perception model instead learns from revised pseudo-labels selected through competition among candidates compatible with knowledge. Early in training, the model tends to assign higher posterior probabilities to majority concepts, causing majority-dominated candidates to be selected even when their concept assignments are incorrect. Once reused as pseudo-labels, these candidates further reduce the supervision available to minority concepts and widen their score disadvantage in subsequent it-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

erations. Coupled effects of such "competitive score dilution" and "supervision starvation" lead to a vicious cycle that continuously causes bias in the perception-abduction loop. It will be intensified as longer reasoning chains or more complex knowledge enlarge the abduction space.

Although existing methods have been successfully applied to various applications, they only focus on one aspect of this coupled challenge and are unable to effectively address our problems. Imbalanced learning methods rectify class distributions or decision boundaries by resampling, reweighting, meta-learning, or ensembles (Sağlam & Cengiz, 2022; Guo et al., 2022; Moniz & Cerqueira, 2021; Pirizadeh et al., 2021), but they typically rely on directly observed class labels and therefore cannot be used to tackle competition among candidates in ABL. ABL methods improve consistency search, ambiguity handling, or curriculum construction (Huang et al., 2021; He et al., 2024; Hu et al., 2025b), but they generally assume the concept distribution is uniform and do not explicitly rectify skewness-induced bias in candidate scoring.

To solve these problems, we propose distribution-prior-rectified progressive abductive learning (DPRP-ABL). DPRP-ABL uses structure-progressive scheduling to perform early abduction with nested knowledge sub-bases, thereby narrowing the number of candidates before gradually recovering the full abduction space. Within each active space, it applies distribution-prior-aware rectification, increasing the competitiveness of candidates that contain rare concepts. These components are complementary. Progression limits early candidate competition, while rectification mitigates skewness-induced scoring bias among candidates. Theoretically, we analyze how candidate-size contraction and score rectification affect selection at a fixed model state and establish conditions under which the ground-truth concept assignment remains feasible across nested phases. Experiments on the skewed digit addition and judicial sentencing tasks demonstrate gains in reasoning-output performance and concept recognition. Main contributions are summarized as follows.

- We formulate concept skewness in ABL and identify a vicious cycle in which competitive score dilution and supervision starvation continuously cause bias.
- DPRP-ABL combines structure-progressive knowledge scheduling with distribution-prior-aware rectification, alleviating skewness-induced bias in the perception-abduction loop.
- We characterize the fixed-state effects of candidate-size contraction and score rectification and establish conditions under which nested phases preserve ground-truth feasibility.
- Experiments verify the superiority of DPRP-ABL, indicating that jointly controlling candidate size and candidate scoring can improve performance in ABL with concept skewness.

2 PRELIMINARY

Knowledge Base. Knowledge base ($\mathcal{KB}$) serves as a critical bridge connecting raw data perception with high-level logical reasoning (Zhou, 2019). It explicitly embeds structured domain knowledge that is typically expressed via first-order logic or logic programming, which offers strong symbolic interpretability while maintaining computational tractability for reasoning (Dai et al., 2019). Formally, a classical $\mathcal{KB}$ is represented as a set of rules in the Horn clause form (Horn, 1951).

$$A \leftarrow B_1 \wedge B_2 \wedge \cdots \wedge B_m, \tag{1}$$

where "$\leftarrow$" denotes logical implication, indicating that the head of the rule, A, holds when the conjunction of the predicates in the rule body $\bigwedge_{i=1}^{m} B_i$ is satisfied. With such symbolic association provided by $\mathcal{KB}$, forward deduction, deriving A from the premises $\bigwedge_{i=1}^{m} B_i$, and backward abduction, seeking possible explanations $\bigwedge_{i=1}^{m} B_i$ for the observation of A, can be performed.

Standard ABL Formulation. Denote $\mathbf{X} = (\boldsymbol{x}_1, \cdots, \boldsymbol{x}_m)$ as the input sequence of raw observations, where $\boldsymbol{x}_i \in \mathcal{X}$ corresponds to concept label $z_i \in \mathcal{Z} = \{z^1, \cdots, z^C\}$, which is typically unavailable. $\mathcal{KB}$ defines how the concept can entail the target label $y \in \mathcal{Y}$, where each concept z plays the role of a predicate in the rule body and target y is in the rule head. The goal is to learn a perception model $f_\theta : \mathcal{X} \to \mathcal{Z}$ that maps each $\boldsymbol{x}_i$ to its concept z_i, such that the resulting sequence $\hat{\mathbf{z}} = (\hat{z}_1, \cdots, \hat{z}_m)$ satisfies $\hat{\mathbf{z}} \wedge \mathcal{KB} \models y$. During training, z cannot be used as direct supervision. Instead, the perception model predicts $P_\theta(z_i \mid \boldsymbol{x}_i)$, and abduction space $\mathcal{S} = \{\mathbf{z} \in \mathcal{Z}^m | \mathbf{z} \wedge \mathcal{KB} \models y\}$ is constructed, where $\mathbf{z}$ also can be called as "candidate". Standard consistency-based ABL selects

$$\bar{\mathbf{z}} = \arg\max_{\mathbf{z} \in \mathcal{S}} \prod_{i=1}^{m} P_\theta(z_i \mid \boldsymbol{x}_i), \tag{2}$$

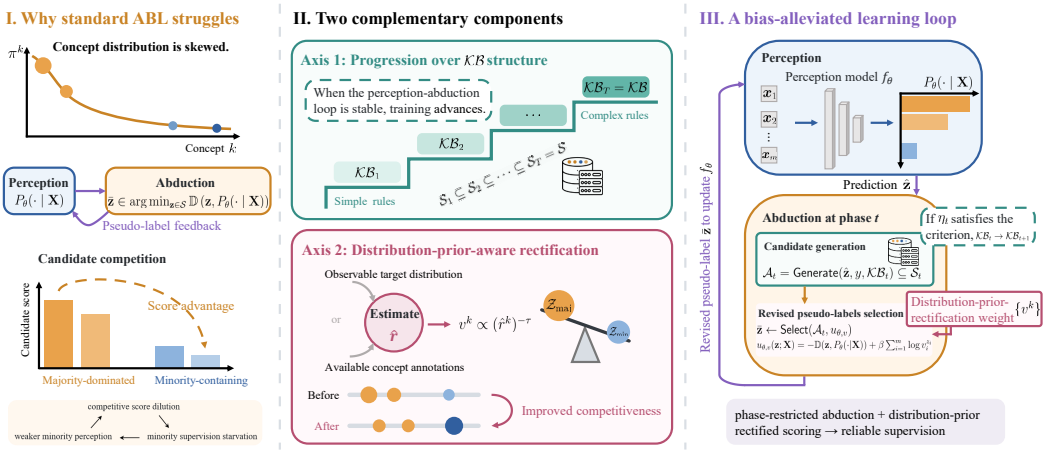

Figure 1: The illustration of DPRP-ABL. In axis 1, nested sub-bases restrict early abduction to simpler rules and fewer candidates. In axis 2, it estimates concept marginals from available supervision and converts them into distribution-prior-rectification weights.

as the revised pseudo-label to update f_θ. This perception-abduction loop iterates until convergence.

Concept Skewness and Challenges. Let $\pi^k = \Pr(z = z^k)$ and $\boldsymbol{\pi} = (\pi^1, \ldots, \pi^C)$. A task exhibits concept skewness when the majority concept subset $\mathcal{Z}_{\text{maj}}$ accounts for most observations and the remaining concepts in $\mathcal{Z}_{\text{min}} = \mathcal{Z} \setminus \mathcal{Z}_{\text{maj}}$ have substantially smaller marginals. Concept skewness directly affects the perception-abduction loop. Specifically, in early training, predictions often reflect the empirical marginals before discriminative features are learned. Thus, for a majority-dominated candidate $\mathbf{z}_b$ and a minority-containing candidate $\mathbf{z}_a$, their approximate score ratio is

$$\frac{\prod_i P_\theta(z_{b,i} \mid \boldsymbol{x}_i)}{\prod_i P_\theta(z_{a,i} \mid \boldsymbol{x}_i)} \approx \frac{\prod_i \pi^{z_{b,i}}}{\prod_i \pi^{z_{a,i}}} \gg 1. \tag{3}$$

Although $\mathbf{z}_a$ is correct, it will not be selected in abduction. We call such loss of competitiveness for the correct minority-containing candidate "competitive score dilution". Once the incorrect majority-dominated candidate is reused as pseudo-labels, it is difficult for the perception model to obtain sufficient supervision for minority concepts, causing "supervision starvation". This vicious cycle is intensified when a larger abduction space introduces more majority-dominated competitors (Hu et al., 2025b), motivating us to control both the abduction space and candidate scoring.

3 Method

3.1 Overview

The main framework of the proposed DPRP-ABL is presented in Fig. 1. The core insight is to jointly control the complexity of $\mathcal{KB}$ to regulate the abduction space and estimate marginals to rectify candidate scoring, thereby limiting early candidate competition and mitigating skewness-induced bias in the perception-abduction loop. Formally, we define the utility of candidate $\mathbf{z}$ as

$$u_{\theta,v}(\mathbf{z}; \mathbf{X}) = -\mathbb{D}\left(\mathbf{z}, P_\theta(\cdot|\mathbf{X})\right) + \beta \sum_{i=1}^{m} \log v_i^{z_i}, \tag{4}$$

where $\mathbb{D}$ measures disagreement with the perception model, $v_i^{z_i} > 0$ is the distribution-prior-rectification weight for the concept at position i, and $\beta \geq 0$ is the hyperparameter controlling rectification strength. The first term follows the idea of standard consistency-based selection, whereas the second changes the relative scores of concepts with different marginals. At phase t, the abductor selects revised pseudo-label $\bar{\mathbf{z}}^{(t)} = \arg\max_{\mathbf{z} \in \mathcal{S}_t} u_{\theta,v}(\mathbf{z}; \mathbf{X})$, which is subsequently used to update the perception model by

$$\theta \leftarrow \arg\min_\theta \frac{1}{mN_t} \sum_{n=1}^{N_t} \sum_{i=1}^{m} L\left(\bar{z}_{n,i}^{(t)}, f_\theta(\boldsymbol{x}_{n,i})\right). \tag{5}$$

Thus, the distribution-prior-rectification weights affect learning through pseudo-label selection, rather than acting as a differentiable regularizer on θ.

3.2 Axis 1: Progression over $\mathcal{KB}$ Structure

Instead of exposing the perception model to the full $\mathcal{KB}$ all at once, DPRP-ABL constructs nested sub-bases $\mathcal{KB}_1 \subseteq \cdots \subseteq \mathcal{KB}_T = \mathcal{KB}$, guided by which the training is conducted in multiple phases. Constructing follows three intuitive principles (Hu et al., 2025b) that together ensure each sub-base builds upon the last in a stable and coherent manner.

P1: Stepwise structural complexity. According to structural complexity, simpler rules are placed in earlier phases, while more complex rules are introduced later. This conforms to human learning law, where we first master the basics under simpler instruction before advancing to harder content.

P2: Dependency cohesion. Rules that usually co-occur or form tight reasoning chains are grouped into the same phase, ensuring that learning is in a coherent and locally complete context rather than encountering isolated rule fragments.

P3: Self-contained reasoning. The sub-base must be sufficient to derive conclusions without relying on rules not yet introduced. This guarantees that rules within each phase are logically sound and produce supervision signals of the same quality as the full $\mathcal{KB}$, merely over a smaller size.

Denote $\mathbf{a}_n$ as task-specific structural metadata available to the offline progression constructor, and the difficulty stratification function $\phi(\mathbf{a}_n, y_n) \in \{1, \ldots, T\}$ assigns the n-th sample to its earliest admissible phase. There are the active dataset $\mathcal{D}_t = \{(\mathbf{X}_n, y_n) : \phi(\mathbf{a}_n, y_n) \leq t\}$ and $\mathcal{KB}_t = \mathcal{KB}_{|\mathcal{Z}_t}$ at the t-th phase. Note that the metadata $\mathbf{a}_n$ determine only the offline phase assignment and active rules. They are not inputs to f_θ, prediction targets, loss terms, or backpropagation signals.

Besides, phase advancement is based on the agreement between current predictions of the perception model and revised pseudo-labels, i.e.,

$$\eta_t = \mathbb{M}\left(\hat{\mathbf{z}}_t, \bar{\mathbf{z}}_t^{(t)}\right), \tag{6}$$

where $\mathbb{M}(\cdot, \cdot)$ is the agreement metric. When it satisfies the phase-specific criterion for consecutive epochs, it indicates the current perception-abduction loop is stable and is ready to advance to benefit from increased complexity.

3.3 Axis 2: Distribution-Prior-Aware Rectification

As a complement, DPRP-ABL alleviates the skewness-induced bias by distribution-prior-aware rectification. Firstly, it needs the estimate $\hat{\boldsymbol{r}} = (\hat{r}^1, \cdots, \hat{r}^C)$ of concept marginals. The estimator depends on available supervision. We provide two common cases encountered in practical tasks.

Case 1: Estimation from the target distribution. When the symbolic transition $\sigma : \mathcal{Z}^m \to \mathcal{Y}$ is an analytic and enumerable function, i.e., $\sigma^{-1}(y)$ can be explicitly listed, the observed target distribution $q(y)$ determines $\hat{\boldsymbol{r}}$ with the assumption that concepts are independent and identically distributed by

$$\hat{\boldsymbol{r}} = \underset{\boldsymbol{r} \in \Delta^{C-1}}{\arg\min} - \sum_y q(y) \log \sum_{\mathbf{z} \in \sigma^{-1}(y)} \prod_{i=1}^{m} r^{z_i}, \tag{7}$$

where $\Delta^{C-1} = \{\boldsymbol{r} \in \mathbb{R}^C | r^k \geq 0, \sum_{k=1}^{C} r^k = 1\}$ denotes the probability simplex over C concepts.

Case 2: Estimation from the labeled subset. When a small concept-labeled subset is available, $\hat{r}^k$ can be calculated as the Laplace-smoothed marginal (Manning & Schütze, 1999) of concept z^k, i.e.,

$$\hat{r}^k = \frac{n_k + \lambda}{\sum_{j=1}^{C} n_j + \lambda C}, \tag{8}$$

where n_k denotes the cumulative number of concept z^k in the current estimation pool, which may be updated with revised pseudo-labels, and λ is the Laplace-smoothing parameter.

Then, the distribution-prior-rectification weight is calculated by a tempered inverse transform, i.e.,

$$v^k = \frac{(\hat{r}^k)^{-\tau}}{\frac{1}{C}\sum_{j=1}^{C}(\hat{r}^j)^{-\tau}}, \quad k = 1, \cdots, C \tag{9}$$

where $\tau \in [0, 1]$ is the temperature hyperparameter. Setting $\tau = 0$ disables rectification, while larger τ assigns stronger compensation to rarer concepts.

3.4 Training Procedure

During offline preparation, DPRP-ABL constructs the nested sub-bases $\{\mathcal{KB}_t\}_{t=1}^{T}$ and sub-datasets $\{\mathcal{D}_t\}_{t=1}^{T}$ and initializes the marginal estimate $\hat{\boldsymbol{r}}$, from which the rectification weights $\boldsymbol{v}$ are derived. Training then proceeds through T structure-progressive phases. At phase t, the perception model first produces concept predictions on $\mathcal{D}_t$. The abductor enumerates candidates under $\mathcal{KB}_t$ and selects revised pseudo-labels using the candidate utility in Eq. (4). These pseudo-labels subsequently supervise the update of the perception model. The phase advances when η_t in Eq. (6) meets the stability criterion or the maximum phase budget is reached. Training terminates when the final phase completes its stopping condition. Algorithm 1 summarizes the training procedure.

4 Theoretical Analysis

To provide formal analysis of DPRP-ABL, we establish several complementary guarantees. All proofs are in Appendix G.

4.1 Effect of Candidate-Size Contraction

For a finite non-empty candidate set $\mathcal{A}$, define the temperature-indexed selection distribution

$$p_\gamma(\mathbf{z}; \mathcal{A}) = \begin{cases} \dfrac{\exp\left(u_{\theta,v}(\mathbf{z}; \mathbf{X})/\gamma\right)}{\sum_{\mathbf{z}' \in \mathcal{A}} \exp\left(u_{\theta,v}(\mathbf{z}'; \mathbf{X})/\gamma\right)}, & \gamma > 0, \\ \dfrac{\mathbb{I}[\mathbf{z} \in \mathcal{M}(\mathcal{A})]}{|\mathcal{M}(\mathcal{A})|}, & \gamma = 0, \end{cases} \tag{10}$$

where $\mathcal{M}(\mathcal{A}) = \arg\max_{\mathbf{z}' \in \mathcal{A}} u_{\theta,v}(\mathbf{z}'; \mathbf{X})$. The case $\gamma = 0$ corresponds to deterministic maximization with uniform tie-breaking and is also the zero-temperature limit of the softmax selector.

Theorem 1. *At the t-th phase, for every candidate $\mathbf{z} \in \mathcal{A}_t \subseteq \mathcal{A}$ and every $\gamma \geq 0$,*

$$p_\gamma(\mathbf{z}; \mathcal{A}_t) \geq p_\gamma(\mathbf{z}; \mathcal{A}). \tag{11}$$

For $\gamma > 0$, the exact probability ratio is

$$\frac{p_\gamma(\mathbf{z}; \mathcal{A}_t)}{p_\gamma(\mathbf{z}; \mathcal{A})} = \frac{Z_\gamma(\mathcal{A})}{Z_\gamma(\mathcal{A}_t)} = 1 + \frac{Z_\gamma(\mathcal{A} \setminus \mathcal{A}_t)}{Z_\gamma(\mathcal{A}_t)}, \tag{12}$$

where $Z_\gamma(\cdot) = \sum_{\mathbf{z}' \in \cdot} \exp\left(u_{\theta,v}(\mathbf{z}'; \mathbf{X})/\gamma\right)$.

4.2 Effect of Score Rectification

Let $P_{\theta,i}(k) = P_\theta(z_i = k|\mathbf{X}), k \in \mathcal{Z}_i$, where $\mathcal{Z}_i$ is the finite concept space at the i-th position, and define the negative log-likelihood disagreement $\mathbb{D}_{\mathrm{NLL}}(\mathbf{z}, P_\theta) = -\sum_{i=1}^{m} \log P_{\theta,i}(z_i)$.

Theorem 2. *Assume that $P_{\theta,i}(k) > 0$ and $v_i^k > 0$ for all i and $k \in \mathcal{Z}_i$. For $\beta \geq 0$, define $Z_i = \sum_{k \in \mathcal{Z}_i} P_{\theta,i}(k)(v_i^k)^\beta$ and the probability $\widetilde{P}_{\theta,i}(k) = \frac{P_{\theta,i}(k)(v_i^k)^\beta}{Z_i}$. Then, for every $\mathbf{z}$, there is*

$$\mathbb{D}_{\mathrm{NLL}}(\mathbf{z}, P_\theta) - \beta \sum_{i=1}^{m} \log v_i^{z_i} = -\sum_{i=1}^{m} \log \widetilde{P}_{\theta,i}(z_i) - \sum_{i=1}^{m} \log Z_i. \tag{13}$$

Consequently, for any $\mathcal{A}$, there is

$$\arg\min_{\mathbf{z} \in \mathcal{A}} \left[\mathbb{D}_{\mathrm{NLL}}(\mathbf{z}, P_\theta) - \beta \sum_{i=1}^{m} \log v_i^{z_i} \right] = \arg\min_{\mathbf{z} \in \mathcal{A}} \left[-\sum_{i=1}^{m} \log \widetilde{P}_{\theta,i}(z_i) \right]. \tag{14}$$

4.3 Feasibility and Ground-Truth Coverage under Phase Transitions

For any sample $(\mathbf{X}, y)$ with ground-truth concept assignment $\mathbf{z}^*$, let $\kappa : \mathcal{A} \to \mathbb{R}_{\geq 0}$ be a structural-complexity measure and $b_1 \leq b_2 \leq \cdots \leq b_T$ be thresholds. There is $\mathcal{A}_t = \{\mathbf{z} \in \mathcal{A} : \kappa(\mathbf{z}) \leq b_t\}$.

Theorem 3. *For any sample $(\mathbf{X}, y)$ and abduction state, there is*

$$\mathcal{A}_1 \subseteq \mathcal{A}_2 \subseteq \cdots \subseteq \mathcal{A}_T \subseteq \mathcal{A} \subseteq \mathcal{S}. \tag{15}$$

If the ground-truth concept assignment satisfies $\mathbf{z}^* \in \mathcal{A}$ and $\kappa(\mathbf{z}^*) \leq b_t$, then

$$\mathbf{z}^* \in \mathcal{A}_t \subseteq \mathcal{S}. \tag{16}$$

Consequently, every candidate selected from $\mathcal{A}_t$ is consistent with the full constraints of the task, and the ground-truth concept assignment is preserved whenever the base-coverage and phase-threshold conditions are met.

5 Experiments

We evaluate DPRP-ABL [1] on two tasks that instantiate the discrete and continuous $\mathcal{KB}$ that are detailed in Appendix D, respectively. Experiments and analysis focus on answering questions:

- Does DPRP-ABL offer performance improvements compared to existing methods?
- Does DPRP-ABL recognize concepts or yield gains for minority concept learning?
- Are both axes indispensable?

5.1 Settings

Tasks. **Digit addition task**, which has discrete fixed $\mathcal{KB}$, is the most widely used benchmark in NeSy. Taking a sequence of $2d$ digit images as input and the sum of two d-digit numbers as the target, the perception model learns to recognize the digit concept of each image by serving the symbolic addition rules as $\mathcal{KB}$. **Judicial sentencing task**, which has parametric continuous $\mathcal{KB}$, is a real-world application that can demonstrate the practical applicability of DPRP-ABL in complex domains. Taking a written legal case description and a monetary amount as input, with sentenced months as the target, the perception model learns to extract the attribute concept about legal factors such as "surrender", "young", "forgive", and so on. The $\mathcal{KB}$ couples legal rules with the parametric sentencing function whose parameters are co-optimized with the perception model.

Datasets. **For digit addition task**, MNIST(LeCun et al., 1998), SVHN(Netzer et al., 2011), and CIFAR(Krizhevsky & Hinton, 2009), which covers grayscale handwritten digits, colored street scene digits, and colored physical entities to increase the difficulty of perception, are adopted as the base image source. To simulate concept skewness, the digit concept set $\{0, \cdots, 9\}$ is partitioned into a majority subset $\mathcal{Z}_{\text{maj}} = \{0, 1, 2, 3, 4\}$ and a minority subset $\mathcal{Z}_{\text{min}} = \{5, 6, 7, 8, 9\}$. Digit images belonging to $\mathcal{Z}_{\text{min}}$ are downsampled and retained only samples amounting to 5% of the majority size. We evaluate three sequence lengths $d \in \{1, 2, 3\}$ corresponding to addition with up to three-digit operands. **For judicial sentencing task**, the dataset released by (Huang et al., 2020) is used, which contains 687 criminal judgment records for theft cases adjudicated with eight binary attributes and the sentenced months. The concept distribution is naturally skewed, where rare attributes such as "young" and "surrender" appear in fewer than 10% of cases, while common attributes such as "no damage", "again", and "attitude" appear in over 50%. In training, a small subset of cases are assigned full attribute annotation for perception pre-training, $\mathcal{KB}$ learning, and estimator seeding.

Methods and metrics. Several methods are selected for comparison. There are NeSy methods, including NeurASP (Yang et al., 2020), DeepStochLog (Winters et al., 2021), and ImbNSL (Tsamoura et al., 2025), and ABL methods, including the standard consistency-based ABL method, that is ABL_{ham} measuring based on Hamming distance (Dai et al., 2019) and ABL_{con} measuring based on maximum likelihood (Dai et al., 2019), ABLSim (Huang et al., 2021), A^3BL (He et al., 2024), SS-ABL (Huang et al., 2020), and C-ABL (Hu et al., 2025b). Different methods are suited to different tasks, detailed introductions of which can be found in Appendix J. To quantify, several

[1] Codes: `https://anonymous.4open.science/r/DPRP-ABLcodes_iclr27id32415`.

metrics are selected for evaluation. There are accuracy (ACC), area under the curve (AUC), recall (RECM), precision (PREM), F1 score (F1), mean average precision (MAP), geometric mean (G-mean), Jensen-Shannon divergence (JS), mean squared error (MSE), and mean absolute error (MAE), which are elaborately described in Appendix H.2. In the digit addition task, the perception model is the LeNet-5 architecture (LeCun et al., 1998) for grayscale images and the ResNet-18 architecture (He et al., 2015) for colored images. In the judicial sentencing task, the perception model is BERT-base-Chinese (Devlin et al., 2019).

5.2 Task 1: d-Digit Addition[2]

Performance analysis. To demonstrate the superiority of DPRP-ABL, we first analyze classification performance at the sum target level. Results are recorded in Table 1, from which we have the following observations:

- DPRP-ABL demonstrates superior performance in most cases. On the SVHN dataset, its ACC exceeds the second-best methods by 3.44%, 3.40%, and 2.52% at $d = 1, 2, 3$, respectively. On the CIFAR dataset, it trails the best method at $d = 1$ but ranks first at $d = 2$ and $d = 3$, with margins of 0.26% and 3.25%, respectively.
- DeepStochLog achieves the highest ACC on the MNIST dataset and the CIFAR dataset with $d = 1$. In contrast, DPRP-ABL leads across all SVHN settings and on multi-digit CIFAR, indicating that its advantage is more pronounced when perceptual difficulty is accompanied by longer reasoning chains.
- DPRP-ABL achieves the lowest or near-lowest per-epoch time in most cases, maintaining competitive performance. Several methods with excellent classification performance become more expensive as d increases. Our efficiency is consistent with restricting early abduction to $\mathcal{S}_t \subseteq \mathcal{S}$ while recovering the full knowledge base in the final phase.

Table 1: Comparison "$\mathrm{Avg}_{\mathrm{std}}$" of sum target ACC (%, top) and epoch training time (seconds, bottom) on the d-digit addition task. "N/A" indicates that the result is not available within 30 hours.

Methods	MNIST			SVHN			CIFAR		
	$d=1$	$d=2$	$d=3$	$d=1$	$d=2$	$d=3$	$d=1$	$d=2$	$d=3$
NeurASP	$95.74_{0.25}$	N/A	N/A	$48.94_{2.17}$	N/A	N/A	$15.91_{0.87}$	N/A	N/A
DeepStochLog	$\mathbf{95.90_{0.15}}$	$\mathbf{93.56_{0.22}}$	$\mathbf{91.59_{0.47}}$	$75.70_{1.02}$	$65.80_{1.64}$	$57.10_{1.35}$	$\mathbf{28.87_{4.49}}$	$11.82_{2.96}$	$4.82_{1.16}$
ImbNSL	$81.64_{18.75}$	$63.39_{27.97}$	$22.81_{25.66}$	$64.45_{10.11}$	$62.45_{9.23}$	$17.84_{9.01}$	$12.07_{12.45}$	$5.75_{4.17}$	$0.81_{0.03}$
$\mathrm{ABL_{ham}}$	$32.34_{5.02}$	$54.17_{4.73}$	$52.86_{2.21}$	$46.24_{7.29}$	$12.45_{0.72}$	$0.59_{0.32}$	$6.47_{1.18}$	$0.57_{0.17}$	$0.07_{0.02}$
$\mathrm{ABL_{con}}$	$93.29_{1.72}$	$91.61_{1.03}$	$88.88_{0.71}$	$76.05_{1.07}$	$65.09_{1.24}$	$54.75_{2.11}$	$12.96_{4.98}$	$13.66_{5.46}$	$9.30_{5.51}$
ABLSim	$93.66_{1.01}$	$91.68_{1.01}$	$88.14_{0.51}$	$73.21_{1.72}$	$70.20_{1.39}$	$62.16_{2.01}$	$21.89_{1.86}$	$14.89_{1.98}$	$11.72_{2.42}$
$\mathrm{A^3BL}$	$94.22_{0.80}$	$91.89_{0.33}$	$88.87_{2.32}$	$78.05_{1.13}$	$62.61_{1.89}$	$53.38_{1.98}$	$26.29_{2.71}$	$8.78_{8.65}$	$3.45_{4.68}$
DPRP-ABL	$94.61_{0.27}$	$92.06_{1.01}$	$89.34_{0.78}$	$\mathbf{81.49_{0.05}}$	$\mathbf{73.60_{1.40}}$	$\mathbf{64.68_{0.75}}$	$15.88_{4.50}$	$\mathbf{15.15_{2.92}}$	$\mathbf{14.97_{1.17}}$
NeurASP	$380.82_{185.63}$	N/A	N/A	$2185.09_{720.96}$	N/A	N/A	$760.10_{54.72}$	N/A	N/A
DeepStochLog	$57.43_{2.29}$	$173.78_{5.67}$	$282.08_{35.59}$	$50.56_{9.31}$	$146.64_{5.12}$	$272.99_{31.22}$	$49.52_{0.80}$	$304.46_{73.96}$	$210.89_{0.84}$
ImbNSL	$26.85_{0.96}$	$297.27_{32.68}$	$5608.02_{40.18}$	$33.76_{2.31}$	$340.76_{31.25}$	$4765.50_{98.21}$	$72.04_{3.16}$	$514.00_{4.48}$	$5517.52_{37.26}$
$\mathrm{ABL_{ham}}$	$3.52_{0.57}$	$15.35_{12.48}$	$\mathbf{16.35_{0.16}}$	$28.74_{19.64}$	$40.43_{2.98}$	$41.53_{1.84}$	$35.88_{18.64}$	$40.20_{10.73}$	$41.67_{16.26}$
$\mathrm{ABL_{con}}$	$3.92_{0.65}$	$16.55_{0.41}$	$17.11_{0.53}$	$24.44_{2.36}$	$18.78_{1.34}$	$34.79_{0.65}$	$17.42_{12.39}$	$29.20_{13.55}$	$38.39_{6.17}$
ABLSim	$38.50_{0.24}$	$477.34_{1.02}$	$5469.53_{105.38}$	$44.94_{1.42}$	$493.39_{4.73}$	$3502.30_{30.90}$	$49.90_{4.24}$	$512.60_{8.69}$	$3509.62_{21.67}$
$\mathrm{A^3BL}$	$7.48_{5.51}$	$15.31_{2.06}$	$30.69_{0.45}$	$\mathbf{9.65_{0.07}}$	$19.73_{5.32}$	$89.07_{8.33}$	$12.20_{0.24}$	$31.73_{1.93}$	$48.73_{2.54}$
DPRP-ABL	$\mathbf{2.58_{0.02}}$	$\mathbf{15.21_{1.77}}$	$27.68_{0.76}$	$10.19_{0.00}$	$\mathbf{18.16_{5.40}}$	$\mathbf{34.42_{2.97}}$	$\mathbf{10.78_{2.94}}$	$\mathbf{28.86_{14.00}}$	$\mathbf{37.37_{3.15}}$

Then, Fig. 2 shows comparisons at the digit concept level on the SVHN dataset, and more results are in Appendix I.1, from which we have the following observations:

- In every sequence length and most metrics, DPRP-ABL achieves the best performance in concept labeling. No other method rivals DPRP-ABL on both RECM and PREM, indicating that its recall improvement is accompanied by at least a stable level of precision.
- The performance discrepancy among methods increases with d. From $d = 1$ to $d = 3$, the across-method range widens, such as from 0.2285 to 0.5538 for ACC and from 0.1521 to 0.4157 for G-mean. It is consistent with the intensifying candidate competition in longer sequences. By restricting $\mathcal{S}_t \subseteq \mathcal{S}$, DPRP-ABL can perform better at concept labeling.

To investigate the digit recognition ability of the perception model, DPRP-ABL is compared with two fully supervised references trained using ground-truth digit labels. SUP uses cross-entropy loss (Mao et al., 2023), which is the same as our DPRP-ABL, whereas SUP-F uses focal loss to emphasize challenging examples (Lin et al., 2020). SUP and SUP-F are included as diagnostic references rather than supervision-matched baselines, because DPRP-ABL does not use digit labels as optimization signals.

[2]The task-specific construction is detailed in Appendix E.1.

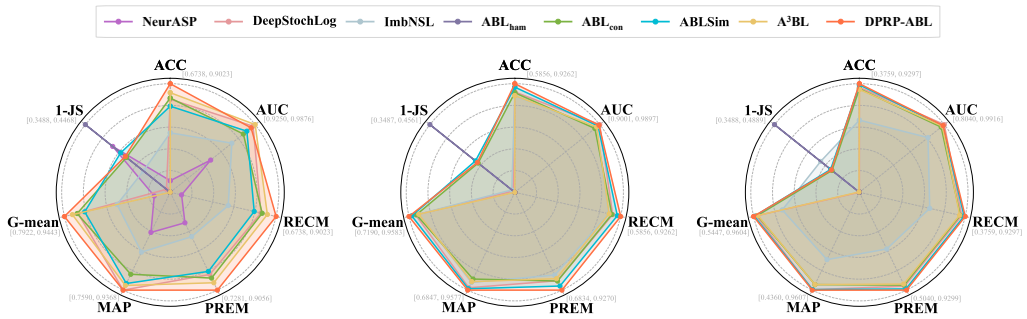

(a) $d = 1$ (b) $d = 2$ (c) $d = 3$

Figure 2: Comparison of digit concept labeling performance on the SVHN dataset. The missing item indicates that the result is not available within 30 hours.

As presented in Tabel 2, DPRP-ABL achieves a higher AUC than SUP on the SVHN dataset with both $d = 2$ and $d = 3$, as well as on CIFAR and MNIST with $d = 3$. SUP-F achieves the highest AUC in every setting. On the CIFAR dataset, the gap between DPRP-ABL and SUP-F narrows from 0.1783 at $d = 1$ to 0.0319 at $d = 3$. These results indicate that DPRP-ABL can approach fully supervised concept recognition in several challenging settings.

Table 2: Comparison "$\mathrm{Avg}_{\mathrm{std}}$" of AUC among DPRP-ABL and methods training with full supervision.

Methods	MNIST			SVHN			CIFAR		
	$d=1$	$d=2$	$d=3$	$d=1$	$d=2$	$d=3$	$d=1$	$d=2$	$d=3$
DPRP-ABL	$.9987_{.0003}$	$.9990_{.0006}$	$.9994_{.0002}$	$.9846_{.0005}$	$.9889_{.0010}$	$.9890_{.0009}$	$.7622_{.0391}$	$.8865_{.0161}$	$.9251_{.0084}$
SUP	$.9989_{.0002}$	$.9993_{.0001}$	$.9994_{.0002}$	$.9891_{.0004}$	$.9840_{.0009}$	$.9869_{.0007}$	$.8602_{.0101}$	$.8969_{.0099}$	$.9084_{.0048}$
SUP-F	$\mathbf{.9993}_{.0001}$	$\mathbf{.9996}_{.0001}$	$\mathbf{.9996}_{.0000}$	$\mathbf{.9915}_{.0005}$	$\mathbf{.9911}_{.0008}$	$\mathbf{.9926}_{.0003}$	$\mathbf{.9405}_{.0048}$	$\mathbf{.9503}_{.0010}$	$\mathbf{.9570}_{.0012}$

Ablation study. To assess the individual and joint contributions of the two axes in DPRP-ABL, we conduct 2×2 ablation. Specifically, S_1 denotes the progressive schedule, disabling which replaces the nested $\mathcal{KB}_t$ with the full $\mathcal{KB}$ throughout training. S_2 denotes distribution-prior-aware rectification, disabling which sets $\beta = 0$ in Eq. (4). Results on the SVHN dataset are reported in Table 3. More results are in Appendix I.1. Enabling either S_1 or S_2 improves performance to some extent over the no-axis variant at every sequence length. Combining both components achieves overall superiority. It suggests that progression over $\mathcal{KB}$ and distribution-prior-aware rectification provide distinct but complementary benefits. S_1 creates a focused abduction space to limit early candidate competition. S_2 rectifies the candidate scoring to ensure the relative fair competitiveness.

Table 3: Ablation "$\mathrm{Avg}_{\mathrm{std}}$" of two axes in DPRP-ABL on the SVHN dataset. "S_1" and "S_2" indicate the operation of progression over $\mathcal{KB}$ and distribution-prior-aware rectification, respectively. "z_ACC" and "y_ACC" represent concept-level and target-level ACC.

Length	S_1	S_2	z_ACC (↑)	AUC (↑)	RECM (↑)	PREM (↑)	MAP (↑)	G-mean (↑)	JS (↓)	y_ACC (↑)
$d=1$	✗	✗	$.8722_{.0063}$	$.9789_{.0015}$	$.8722_{.0063}$	$.8837_{.0055}$	$.9080_{.0078}$	$.9260_{.0038}$	$.6007_{.0020}$	$.7605_{.0107}$
	✓	✗	$.9014_{.0059}$	$.9842_{.0009}$	$.9014_{.0059}$	$.9054_{.0046}$	$.9364_{.0047}$	$.9437_{.0035}$	$.5996_{.0009}$	$.8130_{.0102}$
	✗	✓	$.8988_{.0085}$	$.9843_{.0023}$	$.8988_{.0085}$	$.9044_{.0060}$	$.9329_{.0119}$	$.9421_{.0051}$	$.6005_{.0018}$	$.8084_{.0150}$
	✓	✓	$\mathbf{.9023}_{.0002}$	$\mathbf{.9846}_{.0005}$	$\mathbf{.9023}_{.0002}$	$\mathbf{.9056}_{.0008}$	$\mathbf{.9368}_{.0014}$	$\mathbf{.9443}_{.0001}$	$\mathbf{.5993}_{.0001}$	$\mathbf{.8149}_{.0005}$
$d=2$	✗	✗	$.8978_{.0043}$	$.9846_{.0024}$	$.8978_{.0043}$	$.9043_{.0022}$	$.9268_{.0159}$	$.9415_{.0026}$	$.6049_{.0016}$	$.6509_{.0124}$
	✓	✗	$.9176_{.0026}$	$.9875_{.0010}$	$.9176_{.0026}$	$.9187_{.0021}$	$.9527_{.0028}$	$.9534_{.0015}$	$.6025_{.0012}$	$.7099_{.0086}$
	✗	✓	$.9145_{.0092}$	$.9870_{.0015}$	$.9145_{.0092}$	$.9158_{.0083}$	$.9506_{.0057}$	$.9516_{.0054}$	$\mathbf{.6011}_{.0016}$	$.7009_{.0283}$
	✓	✓	$\mathbf{.9262}_{.0046}$	$\mathbf{.9889}_{.0010}$	$\mathbf{.9262}_{.0046}$	$\mathbf{.9270}_{.0045}$	$\mathbf{.9577}_{.0040}$	$\mathbf{.9583}_{.0027}$	$.6039_{.0009}$	$\mathbf{.7360}_{.0140}$
$d=3$	✗	✗	$.9040_{.0058}$	$.9847_{.0019}$	$.9040_{.0058}$	$.9081_{.0051}$	$.9310_{.0066}$	$.9453_{.0034}$	$.6060_{.0028}$	$.5475_{.0211}$
	✓	✗	$.9244_{.0009}$	$.9859_{.0008}$	$.9244_{.0009}$	$.9250_{.0009}$	$.9460_{.0063}$	$.9573_{.0005}$	$.6055_{.0014}$	$.6247_{.0028}$
	✗	✓	$.9184_{.0179}$	$.9877_{.0015}$	$.9184_{.0179}$	$.9207_{.0152}$	$.9488_{.0175}$	$.9538_{.0104}$	$\mathbf{.6041}_{.0007}$	$.6020_{.0705}$
	✓	✓	$\mathbf{.9297}_{.0018}$	$\mathbf{.9890}_{.0009}$	$\mathbf{.9297}_{.0018}$	$\mathbf{.9299}_{.0019}$	$\mathbf{.9607}_{.0026}$	$\mathbf{.9604}_{.0010}$	$.6049_{.0005}$	$\mathbf{.6468}_{.0075}$

5.3 Task 2: Judicial Sentencing[3]

Performance analysis. We compare DPRP-ABL with SS-ABL and C-ABL, two ABL methods applicable to this task, under semi-supervised settings, where 10% and 50% of the training samples have attribute annotations. As presented in Table 4, DPRP-ABL achieves higher F1 and lower MSE and MAE and requires fewer tokens to converge, showing its superior performance and excellent efficiency in the real-world task.

[3]The task-specific construction is detailed in Appendix E.2.

To further investigate our ability to recognize the minority, concept-level gains are recorded in Fig. 3. As we can see, DPRP-ABL can maintain excellent performance on majority attributes and has a superior ability to recognize minority ones, thereby demonstrating its overall outstanding classification performance. As for some extreme minority attributes, such as "surrender" and "young", almost saturate at high F1 across all methods because they correspond to legally defined actions with unambiguous textual markers in judicial documents, making them distinguishable. In contrast, for minority attributes requiring contextual inference from case facts, such as "theft" and "indoor", DPRP-ABL obtains the largest gains.

Table 4: Comparison "$\mathrm{Avg}_{\mathrm{std}}$" of F1 and tokens to converge (the point where F1 reaches 99% of its final value).

Labeled %	Methods	F1 (↑)	MSE (↓)	MAE (↓)	Tokens
10%	SS-ABL	$.8862_{.0283}$	$.9958_{.1422}$	$.7910_{.0362}$	6.645M
	C-ABL	$.8935_{.0185}$	$.9313_{.0773}$	$.7835_{.0252}$	5.022M
	DPRP-ABL	$\mathbf{.9265_{.0068}}$	$\mathbf{.8479_{.0854}}$	$\mathbf{.7638_{.0322}}$	**4.840M**
50%	SS-ABL	$.9116_{.0103}$	$.8366_{.0814}$	$.7526_{.0384}$	6.645M
	C-ABL	$.9253_{.0148}$	$.8270_{.1199}$	$.7426_{.0465}$	4.671M
	DPRP-ABL	$\mathbf{.9460_{.0051}}$	$\mathbf{.7804_{.0447}}$	$\mathbf{.7397_{.0235}}$	**3.359M**

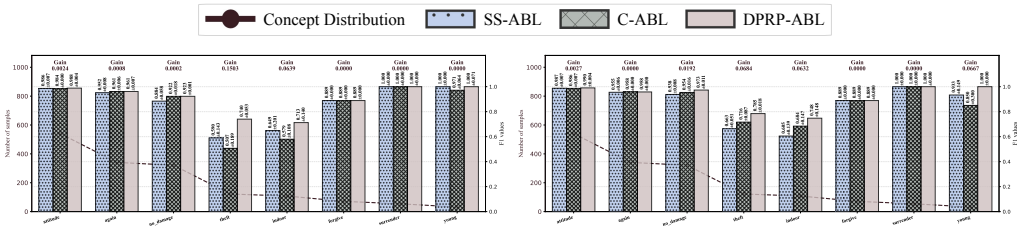

(a) 10% Labeled Data (b) 50% Labeled Data

Figure 3: Comparison "$\mathrm{Avg}_{\mathrm{std}}$" of F1 and concept-level gains on the judicial sentencing task.

Ablation study. We conduct four variants and record results in Table 5. S_2 provides larger standalone improvements in F1 than S_1. It is consistent with the task structure that rectification is conducted at the attribute level, while the sentencing constraint couples attributes within complete feasible vectors. A compensation for a rare attribute can therefore change the ranking of entire candidates. Progression provides an additional, although smaller, contribution. Enabling both axes undoubtedly yields the best overall F1, MSE, and MAE.

Table 5: Ablation "$\mathrm{Avg}_{\mathrm{std}}$" of two axes on the judicial sentencing task. "S_1" and "S_2" indicate progression over $\mathcal{KB}$ and distribution-prior-aware rectification, respectively.

Labeled %	S_1	S_2	F1 (↑)	MSE (↓)	MAE (↓)
10%	✗	✗	$.8695_{.0165}$	$.8974_{.0997}$	$.7706_{.0364}$
	✓	✗	$.8820_{.0244}$	$.8747_{.0398}$	$.7688_{.0219}$
	✗	✓	$.9097_{.0285}$	$.8855_{.0445}$	$.7657_{.0227}$
	✓	✓	$\mathbf{.9265_{.0068}}$	$\mathbf{.8479_{.0854}}$	$\mathbf{.7638_{.0322}}$
50%	✗	✗	$.9281_{.0091}$	$.8235_{.0266}$	$.7587_{.0138}$
	✓	✗	$.9313_{.0029}$	$.8084_{.0829}$	$.7526_{.0285}$
	✗	✓	$.9422_{.0102}$	$.7969_{.0374}$	$.7488_{.0154}$
	✓	✓	$\mathbf{.9460_{.0051}}$	$\mathbf{.7804_{.0447}}$	$\mathbf{.7397_{.0235}}$

6 Related Work

NeSy methods combine neural perception with logical or probabilistic reasoning (Raedt et al., 2020; Hitzler & Sarker, 2022; Manhaeve et al., 2018; Yang et al., 2020; Winters et al., 2021; Tsamoura et al., 2025). Representatively, ABL uses logical abduction to decipher symbolic concepts from outputs and knowledge and conducts reasoning (Zhou, 2019; Dai et al., 2019). Existing ABL methods improve consistency search, efficiency, ambiguity handling, and iterative revision (Huang et al., 2021; He et al., 2024; Hu et al., 2025a; Jia et al., 2025). Curriculum-based methods order training samples, as in CurABL, or gradually expand nested $\mathcal{KB}$, as in C-ABL (Wei et al., 2025; Hu et al., 2025b). A broader discussion of related work is provided in Appendix J.

7 Conclusion

We investigate ABL under concept skewness and propose DPRP-ABL. Unlike existing ABL methods that treat $\mathcal{KB}$ as a static black box or inherently assume a uniform concept distribution, DPRP-ABL combines structure-progressive $\mathcal{KB}$ scheduling with distribution-prior-aware rectification. It restricts early candidate competition and adjusts majority-favoring scoring, alleviating skewness-induced bias in the perception-abduction loop. Our analysis establishes fixed-state and conditional properties of these components. Experiments on the digit addition and judicial sentencing tasks show its superiority. By coordinating the knowledge structure with distribution information, DPRP-ABL offers a practical strategy for ABL under concept skewness.

AI USE STATEMENT

In this work, we used the generative AI tool for the polishing of the manuscript. To refine the language, improve readability, and ensure clarity, it assists in sentence rephrasing, grammar checking, and enhancing the overall flow of the text. It is important to note that the contributions of the generative AI tool were solely focused on improving the linguistic quality of the paper, with no involvement in the scientific content or data analysis. All research ideation, methodology, and analyses were developed and conducted by the authors. We have ensured that the use of the generative AI tool adheres to ethical guidelines. The authors take full responsibility for the final text, claims, code, and results.

ETHICS STATEMENT

Our study develops a machine-learning method. The experiments use established vision benchmarks and a previously released judicial-sentencing dataset (Huang et al., 2020). Judicial records and sentencing labels may reflect historical, institutional, demographic, or annotation biases. Thus, DPRP-ABL is evaluated solely as a research method and should not be used for autonomous sentencing.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our findings, we provide comprehensive technical details. The formulation and components of DPRP-ABL are described in Section 3, and the complete training procedure is provided in Algorithm 1. Theoretical results are stated in Section 4, with complete proofs in Appendix G. Task-specific instantiations are provided in Appendix E, while dataset introduction, evaluation metrics, and implementation details are documented in Appendix H. The code will be released on GitHub under the MIT License to facilitate future research.

REFERENCES

Emmanuel Abbe, Samy Bengio, Aryo Lotfi, and Kevin Rizk. Generalization on the unseen, logic reasoning and degree curriculum. In *Proceedings of the International Conference on Machine Learning*, volume 202, pp. 31–60. PMLR, 2023. URL https://proceedings.mlr.press/v202/abbe23a.html.

Samy Badreddine, Artur d'Avila Garcez, Luciano Serafini, and Michael Spranger. Logic tensor networks. *Artificial Intelligence*, 303:103649, 2022. ISSN 0004-3702. doi: https://doi.org/10.1016/j.artint.2021.103649.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the International Conference on Machine Learning*, pp. 41–48, 2009.

Matko Bošnjak, Tim Rocktäschel, Jason Naradowsky, and Sebastian Riedel. Programming with a differentiable forth interpreter, 2017. URL https://arxiv.org/abs/1605.06640.

Kai-Di Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Teng-Yu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/621461af90cadfdaf0e8d4cc25129f91-Paper.pdf.

Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.

Nitesh V. Chawla, Aleksandar Lazarevic, Lawrence O. Hall, and Kevin W. Bowyer. SMOTEBoost: Improving prediction of the minority class in boosting. In *Proceedings of the Knowledge Discovery in Databases*, pp. 107–119, Berlin, Heidelberg, 2003. Springer.

Sean X Chen. General properties and estimation of conditional Bernoulli models. *Journal of Multivariate Analysis*, 74(1):69–87, 2000. ISSN 0047-259X. doi: https://doi.org/10.1006/jmva.1999.1872. URL https://www.sciencedirect.com/science/article/pii/S0047259X99918729.

Brandon C. Colelough and William Regli. Neuro-symbolic ai in 2024: A systematic review, 2025. URL https://arxiv.org/abs/2501.05435.

Wang-Zhou Dai, Qiu-Ling Xu, Yang Yu, and Zhi-Hua Zhou. Bridging machine learning and logical reasoning by abductive learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, volume 32, pp. 2811–2822, 2019.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 4171–4186, 2019.

Richard Evans and Edward Grefenstette. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61(1):1–64, 2018. ISSN 1076-9757.

En-Hao Gao, Yu-Xuan Huang, Wen-Chao Hu, Xiao-Han Zhu, and Wang-Zhou Dai. Knowledge-enhanced historical document segmentation and recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 8409–8416, 2024. URL https://doi.org/10.1609/aaai.v38i8.28683.

Artur S d'Avila Garcez, Krysia Broda, and Dov M Gabbay. *Neural-symbolic learning systems: foundations and applications*. Springer Science & Business Media, London, 2002. ISBN 1-85233-512-2.

Artur S. d'Avila Garcez, Dov M. Gabbay, Oliver Ray, and John Woods. Abductive reasoning in neural-symbolic systems. *Topoi-An International Review Of Philosophy*, 26(1):37–49, 2007.

Alexander L. Gaunt, Marc Brockschmidt, Nate Kushman, and Daniel Tarlow. Differentiable programs with neural libraries. In *Proceedings of the International Conference on Machine Learning*, volume 70, pp. 1213–1222. PMLR, 2017. URL https://proceedings.mlr.press/v70/gaunt17a.html.

Lise Getoor and Ben Taskar. *Introduction to statistical relational learning*. The MIT Press, 2007. ISBN 9780262256230.

Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, Adrià Puigdomènech Badia, Karl Moritz Hermann, Yori Zwols, Georg Ostrovski, Adam Cain, Helen King, Christopher Summerfield, Phil Blunsom, Koray Kavukcuoglu, and Demis Hassabis. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538:471–476, 2016. ISSN 1476-4687. doi: 10.1038/nature20101. URL https://doi.org/10.1038/nature20101.

Jian Guan, Jia-Bei Liu, Jian-Guo Sun, Peng-Ming Feng, Shuai Tong, and Wen-Wu Wang. Meta metric learning for highly imbalanced aerial scene classification. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 4047–4051, Barcelona, Spain, 2020. IEEE.

Dan-Dan Guo, Zhuo Li, Mei-Xi Zheng, He Zhao, Ming-Yuan Zhou, and Hong-Yuan Zha. Learning to re-weight examples with optimal transport for imbalanced classification. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, volume 35, pp. 25517–25530, New Orleans, LA, USA, 2022. MIT Press.

Jing-Jing Guo, Qian Wang, Yi-Ting Li, and Peng-Kun Liu. Façade defects classification from imbalanced dataset using meta learning-based convolutional neural network. *Computer-Aided Civil and Infrastructure Engineering*, 35(12):1403–1418, 2020.

Angélica Guzmán-Ponce, José Salvador Sánchez, Rosa Maria Valdovinos, and José Raymundo Marcial-Romero. DBIG-US: A two-stage under-sampling algorithm to face the class imbalance problem. *Expert Systems with Applications*, 168:1–12, 2021.

Hai-Bo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, 2009. doi: 10.1109/TKDE.2008.239.

Hao-Yuan He, Hui Sun, Zheng Xie, and Ming Li. Ambiguity-aware abductive learning. In *Proceedings of the International Conference on Machine Learning*, volume 235, pp. 18019–18042. PMLR, 2024.

Kai-Ming He, Xiang-Yu Zhang, Shao-Qing Ren, and Jian Sun. Deep residual learning for image recognition, 2015. URL https://arxiv.org/abs/1512.03385.

Pascal Hitzler and Md Kamruzzaman Sarker. *Neuro-symbolic artificial intelligence: The state of the art*. IOS press, 2022. ISBN 978-1-64368-244-0.

Alfred Horn. On sentences which are true of direct unions of algebras. *Journal of Symbolic Logic*, 16(1):14–21, 1951. doi: 10.2307/2268661.

Wen-Chao Hu, Wang-Zhou Dai, Yuan Jiang, and Zhi-Hua Zhou. Efficient rectification of neuro-symbolic reasoning inconsistencies by abductive reflection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 17333–17341, 2025a. URL https://doi.org/10.1609/aaai.v39i16.33905.

Wen-Chao Hu, Qi-Jie Li, Lin-Han Jia, Cunjing Ge, Yu-Feng Li, Yuan Jiang, and Zhi-Hua Zhou. Curriculum abductive learning, 2025b. URL https://arxiv.org/abs/2505.12275.

Hui Huang, Shi-Hao Wu, Ming-Lun Gong, Daniel Cohen-Or, Uri Ascher, and Hao Zhang. Edge-aware point set resampling. *ACM Transactions on Graphics*, 32(9):1–12, 2013.

Yu-Xuan Huang, Wang-Zhou Dai, Jian Yang, Le-Wen Cai, Shaofen Cheng, Ruizhang Huang, Yu-Feng Li, and Zhi-Hua Zhou. Semi-supervised abductive learning and its application to theft judicial sentencing. In *Proceedings of the International Conference on Data Mining*, pp. 1070–1075, 2020.

Yu-Xuan Huang, Wang-Zhou Dai, Le-Wen Cai, Stephen H Muggleton, and Yuan Jiang. Fast abductive learning by similarity-based consistency optimization. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, volume 34, pp. 26574–26584, 2021.

Yu-Xuan Huang, Wang-Zhou Dai, Lian-Wen Cai, Yuan Jiang, and Zhi-Hua Zhou. Enabling knowledge refinement upon new concepts in abductive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 7928–7935, 2023. URL https://doi.org/10.1609/aaai.v37i7.25959.

Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Li-Qiang Wang, and Bo-Qing Gong. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7610–7619, Virtual, 2020. IEEE.

Lin-Han Jia, Si-Yu Han, Lan-Zhe Guo, Zhi Zhou, Zhao-Long Li, Yu-Feng Li, and Zhi-Hua Zhou. A smooth transition between induction and deduction: Fast abductive learning based on probabilistic symbol perception, 2025. URL https://arxiv.org/abs/2502.12919.

James M. Joyce. *Kullback-Leibler divergence*, pp. 720–722. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. ISBN 978-3-642-04898-2. doi: 10.1007/978-3-642-04898-2_327. URL https://doi.org/10.1007/978-3-642-04898-2_327.

Michał Koziarski, Michał Woźniak, and Bartosz Krawczyk. Combined cleaning and resampling algorithm for multi-class imbalanced data with label noise. *Knowledge-Based Systems*, 204:1–18, 2020.

Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.

Tejas D Kulkarni, Pushmeet Kohli, Joshua B Tenenbaum, and Vikash Mansinghka. Picture: A probabilistic programming language for scene perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4390–4399, 2015. doi: 10.1109/CVPR.2015.7299068.

Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015. doi: 10.1126/science.aab3050. URL https://www.science.org/doi/abs/10.1126/science.aab3050.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Stan Z. Li and Anil Jain (eds.). *Hamming distance*, pp. 668. Springer US, Boston, MA, 2009. ISBN 978-0-387-73003-5. doi: 10.1007/978-0-387-73003-5_956. URL https://doi.org/10.1007/978-0-387-73003-5_956.

Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kai-Ming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):318–327, 2020.

Jia-Hang Liu, Rui-Lei Feng, Peng Chen, Xiao-Zhen Wang, and Yue Ni. Dynamic loss reweighting method based on cumulative classification scores for long-tailed remote sensing image classification. *Remote Sensing*, 15(394):1–26, 2023.

Xu-Ying Liu, Jian-Xin Wu, and Zhi-Hua Zhou. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2): 539–550, 2008.

Zhi-Ning Liu, Wei Cao, Zhi-Feng Gao, Jiang Bian, He-Chang Chen, Yi Chang, and Tie-Yan Liu. Self-paced ensemble for highly imbalanced massive data classification. In *Proceedings of the IEEE International Conference on Data Engineering*, pp. 841–852, Dallas, Texas, USA, 2020. IEEE.

Luca Salvatore Lorello, Marco Lippi, and Stefano Melacci. The KANDY benchmark: Incremental neuro-symbolic learning and reasoning with Kandinsky patterns. *Machine Learning*, 114(7), 2025. ISSN 0885-6125. doi: 10.1007/s10994-025-06798-x. URL https://doi.org/10.1007/s10994-025-06798-x.

Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. DeepProbLog: Neural probabilistic logic programming, 2018. URL https://arxiv.org/abs/1805.10872.

Christopher D. Manning and Hinrich Schütze. *Foundations of statistical natural language processing*. MIT Press, Cambridge, MA, 1999.

An-Qi Mao, Mehryar Mohri, and Yu-Tao Zhong. Cross-entropy loss functions: Theoretical analysis and applications. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 23803–23828. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/mao23b.html.

Emanuele Marconato, Gianpaolo Bontempo, Elisa Ficarra, Simone Calderara, Andrea Passerini, and Stefano Teso. Neuro-symbolic continual learning: Knowledge, reasoning shortcuts and concept rehearsal. In *Proceedings of the International Conference on Machine Learning*, volume 202, pp. 23915–23936. PMLR, 2023. URL https://proceedings.mlr.press/v202/marconato23a.html.

Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment, 2021. URL https://arxiv.org/abs/2007.07314.

Nuno Moniz and Vitor Cerqueira. Automated imbalanced classification via meta-learning. *Expert Systems with Applications*, 178:1–14, 2021.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, volume 2011, pp. 7, 2011.

Mohsen Pirizadeh, Nafiseh Alemohammad, Mohammad Manthouri, and Meysam Pirizadeh. A new machine learning ensemble model for class imbalance problem of screening enhanced oil recovery methods. *Journal of Petroleum Science and Engineering*, 198:1–22, 2021.

Luc De Raedt and Angelika Kimmig. Probabilistic (logic) programming concepts. *Machine Learning*, 100(1):5–47, 2015. URL https://api.semanticscholar.org/CorpusID:3166992.

Luc De Raedt, Sebastijan Dumančić, Robin Manhaeve, and Giuseppe Marra. From statistical relational to neuro-symbolic artificial intelligence, 2020. URL https://arxiv.org/abs/2003.08316.

Stuart Russell. Unifying logic and probability. *Communications of the ACM*, 58(7):88–97, 2015.

Fatih Sağlam and Mehmet Ali Cengiz. A novel SMOTE-based resampling technique trough noise detection and the boosting procedure. *Expert Systems with Applications*, 200:1–12, 2022.

R.L. Solso, O.H. MacLin, and M.K. MacLin. *Cognitive psychology*. Pearson/Allyn and Bacon, 2008. ISBN 9780205521081. URL https://books.google.pl/books?id=fq9LcAAACAAJ.

Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum learning: A survey. *International Journal of Computer Vision*, 130(6):1526–1565, 2022. ISSN 0920-5691. doi: 10.1007/s11263-022-01611-x. URL https://doi.org/10.1007/s11263-022-01611-x.

Joshua B. Tenenbaum, Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022):1279–1285, 2011. doi: 10.1126/science.1192788. URL https://www.science.org/doi/abs/10.1126/science.1192788.

Geoffrey G. Towell and Jude W. Shavlik. Knowledge-based artificial neural networks. *Artificial Intelligence*, 70(1):119–165, 1994. ISSN 0004-3702. doi: https://doi.org/10.1016/0004-3702(94)90105-8. URL https://www.sciencedirect.com/science/article/pii/0004370294901058.

Efthymia Tsamoura, Kai-Fu Wang, and Dan Roth. Imbalances in neurosymbolic learning: Characterization and mitigating strategies. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, pp. 152087–152123. Curran Associates, Inc., 2025. doi: 10.52202/085713-5086. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/df46fcdb724ab86d36e5822b4d657ea9-Paper-Conference.pdf.

Jia-Chen Wang, Da-Zhen Deng, Xiao Xie, Xin-Huan Shu, Yu-Xuan Huang, Le-Wen Cai, Hui Zhang, Min-Ling Zhang, Zhi-Hua Zhou, and Ying-Cai Wu. Tac-Valuer: Knowledge-based stroke evaluation in table tennis. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, KDD'21, pp. 3688–3696, New York, NY, USA, 2021a. Association for Computing Machinery. ISBN 9781450383325. doi: 10.1145/3447548.3467104. URL https://doi.org/10.1145/3447548.3467104.

Xin Wang, Yu-Dong Chen, and Wen-Wu Zhu. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4555–4576, 2021b.

Wen-Da Wei, Xiao-Wen Yang, Jie-Jing Shao, and Ling-Zhi Guo. Curriculum abductive learning for mitigating reasoning shortcuts. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 6533–6541. International Joint Conferences on Artificial Intelligence Organization, 2025. doi: 10.24963/ijcai.2025/727. URL https://doi.org/10.24963/ijcai.2025/727.

Thomas Winters, Giuseppe Marra, Robin Manhaeve, and Luc De Raedt. DeepStochLog: Neural stochastic logic programming, 2021. URL https://arxiv.org/abs/2106.12574.

Zi-Yuan Wu, Wei Guo, Wei Zhou, Han-Jia Ye, Yuan Jiang, Hou-Xuan Li, and Zhi-Hua Zhou. Abductive multi-instance multi-label learning for periodontal disease classification with prior domain knowledge. *Medical Image Analysis*, 101:103452, 2025. ISSN 1361–8415. doi: https://doi.org/10.1016/j.media.2024.103452. URL https://www.sciencedirect.com/science/article/pii/S1361841524003797.

Xiao-Wen Yang, Jie-Jing Shao, Wei-Wei Tu, Yu-Feng Li, Wang-Zhou Dai, and Zhi-Hua Zhou. Safe abductive learning in the presence of inaccurate rules. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 16361–16369, 2024a. URL https://doi.org/10.1609/aaai.v38i15.29572.

Xiao-Wen Yang, Wen-Da Wei, Jie-Jing Shao, Yu-Feng Li, and Zhi-Hua Zhou. Analysis for abductive learning and neural-symbolic reasoning shortcuts. In *Proceedings of the International Conference on Machine Learning*, volume 235, pp. 56524–56541. PMLR, 2024b. URL https://proceedings.mlr.press/v235/yang24ac.html.

Zhun Yang, Adam Ishay, and Joohyung Lee. NeurASP: Embracing neural networks into answer set programming. In Christian Bessiere (ed.), *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 1755–1762. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/243. URL https://doi.org/10.24963/ijcai.2020/243. Main track.

L.A. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965. ISSN 0019-9958. doi: https://doi.org/10.1016/S0019-9958(65)90241-X. URL https://www.sciencedirect.com/science/article/pii/S001999586590241X.

Hao Zhou, Jun Zhang, Ting-Jin Luo, Ya-Zhou Yang, and Jun Lei. Debiased scene graph generation for dual imbalance learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4274–4288, 2023.

Zhi-Hua Zhou. Abductive learning: Towards bridging machine learning and logical reasoning. *Science China Information Sciences*, 62(7):1–3, 2019.

In the appendix, we mainly provide some details that are not explicitly covered in the paper.

A LIMITATIONS

While DPRP-ABL demonstrates substantial improvements in performance and efficiency, it has several limitations. First, its current instantiations assume access to task-specific structural metadata for offline progression construction. In our experiments, these metadata are used only to assign samples and knowledge constraints to phases and are never included as prediction targets, loss terms, or backpropagation signals. Suitable phase information may nevertheless be unavailable or costly in less curated applications. Second, distribution-prior-aware rectification relies on the estimate of concept marginals. Estimates obtained from target distributions, small labeled subsets, or revised pseudo-labels may be affected by sampling noise or error propagation. Third, DPRP-ABL assumes a largely reliable $\mathcal{KB}$ with a meaningful progressive structure. Densely entangled, noisy, incomplete, or open-world knowledge may not yield effective nested sub-bases. Thus, automatically constructing progression from observable inputs and reasoning outcomes, estimating marginals under uncertainty, and improving robustness to imperfect knowledge remain important directions for future work (He et al., 2024; Yang et al., 2024a; Huang et al., 2023).

B NOTATIONS

Notations usually used are summarized in Table 6.

C TRAINING PROCEDURE

Algorithm 1 summarizes the complete training procedure.

Algorithm 1 The training procedure of DPRP-ABL.

Input: Dataset $\mathcal{D} = \{(\mathbf{X}_n, y_n)\}_{n=1}^{N}$, knowledge base $\mathcal{KB}$, optional concept-labeled subset $\mathcal{D}_L$, difficulty stratification function ϕ, and hyperparameters β and τ.
Output: Trained perception model f_θ.

offline preparation
Construct the nested $\{\mathcal{KB}_t\}_{t=1}^{T}$ and $\{\mathcal{D}_t\}_{t=1}^{T}$ via the difficulty stratification function ϕ.
Initialize the marginal estimate $\hat{\boldsymbol{r}}$ from the available supervision.
Compute the distribution-prior-rectification weights $\boldsymbol{v}$ from $\hat{\boldsymbol{r}}$.
# *progressive training*
for phase $t = 1$ to T **do**
  repeat
    $\bar{\mathcal{D}}_t \leftarrow \varnothing$.
    for each $(\mathbf{X}, y) \in \mathcal{D}_t$ **do**
      $\hat{\mathbf{z}} \leftarrow f_\theta(\mathbf{X})$.
      $\mathcal{S}_t \leftarrow \{\mathbf{z} | \mathbf{z} \wedge \mathcal{KB}_t \models y\}$.
      $\mathcal{A}_t \leftarrow \mathsf{Generate}(\hat{\mathbf{z}}, \mathcal{S}_t)$.
      if $\mathcal{A}_t \neq \varnothing$ **then**
        $\bar{\mathbf{z}} \leftarrow \mathsf{Select}\big(\mathcal{A}_t, u_{\theta,v}(\cdot; \mathbf{X})\big)$.
        $\bar{\mathcal{D}}_t \leftarrow \bar{\mathcal{D}}_t \cup \{(\mathbf{X}, \bar{\mathbf{z}})\}$.
      end if
    end for
    Compute η_t using Eq. (6).
    $\hat{\boldsymbol{r}} \leftarrow \mathsf{UpdatePrior}(\hat{\boldsymbol{r}}, \mathcal{D}_L, \bar{\mathcal{D}}_t)$.
    Recompute $\boldsymbol{v}$ from $\hat{\boldsymbol{r}}$.
    Update θ using $\bar{\mathcal{D}}_t$ as supervision.
  until Advancement criterion for phase t is met, or termination condition is met at $t = T$.
end for

Note that UpdatePrior can be the identity mapping so that $\hat{\boldsymbol{r}}$ and $\boldsymbol{v}$ remain fixed throughout progressive training.

Table 6: Notations.

Notations	Description
$\boldsymbol{x}$	Raw observation
$\mathbf{X}$	Input sequence
$\mathcal{X}$	Input space
y	Target label
$\mathcal{Y}$	Target space
z	Concept label that is unavailable in training
$\mathbf{z}$	Concept assignment that is also can be "candidate"
$\bar{\mathbf{z}}$	Revised pseudo-labels of concepts
$\hat{\mathbf{z}}$	Model predictions
$\mathbf{z}^*$	Ground-truth concept assignment
$\mathcal{Z}$	Concept space
$\mathcal{Z}_{\text{maj}}$	Majority concept space
$\mathcal{Z}_{\text{min}}$	Minority concept space
$\mathcal{Z}_i$	Finite concept space at the i-th position
f_θ	Perception model
θ	Parameter of f
$\mathcal{KB}$	Knowledge base
$\mathcal{KB}_t$	Knowledge sub-base at the t-th phase
$\mathcal{S}$	Abduction space
$\mathcal{S}_t$	Phase-specific abduction space at the t-th phase
$\mathcal{A}$	Reference candidate set
$\mathcal{A}_t$	Effective candidate set at the t-th phase
$\mathcal{D}$	Training dataset
$\mathcal{D}_t$	Set of training samples activated at the t-th phase
$\bar{\mathcal{D}}_t$	Revised pseudo-labeled dataset at the t-th phase
N	Number of samples
N_t	Number of samples activated at the t-th phase
m	Length of the input sequence
C	Dimension of concept space
T	Number of phases of progression
κ	Structural complexity measure
ϕ	Difficulty stratification function
σ	Symbolic transition function
π	Marginal probability
$\boldsymbol{\pi}$	Marginal probability vector
$\hat{\boldsymbol{r}}$	Estimate of $\boldsymbol{\pi}$
v	Distribution-prior-aware rectification weight
$u_{\theta,v}$	Candidate utility
L	Loss function
$\mathbb{D}$	Evaluator of disagreement the candidate with the current prediction
$\mathbb{M}$	Evaluator of agreement the revised pseudo-label with the current prediction
Generate	A general competitive candidate set generating policy
Select	A general abductor
UpdatePrior	A general prior-maintenance policy

D KNOWLEDGE BASE

To accommodate more real-world applications and manage collaborative optimization in more complex reasoning tasks, we generalize the formal definition of $\mathcal{KB}$ into two distinct representation forms.

- **Discrete fixed** $\mathcal{KB}$. Aligning with traditional deterministic symbolic logic systems, its rule set is entirely discrete and statically fixed. Logical deduction and abduction are executed

within a discrete symbolic space that is highly deterministic but prone to combinatorial explosion as the reasoning chain extends.

- **Parametric continuous** $\mathcal{KB}$. Symbolic rules are relaxed or parameterized into continuous algebraic functions that encompass unknown statistical parameters. Deduction and abduction are converted into a collaborative optimization problem, where unknown parameters are updated with the parameter of the perception model, thereby balancing structural domain knowledge constraints with statistical learning robustness.

E TASK-SPECIFIC INSTANTIATIONS OF DPRP-ABL

We instantiates some components, such as Generate, Select, UpdatePrior, etc., for our two tasks. In both tasks, the same dependency chain is preserved, i.e.,

- The phase-specific knowledge base determines the abduction space $\mathcal{S}_t$.
- The candidate generator extracts the effective candidate set $\mathcal{A}_t \subseteq \mathcal{S}_t$.
- The estimated concept marginals determine the rectification weights.
- The candidate utility in Eq. (4) drives the final selection.

The two tasks differ in their concept structures, sources of concept distribution prior, and candidate selectors, providing two complementary instantiations of a unified framework.

E.1 INSTANTIATION ON d-DIGIT ADDITION

In the d-digit addition task, an input $\mathbf{X} = (\boldsymbol{x}_1, \cdots, \boldsymbol{x}_{2d})$ contains the images forming two d-digit operands, and $\mathbf{z} = (z_1, \cdots, z_{2d}) \in \{0, \cdots, 9\}^{2d}$ specifies their digit concepts. The observable target label is their sum, i.e., $y = \sigma_{\text{add}}(\mathbf{z})$. We instantiate structural complexity by the number of carry propagations, i.e., $\kappa_{\text{add}}(\mathbf{z}) = \text{carry}(\mathbf{z})$. Accordingly, abduction space at the t-th phase is

$$\mathcal{S}_t^{\text{add}}(y) = \left\{\mathbf{z} \in \{0, \cdots, 9\}^{2d} | \sigma_{\text{add}}(\mathbf{z}) = y, \kappa_{\text{add}}(\mathbf{z}) \leq b_t\right\}. \tag{17}$$

Since $b_1 \leq \cdots \leq b_T$, these spaces form a nested sequence ending in the full addition space.

Phase advancement is determined by the agreement between digit predictions and their revised pseudo-labels. For each sample, we instantiate the agreement metric in Eq. (6) as digit accuracy, i.e.,

$$\mathbb{M}_{\text{add}}\left(\hat{\mathbf{z}}_n, \bar{\mathbf{z}}_n^{(t)}\right) = \frac{1}{2d} \sum_{i=1}^{2d} \mathbb{I}\left[\hat{z}_{n,i} = \bar{z}_{n,i}^{(t)}\right]. \tag{18}$$

The abductor further favors minimum revisions of the current prediction $\hat{\mathbf{z}}$. For $\mathcal{S}_t^{\text{add}}(y)$, define

$$h_t(\hat{\mathbf{z}}, y) = \min_{\mathbf{z} \in \mathcal{S}_t^{\text{add}}(y)} d_{\text{H}}(\mathbf{z}, \hat{\mathbf{z}}), \tag{19}$$

where d_{H} is Hamming distance (Li & Jain, 2009). The effective candidate set returned by Generate is

$$\mathcal{A}_t^{\text{add}} = \left\{\mathbf{z} \in \mathcal{S}_t^{\text{add}}(y) | d_{\text{H}}(\mathbf{z}, \hat{\mathbf{z}}) = h_t(\hat{\mathbf{z}}, y)\right\}. \tag{20}$$

That is to say, $\mathcal{S}_t^{\text{add}}$ represents all candidates permitted by the phase-specific addition rules, whereas $\mathcal{A}_t^{\text{add}}$ indicates the minimum-revision candidates that actually be considered in selection. In particular, if $\hat{\mathbf{z}}$ already satisfies the phase-specific constraint, $h_t = 0$ and $\mathcal{A}_t^{\text{add}} = \{\hat{\mathbf{z}}\}$.

Since digit annotations are unavailable in this task, we estimate the digit marginal $\hat{\boldsymbol{r}} \in \Delta^9$ from the empirical distribution of the observable sum target by Eq. (7). The resulting estimate is converted into the rectification weights, i.e.,

$$v^k = \frac{(\hat{r}^k)^{-\tau}}{\frac{1}{10}\sum_{j=0}^{9}(\hat{r}^j)^{-\tau}}, \qquad k \in \{0, \cdots, 9\}. \tag{21}$$

It is estimated before training and then held fixed. Thus, there is $v_i^k = v^k$ for every digit position i.

Based on the negative log-likelihood, the general candidate utility in Eq. (4) is

$$u_{\mathrm{add}}(\mathbf{z};\mathbf{X}) = \sum_{i=1}^{2d} \log P_{\theta,i}(z_i) + \beta \sum_{i=1}^{2d} \log v_i^{z_i}. \tag{22}$$

The addition abductor uses deterministic candidate selection, i.e.,

$$\bar{\mathbf{z}}^{\mathrm{add}} \in \arg\max_{\mathbf{z}\in\mathcal{A}_t^{\mathrm{add}}} u_{\mathrm{add}}(\mathbf{z};\mathbf{X}), \tag{23}$$

corresponding to the $\gamma = 0$ case of Eq. (10). Such implementation equivalently multiplies each per-position probability by $(v^k)^\beta$, normalizes it, and selects the candidate with the largest of rectified likelihood. Deterministic tie resolution agrees with the $\gamma = 0$ formulation whenever the maximizer is unique, while the uniform tie-breaking convention in Theorem 1 is used for the general analysis of non-unique maximizers.

Theorem 1 applies to both levels of contraction in this task. First, taking $\mathcal{A} = \mathcal{S}^{\mathrm{add}}$ and $\mathcal{A}_t = \mathcal{S}_t^{\mathrm{add}}$, it shows that restricting the full addition space by the carry constraint cannot reduce the selection opportunity of the candidate retained in the current phase, provided that the model state and candidate utility are fixed. Second, within a fixed phase, $\mathcal{A}_t^{\mathrm{add}} \subseteq \mathcal{S}_t^{\mathrm{add}}$ provides an additional minimum-revision contraction, to which the same conclusion applies for every retained candidate. These are fixed-state comparisons. They do not assert that the effective minimum-revision sets are nested across training phases, because both the prediction and the minimum revision distance may change.

Theorem 2 gives the theoretical analysis of rectification. Defining

$$\widetilde{P}_{\theta,i}(k) = \frac{P_{\theta,i}(k)(v^k)^\beta}{\sum_{j=0}^{9} P_{\theta,i}(j)(v^j)^\beta}, \tag{24}$$

Eq. (22) differs from $\sum_i \log \widetilde{P}_{\theta,i}(z_i)$ only by a candidate-independent constant. Consequently, distribution-prior-rectified utility maximization and consistency optimization using the rectified digit marginals have the same maximizers. It is precisely the per-position probability transformation used by this abductor.

E.2 Instantiation on Judicial Sentencing

The judicial sentencing task has a different condition. Each case contains a textual description and a monetary amount M, while z_i in $\mathbf{z} = (z_1, \ldots, z_m) \in \{0,1\}^m$ indicates the presence or absence of legal factor at the i-th position. The observable target y is the sentence length. The knowledge base combines legal constraints with a parametric sentencing function, i.e.,

$$g_\omega(M, \mathbf{z}) = (1 + \boldsymbol{w}^\top \mathbf{z})(aM + b), \tag{25}$$

where $\omega = (\boldsymbol{w}, a, b)$ collects parameters in knowledge model. Structural complexity is instantiated as the number of active legal factors, i.e., $\kappa_{\mathrm{jud}}(\mathbf{z}) = \|\mathbf{z}\|_0$. Let $\mathrm{Valid}_\omega(M, \mathbf{z}, y)$ indicate that $\mathbf{z}$ satisfies the applicable legal constraints and is compatible with the observed sentence under the current knowledge base. The phase-specific abduction space is

$$\mathcal{S}_t^{\mathrm{jud}}(M, y) = \{\mathbf{z} \in \{0,1\}^m | \, \mathrm{Valid}_\omega(M, \mathbf{z}, y) = 1, \|\mathbf{z}\|_0 \le b_t\}. \tag{26}$$

Phase advancement is based on the F1 between predictions and revised attributes. Let $\mathcal{I}_t$ index the samples in the current phase for which valid revised pseudo-labels are obtained, and define $\widehat{\mathbf{Z}}_t = [\hat{\mathbf{z}}_n]_{n\in\mathcal{I}_t}$ and $\bar{\mathbf{Z}}_t^{(t)} = \left[\bar{\mathbf{z}}_n^{(t)}\right]_{n\in\mathcal{I}_t}$. There is

$$\eta_t^{\mathrm{jud}} = \mathrm{F1}_{\mathrm{micro}}\left(\mathrm{vec}\left(\widehat{\mathbf{Z}}_t\right), \mathrm{vec}\left(\bar{\mathbf{Z}}_t^{(t)}\right)\right), \tag{27}$$

where $\mathrm{vec}(\cdot)$ concatenates all sample-attribute entries into a binary vector.

Unlike the exhaustive enumeration used for the d-digit addition task, this abductor performs local revision around the current attribute prediction. Given a maximum revision budget R, its effective candidate set is

$$\mathcal{A}_t^{\mathrm{jud}} = \left\{\mathbf{z} \in \mathcal{S}_t^{\mathrm{jud}}(M, y) | d_{\mathrm{H}}(\mathbf{z}, \hat{\mathbf{z}}) \le R, \chi(\mathbf{z}, \mathbf{X}) = 1\right\}, \tag{28}$$

where χ represents any enabled compatibility check between the candidate and textual evidence. When this check is not used, $\chi \equiv 1$. Thus, the abduction space is determined by the legal and phase constraints, while the effective candidate set also reflects the local revision budget and optional textual filtering.

In this task, a subset $\mathcal{D}_L$ provides samples with attribute annotations. Let $\mathcal{P}_e$ be the set of examples whose revised pseudo-labels are successfully obtained in loop e. For attribute i, denote $n_{i,L}^{+}$ and $n_{i,P}^{+(e)}$ as the numbers of positive occurrences in $\mathcal{D}_L$ and $\mathcal{P}_e$, respectively. The positive prevalence used in the next loop is estimated by

$$\hat{r}_i^{(e+1)} = \frac{n_{i,L}^{+} + n_{i,P}^{+(e)} + 1}{|\mathcal{D}_L| + |\mathcal{P}_e| + 2}. \tag{29}$$

This is an attribute-wise Laplace estimator. It is initialized from $\mathcal{D}_L$ and then recomputed from the labeled subset and the latest successfully revised pseudo-labels, rather than repeatedly accumulating every historical pseudo-label.

The distribution-prior-rectification weight of the positive attribute is

$$v_{i,e}^{1} = \frac{(\hat{r}_i^{(e)})^{-\tau}}{\frac{1}{m}\sum_{j=1}^{m}(\hat{r}_j^{(e)})^{-\tau}}, \qquad v_{i,e}^{0} = 1. \tag{30}$$

Setting this weight of the negative attribute to one matches the implementation in which rectification is added only when an attribute is active. Denoting $p_i = P_\theta(z_i = 1|\mathbf{X})$, the candidate utility is

$$u_{\mathrm{jud}}(\mathbf{z}; \mathbf{X}) = \sum_{i=1}^{m} [z_i \log p_i + (1 - z_i) \log(1 - p_i)] + \beta \sum_{i=1}^{m} \log v_{i,e}^{z_i}. \tag{31}$$

Since $v_{i,e}^{0} = 1$, the second term is equivalently $\beta \sum_i z_i \log v_{i,e}^{1}$.

The abductor explores by sampling from the softmax distribution over complete feasible attribute vectors, i.e.,

$$\Pr\left(\bar{\mathbf{z}}^{\mathrm{jud}} = \mathbf{z} | \mathcal{A}_t^{\mathrm{jud}}\right) = \frac{\exp\left(u_{\mathrm{jud}}(\mathbf{z}; \mathbf{X})\right)}{\sum_{\mathbf{z}' \in \mathcal{A}_t^{\mathrm{jud}}} \exp\left(u_{\mathrm{jud}}(\mathbf{z}'; \mathbf{X})\right)}, \tag{32}$$

which is the $\gamma = 1$ instance of Eq. (10).

For fixed θ, ω, and $\boldsymbol{v}$, the phase constraint $\|\mathbf{z}\|_0 \le b_t$ removes candidates from the full judicial candidate set while retaining those compatible with the current complexity level. Therefore, whenever $\mathcal{A}_t^{\mathrm{jud}} \subseteq \mathcal{A}^{\mathrm{jud}}$ and a candidate remains in the restricted set, Theorem 1 applies with $\gamma = 1$. Its exact probability ratio in Eq. (12) quantifies how removing some competing legal-factor assignments increases the softmax probability of the retained candidate. In the same way, this is a fixed-state statement and does not claim monotonicity across training loops, during which the perception model, the parameters in knowledge model, and the estimate can all change.

Theorem 2 also specializes directly to Bernoulli model (Chen, 2000). Define the rectified positive probability

$$\widetilde{p}_{i,e} = \frac{p_i (v_{i,e}^{1})^{\beta}}{(1 - p_i)(v_{i,e}^{0})^{\beta} + p_i (v_{i,e}^{1})^{\beta}} = \mathrm{sigmoid}\left(\log \frac{p_i}{1 - p_i} + \beta \log \frac{v_{i,e}^{1}}{v_{i,e}^{0}}\right). \tag{33}$$

It is an additive shift of $\beta \log(v_{i,e}^{1}/v_{i,e}^{0})$ on the attribute logit. Under the positivity assumptions of Theorem 2, Eq. (31) differs from the Bernoulli log-likelihood constructed from $\widetilde{p}_{i,e}$ only by a candidate-independent constant. Since the selector in this task is stochastic, the theorem establishes more than equality of maximizers, i.e., the complete softmax distribution over $\mathcal{A}_t^{\mathrm{jud}}$ is unchanged whether rectification is implemented as an explicit utility or through rectified probabilities.

To sum up, our two instantiations occupy complementary points in DPRP-ABL. The first task uses categorical concepts, a target-based distribution prior estimator, fixed weights, and deterministic selection. The second task uses binary attributes, a partially supervised adaptive distribution prior estimator, dynamically updated weights, and stochastic candidate sampling. Despite these differences, Theorem 1 characterizes the effect of candidate set contraction in both settings, while Theorem 2 establishes the same equivalence between distribution-prior rectified utility and rectified perceptual probabilities.

F THEOREM REMARK

Remark 1. *Theorem 1 compares two candidate sets while holding the perception model, distribution-prior-rectification weights, and candidate utility fixed. It does not claim that the selection probability is monotonically increasing across actual training phases, because P_θ, the parametric knowledge model, and the current prediction may change during training. Besides, the conclusion is conditional on the ground-truth concept assignment being retained in the restricted set.*

Remark 2. *Theorem 3 establishes candidate-set nesting, full-constraint feasibility, and conditional ground-truth coverage. It does not imply that the abductor necessarily selects $\mathbf{z}^*$ from $\mathcal{A}_t$. Therefore, a selected pseudo-label can be consistent with the observed target while still being incorrect at the concept level. It consequently does not constitute a convergence or optimization stability guarantee.*

G THEOREM PROOF

G.1 PROOF OF THEOREM 1

Proof. We consider the cases $\gamma > 0$ and $\gamma = 0$, respectively.

Case 1: $\gamma > 0$**.** For any candidate $\mathbf{z} \in \mathcal{A}_t$, its unnormalized score $\exp\left(u_{\theta,v}(\mathbf{z};\mathbf{X})/\gamma\right)$ is identical under $\mathcal{A}$ and $\mathcal{A}_t$. Since $\mathcal{A}_t \subseteq \mathcal{A}$, there is $Z_\gamma(\mathcal{A}) = Z_\gamma(\mathcal{A}_t) + Z_\gamma(\mathcal{A} \setminus \mathcal{A}_t) \geq Z_\gamma(\mathcal{A}_t)$. Consequently,

$$p_\gamma(\mathbf{z};\mathcal{A}_t) = \frac{\exp(u_{\theta,v}(\mathbf{z};\mathbf{X})/\gamma)}{Z_\gamma(\mathcal{A}_t)} \geq \frac{\exp(u_{\theta,v}(\mathbf{z};\mathbf{X})/\gamma)}{Z_\gamma(\mathcal{A})} = p_\gamma(\mathbf{z};\mathcal{A}). \tag{34}$$

Dividing gives

$$\frac{p_\gamma(\mathbf{z};\mathcal{A}_t)}{p_\gamma(\mathbf{z};\mathcal{A})} = \frac{Z_\gamma(\mathcal{A})}{Z_\gamma(\mathcal{A}_t)} = 1 + \frac{Z_\gamma(\mathcal{A} \setminus \mathcal{A}_t)}{Z_\gamma(\mathcal{A}_t)}. \tag{35}$$

Case 2: $\gamma = 0$**.** If $\mathbf{z} \notin \mathcal{M}(\mathcal{A})$, then $p_0(\mathbf{z};\mathcal{A}) = 0$, whereas $p_0(\mathbf{z};\mathcal{A}_t) \geq 0$. Thus, the inequality holds. If $\mathbf{z} \in \mathcal{M}(\mathcal{A})$, then no candidate in $\mathcal{A}$ has utility strictly larger than $u_{\theta,v}(\mathbf{z};\mathbf{X})$. Since $\mathbf{z} \in \mathcal{A}_t \subseteq \mathcal{A}$, the candidate $\mathbf{z}$ remains a maximizer in $\mathcal{A}_t$, and $\mathcal{M}(\mathcal{A}_t) = \mathcal{M}(\mathcal{A}) \cap \mathcal{A}_t$. Thus, $|\mathcal{M}(\mathcal{A}_t)| \leq |\mathcal{M}(\mathcal{A})|$, which yields

$$p_0(\mathbf{z};\mathcal{A}_t) = \frac{1}{|\mathcal{M}(\mathcal{A}_t)|} \geq \frac{1}{|\mathcal{M}(\mathcal{A})|} = p_0(\mathbf{z};\mathcal{A}). \tag{36}$$

This completes the proof for all $\gamma \geq 0$. □

G.2 PROOF OF THEOREM 2

Proof. For every position i, there is

$$\log \widetilde{P}_{\theta,i}(z_i) = \log P_{\theta,i}(z_i) + \beta \log v_i^{z_i} - \log Z_i. \tag{37}$$

Thus, there is

$$-\sum_{i=1}^{m} \log \widetilde{P}_{\theta,i}(z_i) = -\sum_{i=1}^{m} \log P_{\theta,i}(z_i) - \beta \sum_{i=1}^{m} \log v_i^{z_i} + \sum_{i=1}^{m} \log Z_i. \tag{38}$$

Rearranging gives

$$\mathbb{D}_{\mathrm{NLL}}(\mathbf{z}, P_\theta) - \beta \sum_{i=1}^{m} \log v_i^{z_i} = -\sum_{i=1}^{m} \log \widetilde{P}_{\theta,i}(z_i) - \sum_{i=1}^{m} \log Z_i. \tag{39}$$

For a fixed sample and a fixed model state, each Z_i depends on $P_{\theta,i}$ and distribution-prior-rectification weights but not on which z_i is selected. Thus, $-\sum_i \log Z_i$ is constant over all $\mathbf{z} \in \mathcal{A}$ and does not change the minimizer.

Equivalently, defining

$$u(\mathbf{z}) = \sum_{i=1}^{m} \log P_{\theta,i}(z_i) + \beta \sum_{i=1}^{m} \log v_i^{z_i} \tag{40}$$

and

$$\widetilde{u}(\mathbf{z}) = \sum_{i=1}^{m} \log \widetilde{P}_{\theta,i}(z_i), \tag{41}$$

we have

$$u(\mathbf{z}) = \widetilde{u}(\mathbf{z}) + \sum_{i=1}^{m} \log Z_i. \tag{42}$$

The two utilities differ by the same candidate-independent constant. They have identical maximizers when $\gamma = 0$ in Eq. (10). When $\gamma > 0$, the common multiplicative factor $\exp(\sum_i \log Z_i/\gamma)$ cancels from the numerator and denominator of the softmax selection distribution. Thus, the entire selection distribution is also identical. □

G.3 Proof of Theorem 3

Proof. For any $\mathbf{z} \in \mathcal{A}_t$, there is $\mathbf{z} \in \mathcal{A}$ and $\kappa(\mathbf{z}) \leq b_t$. Since $b_t \leq b_{t+1}$, $\kappa(\mathbf{z}) \leq b_{t+1}$ and $\mathbf{z} \in \mathcal{A}_{t+1}$. Thus, $\mathcal{A}_t \subseteq \mathcal{A}_{t+1}$. Applying this argument successively yields

$$\mathcal{A}_1 \subseteq \cdots \subseteq \mathcal{A}_T \subseteq \mathcal{A} \subseteq \mathcal{S}. \tag{43}$$

If $\mathbf{z}^* \in \mathcal{A}$ and$\kappa(\mathbf{z}^*) \leq b_t$, then $\mathbf{z}^*$ satisfies the defining conditions of $\mathcal{A}_t$. Thus,

$$\mathbf{z}^* \in \mathcal{A}_t \subseteq \mathcal{A} \subseteq \mathcal{S}. \tag{44}$$

□

H Experiment Settings

H.1 Datasets

For digit addition task, we gather training sequences by representing handwritten symbols utilizing instances from benchmark datasets including MNIST (LeCun et al., 1998), SVHN (Netzer et al., 2011), and CIFAR (Krizhevsky & Hinton, 2009) to evaluate the ability of the model to recognize concepts from varying levels of perception uncertainty and different domain semantics. Specifically,

- MNIST: This dataset has 10 classes of handwritten digits, ranging from 0 to 9. Each instance is a 28×28 grayscale image.
- SVHN: This dataset has 10 classes of digits, one for each digit from 0 to 9. Each instance is a 32×32 RGB image.
- CIFAR: This dataset collects 10 classes of real-world objects (e.g., airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck). Each instance is a 32×32 RGB image.

For judicial sentencing task, the concept distribution is naturally skewed. We list the sizes of each attribute concept in the following Table 7.

Table 7: Distribution of attributes in judicial sentencing task.

Attribute	attitude	again	no_damage	theft	indoor	forgive	surrender	young
Number	541	341	321	122	108	71	54	32

H.2 Metrics

To quantify the performance of all methods, we select several metrics for evaluation.

- **Accuracy (ACC, ↑):** Proportion of correct predictions.
- **Area Under the Curve (AUC, ↑):** The average area under the Receiver Operating Characteristics (ROC) curve.

- **Recall (RECM, ↑)**: It measures the ability of the model to correctly identify positive samples.
- **Precision (PREM, ↑)**: The proportion of true positives among all samples predicted as positive.
- **F1 score (F1, ↑):** The harmonic average of recall and precision for classification.
- **Mean Average Precision (MAP, ↑)**: The average precision across all classes.
- **Geometric Mean (G-mean, ↑)**: It evaluates the equilibrium between the majority and minority.
- **Jensen-Shannon Divergence (JS, ↓)**: A symmetric and bounded version of Kullback-Leibler Divergence (Joyce, 2011), which quantifies the distance between the predicted posterior distribution and the true distribution.
- **Mean Squared Error (MSE, ↓)**: Mean squared difference between predictions and truths. It is sensitive to outliers with large errors.
- **Mean Absolute Error (MAE, ↓)**: The mean of the absolute difference between predictions and truths. It focuses on the average magnitude of error.

H.3 Implementation Details

All experiments are implemented in Python 3.10 using PyTorch 2.6 and run on a server equipped with an NVIDIA A100 80GB PCIe GPU and an Intel Xeon Platinum 8375C CPU (64 CPU cores in total), under CUDA 12.3. For each method, experiments are conducted with the optimal hyperparameter that is tuned on the validation set to ensure a fair comparison. To eliminate the effect of random factors, all experiments are independently repeated five times.

I Experiment Details and Additional Results

I.1 Task 1: d-Digit Addition

Performance analysis. To evaluate the classification performance of DPRP-ABL at the digit concept level, Fig. 4 and Fig. 5 show comparisons on the MNIST and CIFAR dataset.

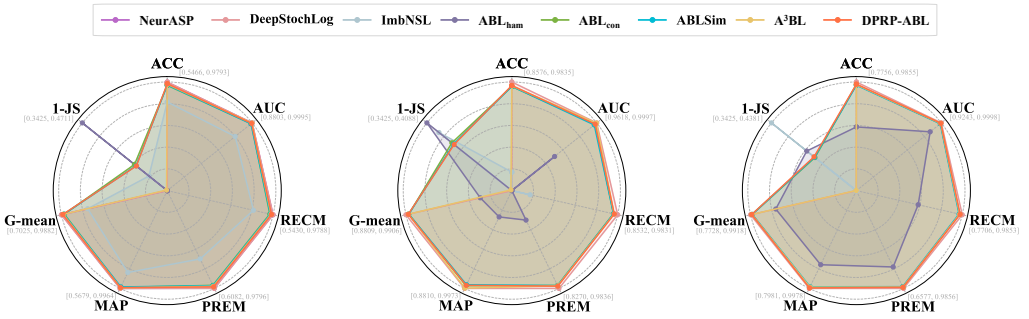

(a) $d = 1$ (b) $d = 2$ (c) $d = 3$

Figure 4: Comparison of digit concept labeling performance on the MNIST dataset. The missing item indicates that the result is not available within 30 hours.

Ablation study. To assess the individual and joint contributions of the two axes in DPRP-ABL, we conduct 2×2 ablation. Results on the MNIST and CIFAR dataset are reported in Table 8 and Table 9.

Convergence behavior. To study the optimization dynamics of DPRP-ABL, we analyze its convergence behavior on the digit addition task with varying digit length d. As shown in Fig 6, the blue and red curves denote the objective value and ACC, respectively. During training, the objective

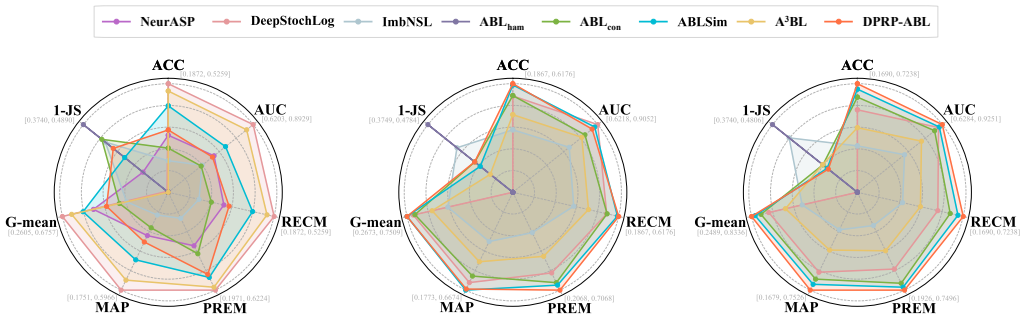

(a) $d = 1$ (b) $d = 2$ (c) $d = 3$

Figure 5: Comparison of digit concept labeling performance on the CIFAR dataset. The missing item indicates that the result is not available within 30 hours.

Table 8: Ablation "$\mathrm{Avg}_{\mathrm{std}}$" of two axes in DPRP-ABL on the MNIST dataset. "S_1" and "S_2" indicate the operation of progression over $\mathcal{KB}$ and distribution-prior-aware rectification, respectively. "z_ACC" and "y_ACC" represent concept-level and target-level ACC.

Length	S_1	S_2	z_ACC (↑)	AUC (↑)	RECM (↑)	PREM (↑)	MAP (↑)	G-mean (↑)	JS (↓)	y_ACC (↑)
$d=1$	✗	✗	$.9660_{.0089}$	$.9982_{.0017}$	$.9652_{.0090}$	$.9677_{.0081}$	$.9917_{.0065}$	$.9804_{.0052}$	$\mathbf{.6079_{.0010}}$	$.9329_{.0172}$
	✓	✗	$.9674_{.0090}$	$.9970_{.0016}$	$.9664_{.0093}$	$.9689_{.0083}$	$.9851_{.0078}$	$.9811_{.0053}$	$.6108_{.0019}$	$.9357_{.0174}$
	✗	✓	$.9705_{.0016}$	$.9980_{.0011}$	$.9697_{.0016}$	$.9714_{.0009}$	$.9916_{.0035}$	$.9830_{.0009}$	$.6097_{.0007}$	$.9419_{.0032}$
	✓	✓	$\mathbf{.9727_{.0014}}$	$\mathbf{.9987_{.0003}}$	$\mathbf{.9721_{.0016}}$	$\mathbf{.9736_{.0012}}$	$\mathbf{.9930_{.0003}}$	$\mathbf{.9844_{.0009}}$	$.6107_{.0019}$	$\mathbf{.9461_{.0027}}$
$d=2$	✗	✗	$.9783_{.0028}$	$.9990_{.0004}$	$.9779_{.0029}$	$.9788_{.0027}$	$\mathbf{.9941_{.0018}}$	$.9876_{.0016}$	$\mathbf{.6107_{.0010}}$	$.9161_{.0103}$
	✓	✗	$.9786_{.0030}$	$.9987_{.0008}$	$.9780_{.0030}$	$.9790_{.0030}$	$.9919_{.0063}$	$.9877_{.0017}$	$.6112_{.0016}$	$.9174_{.0111}$
	✗	✓	$.9750_{.0084}$	$.9979_{.0006}$	$.9744_{.0087}$	$.9755_{.0083}$	$.9903_{.0023}$	$.9856_{.0049}$	$.6112_{.0010}$	$.9040_{.0305}$
	✓	✓	$\mathbf{.9795_{.0027}}$	$\mathbf{.9990_{.0006}}$	$\mathbf{.9789_{.0028}}$	$\mathbf{.9799_{.0025}}$	$.9936_{.0044}$	$\mathbf{.9882_{.0016}}$	$.6124_{.0020}$	$\mathbf{.9206_{.0101}}$
$d=3$	✗	✗	$.9805_{.0013}$	$.9991_{.0003}$	$.9800_{.0013}$	$.9810_{.0013}$	$.9957_{.0011}$	$.9888_{.0007}$	$.6108_{.0043}$	$.8888_{.0071}$
	✓	✗	$.9813_{.0033}$	$.9986_{.0006}$	$.9809_{.0033}$	$.9815_{.0033}$	$.9934_{.0021}$	$.9893_{.0019}$	$.6109_{.0009}$	$.8927_{.0177}$
	✗	✓	$.9809_{.0037}$	$.9986_{.0004}$	$.9805_{.0037}$	$.9813_{.0037}$	$.9938_{.0011}$	$.9891_{.0021}$	$.6119_{.0021}$	$.8908_{.0197}$
	✓	✓	$\mathbf{.9814_{.0014}}$	$\mathbf{.9994_{.0002}}$	$\mathbf{.9809_{.0015}}$	$\mathbf{.9819_{.0012}}$	$\mathbf{.9967_{.0010}}$	$\mathbf{.9893_{.0008}}$	$\mathbf{.6098_{.0017}}$	$\mathbf{.8934_{.0078}}$

Table 9: Ablation "$\mathrm{Avg}_{\mathrm{std}}$" of two axes in DPRP-ABL on the CIFAR dataset. "S_1" and "S_2" indicate the operation of progression over $\mathcal{KB}$ and distribution-prior-aware rectification, respectively. "z_ACC" and "y_ACC" represent concept-level and target-level ACC.

Length	S_1	S_2	z_ACC (↑)	AUC (↑)	RECM (↑)	PREM (↑)	MAP (↑)	G-mean (↑)	JS (↓)	y_ACC (↑)
$d=1$	✗	✗	$.3250_{.0902}$	$.7259_{.0541}$	$.3250_{.0902}$	$.4642_{.1275}$	$.3281_{.0977}$	$.4521_{.0926}$	$\mathbf{.5360_{.0183}}$	$.1296_{.0498}$
	✓	✗	$.3344_{.0977}$	$.7213_{.0713}$	$.3344_{.0977}$	$.4191_{.1732}$	$.3218_{.1359}$	$.4468_{.1128}$	$.5396_{.0238}$	$.1344_{.0648}$
	✗	✓	$.3579_{.0753}$	$.7427_{.0542}$	$.3579_{.0753}$	$.4919_{.1134}$	$.3623_{.1041}$	$.4928_{.0763}$	$.5421_{.0146}$	$.1497_{.0488}$
	✓	✓	$\mathbf{.3817_{.0619}}$	$\mathbf{.7622_{.0391}}$	$\mathbf{.3817_{.0619}}$	$\mathbf{.5539_{.0804}}$	$\mathbf{.3892_{.0848}}$	$\mathbf{.5024_{.0701}}$	$.5515_{.0139}$	$\mathbf{.1588_{.0450}}$
$d=2$	✗	✗	$.5705_{.0263}$	$.8625_{.0111}$	$.5705_{.0263}$	$.6687_{.0129}$	$.5971_{.0222}$	$.7145_{.0206}$	$.5787_{.0028}$	$.1366_{.0546}$
	✓	✗	$.5943_{.0844}$	$.8706_{.0341}$	$.5943_{.0844}$	$.6898_{.0434}$	$.6253_{.0775}$	$.7301_{.0689}$	$.5778_{.0095}$	$.1403_{.0692}$
	✗	✓	$.5923_{.0652}$	$.8738_{.0258}$	$.5923_{.0652}$	$.6848_{.0266}$	$.6294_{.0573}$	$.7287_{.0559}$	$\mathbf{.5747_{.0140}}$	$.1483_{.0587}$
	✓	✓	$\mathbf{.6176_{.0316}}$	$\mathbf{.8865_{.0161}}$	$\mathbf{.6176_{.0316}}$	$\mathbf{.7068_{.0173}}$	$\mathbf{.6619_{.0343}}$	$\mathbf{.7509_{.0253}}$	$.5786_{.0046}$	$\mathbf{.1515_{.0292}}$
$d=3$	✗	✗	$.6554_{.0665}$	$.8982_{.0285}$	$.6554_{.0665}$	$.7105_{.0378}$	$.6869_{.0716}$	$.7816_{.0509}$	$.5827_{.0052}$	$.0930_{.0551}$
	✓	✗	$.6706_{.0832}$	$.9033_{.0309}$	$.6706_{.0832}$	$.7155_{.0550}$	$.6897_{.0767}$	$.7931_{.0651}$	$.5850_{.0103}$	$.1099_{.0483}$
	✗	✓	$.6615_{.0537}$	$.9013_{.0203}$	$.6615_{.0537}$	$.7075_{.0288}$	$.6905_{.0514}$	$.7881_{.0407}$	$\mathbf{.5826_{.0039}}$	$.0927_{.0404}$
	✓	✓	$\mathbf{.7238_{.0095}}$	$\mathbf{.9251_{.0084}}$	$\mathbf{.7238_{.0095}}$	$\mathbf{.7496_{.0020}}$	$\mathbf{.7526_{.0187}}$	$\mathbf{.8336_{.0073}}$	$.5895_{.0023}$	$\mathbf{.1497_{.0117}}$

value exhibits a clear downward trend, while the ACC steadily improves before reaching a stable plateau, indicating stable convergence of DPRP-ABL.

J Extended Related Work

J.1 Neuro-Symbolic and Abductive Learning.

Neuro-symbolic Learning (NeSy) combines data-driven neural networks with knowledge-driven symbolic reasoning to emulate the synergy in the human brain for improving the interpretability and reliability of learning systems (Hitzler & Sarker, 2022; Colelough & Regli, 2025). Various methods such as fuzzy logic (Zadeh, 1965; Badreddine et al., 2022), probabilistic logic programming (Raedt & Kimmig, 2015; Manhaeve et al., 2018; Yang et al., 2020; Winters et al., 2021), and

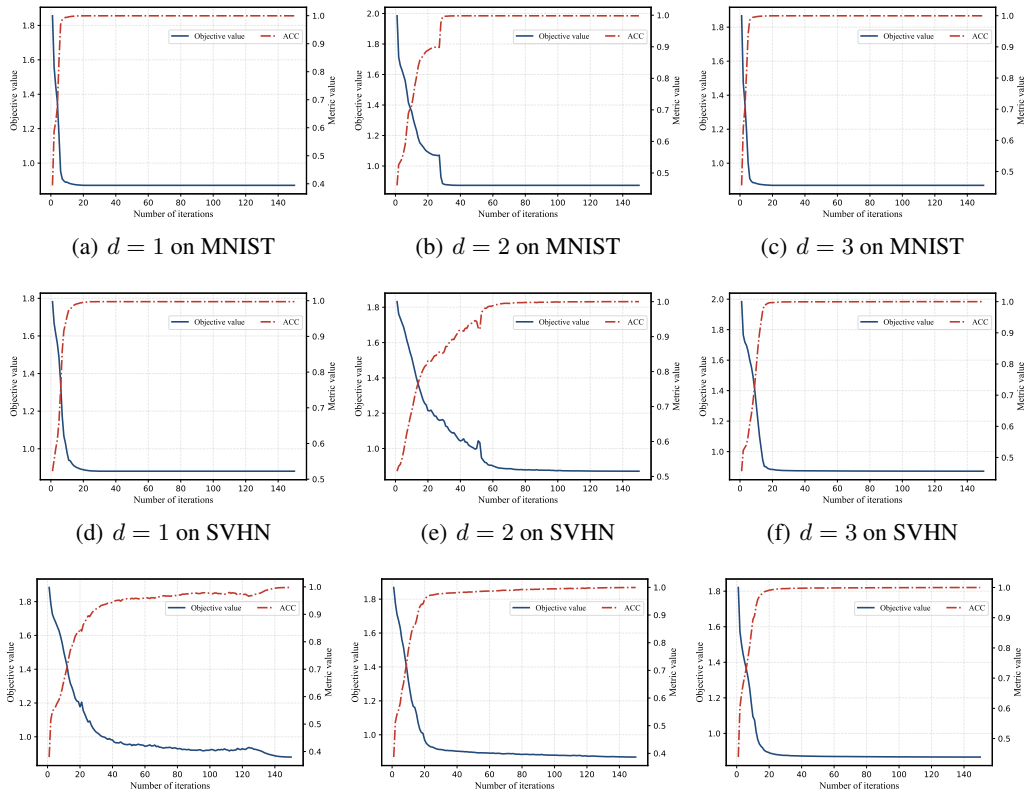

(a) $d = 1$ on MNIST (b) $d = 2$ on MNIST (c) $d = 3$ on MNIST

(d) $d = 1$ on SVHN (e) $d = 2$ on SVHN (f) $d = 3$ on SVHN

(g) $d = 1$ on CIFAR (h) $d = 2$ on CIFAR (i) $d = 3$ on CIFAR

Figure 6: Convergence behavior on the digit addition task with varying digit length d.

statistical relational learning (Getoor & Taskar, 2007; Raedt et al., 2020) have been proposed, while most of which require human-defined symbols as input (Russell, 2015). Serving as an analogy to human cognition for probabilistic reasoning with sub-symbolic primitives, probabilistic programming (Tenenbaum et al., 2011; Kulkarni et al., 2015; Lake et al., 2015) is presented, yet it relies on the assumption that the correspondence between sub-symbolic primitives and symbolic representations used in programming has already existed. Additionally, some research translates logical programs into neural networks (Towell & Shavlik, 1994; Garcez et al., 2002; 2007) or replaces symbolic operations with differentiable functions (Graves et al., 2016; Gaunt et al., 2017; Bošnjak et al., 2017; Evans & Grefenstette, 2018). These methods provide important mechanisms for combining learning and reasoning but typically struggle to exploit structured domain knowledge or require symbolic inputs.

As a particular NeSy paradigm, abductive learning (ABL) deciphers symbolic concepts without explicit annotations by drawing knowledge-consistent supervision signals from logical abduction and then conducts reasoning (Zhou, 2019; Dai et al., 2019). In ABL, perception and reasoning are coupled in a trainable and iterative loop. Existing methods have improved the efficiency, ambiguity handling, and logical consistency of this loop. Huang et al. (Huang et al., 2021) have proposed ABLSim that accelerates abduction by similarity-based consistency optimization. A^{3}BL (He et al., 2024) have been proposed to address the ambiguity problem in ABL by developing a candidate evaluation strategy that accounts for the full abduction space. Hu et al. (Hu et al., 2025a) have examined inconsistency rectification by abductive reflection. Jia et al. (Jia et al., 2025) have achieved smooth probabilistic transitions between inductive and deductive reasoning. Nowadays, ABL has been applied in various domains, including judicial sentencing (Huang et al., 2020), table tennis stroke evaluation (Wang et al., 2021a), knowledge-enhanced document understanding (Gao et al., 2024), and medical diagnosis (Wu et al., 2025).

J.2 Curriculum Learning in NeSy Reasoning.

Curriculum learning organizes training from simpler to more complex examples and can improve convergence and generalization (Bengio et al., 2009; Wang et al., 2021b; Soviany et al., 2022). In NeSy, curriculum strategies have primarily been applied at the data level (Marconato et al., 2023; Abbe et al., 2023; Lorello et al., 2025). Recently, curriculum strategies have been brought into the ABL framework. Wei et al. (Wei et al., 2025) have proposed CurABL, which orders training samples to mitigate reasoning shortcuts (Yang et al., 2024b). C-ABL (Hu et al., 2025b) instead introduces the curriculum over the knowledge base, progressively expanding sub-bases of increasing structural complexity to reduce early abduction space and stabilize training.

J.3 Imbalanced Learning in NeSy Setting.

Imbalanced learning addresses settings in which the label distribution is highly non-uniform. During training, the model tends to learn the characteristics and patterns of majority samples while minority classes receive insufficient supervision, resulting in the bias (He & Garcia, 2009). An increasing number of methods have been proposed to adjust label distributions or decision boundaries and enhance model ability to recognize minority classes. Distribution-prior-aware margin losses (Cao et al., 2019) have been proposed that assign larger decision margins to minority classes proportional to their probabilities. Menon et al. (Menon et al., 2021) have unified several rectification strategies under a logit adjustment framework to compensate for bias. There are other methods such as re-sampling methods (Chawla et al., 2002; Liu et al., 2008; Huang et al., 2013; Koziarski et al., 2020; Guzmán-Ponce et al., 2021; Sağlam & Cengiz, 2022), reweighting methods (Menon et al., 2021; Guo et al., 2022; Liu et al., 2023; Zhou et al., 2023), imbalanced meta-learning methods (Jamal et al., 2020; Guan et al., 2020; Guo et al., 2020), imbalanced ensemble learning methods (Chawla et al., 2003; Liu et al., 2020; Pirizadeh et al., 2021), etc.

In NeSy, the imbalance problem takes a qualitatively different and more complex form. The symbolic transition function introduces additional pathways through which bias is propagated and amplified. Tsamoura et al. (Tsamoura et al., 2025) have provided a formal characterization of imbalance in NeSy, demonstrating that the symbolic transition function can induce class-specific risk differences even under a uniform concept label distribution, which sharply distinguishes imbalance in NeSy from standard supervised learning. They have proposed optimal transport-based score adjustment to rectify the induced bias. Furthermore, such an imbalance issue in ABL has not yet been characterized and resolved.