

# Learning Concept-Driven Logical Rules for Interpretable and Generalizable Medical Image Classification

Yibo Gao<sup>1</sup>, Hangqi Zhou<sup>1</sup>, Zheyao Gao<sup>2</sup>, Bomin Wang<sup>1</sup>,  
Shangqi Gao<sup>3</sup>, Sihan Wang<sup>1</sup>, and Xiahai Zhuang<sup>1,✉</sup>

<sup>1</sup> Fudan University, Shanghai, China

<sup>2</sup> The Chinese University of Hong Kong, Hong Kong, China

<sup>3</sup> University of Cambridge, Cambridge, UK

ybgao22@m.fudan.edu.cn, zzh@fudan.edu.cn

<https://zmiclalab.github.io>

**Abstract.** The pursuit of decision safety in clinical applications highlights the potential of transparent methods in medical imaging. While concept-based models offer local concept explanations (instance-level), they often neglect the global decision logic (dataset-level). Moreover, these models often suffer from *concept leakage*, where unintended information within soft concept representations undermines both interpretability and generalizability. To address these limitations, we propose **Concept Rule Learner** (CRL), a novel framework to learn Boolean logical rules from binary visual concepts. CRL employs logical layers to capture concept correlations and extract clinically meaningful rules, thereby providing both local and global interpretability. The results from two tasks demonstrate that CRL achieves competitive performance with existing interpretable methods while improving generalizability to out-of-distribution data. The code of our work is available at <https://github.com/obiyoag/crl>.

**Keywords:** Explainable-AI · Concept Learning · Rule-based Model.

## 1 Introduction

Deep learning models, especially those operating as black boxes, have shown great promise in medical imaging applications [11,14,13]. However, the high standards of trust and accountability required in healthcare have spurred growing interest in transparent models [24,20], where "*explainability*" and "*logic*" have been emphasized as two aspects by the FDA principles [23]. "*Logic*" refers to the decision rules underlying model predictions, similar to clinical practice, where medical professionals assess symptoms and make decisions based on established clinical guidelines or rules. While recent research increasingly focuses on "*explainability*" via concept explanations, less attention has been paid to "*logic*" rules for medical imaging applications.

---

✉: Xiahai Zhuang is the corresponding author.

Concept Bottleneck Models (CBMs) [12] are a prevalent framework for providing concept explanations. In CBMs, a concept predictor generates explainable concepts, which are then used by a label predictor to make final predictions. Based on CBMs, Concept Embedding Models (CEMs) [7] enhance predictive performance by employing high-dimensional concept embeddings. However, these models only provide explicit local explanations by focusing on individual predictions, detailing how predicted concepts influence each decision [17,27]. *They offer insights for particular instances but may not fully capture the overall decision logic across the entire medical dataset.* For more comprehensive interpretability, integrating both local concept explanations and global logical rules is essential [28,21] for transparent medical decision-making. A recent work, named Deep Concept Reasoner (DCR) [2], explores to extract syntactic logical rules from concept embeddings. However, because DCR relies on fuzzy logic and concept embeddings, its decision-making process is less transparent.

Moreover, the above concept-based models suffer from *concept leakage* [15], where the label predictors inadvertently exploit unintended image information from soft concepts (*i.e.*, probabilities or embeddings). *The leakage compromises both interpretability and generalizability* [16]. For interpretability, predictions are influenced not only by the intended concepts but also by image information encoded within the soft concept representations. The concept predictor no longer needs to faithfully predict the concepts for accurate label predictions. Regarding generalizability, reliance on leaked information may lead to overfitting, rendering the models less robust to distribution shifts. Hard CBMs, which use binary concepts (*i.e.*, 0 and 1) for the label predictor, reduce leakage as these binary values inherently carry less extraneous information. However, they often exhibit poor performance as they predict concepts independently, neglecting the correlations between concepts.

To address the above limitations and achieve transparent medical decision-making, we propose **Concept Rule Learner** (CRL), a framework to learn Boolean logical rules from medical data. Inspired by prior works on neuro-symbolic learning [6,25,26], CRL extends logical layers to capture correlations between binary visual concepts, mitigating the issue of concept leakage. Each logical layer comprises a conjunction layer and a disjunction layer, which perform AND and OR operations, respectively. The connections across these layers generate flexible decision rules, and the final prediction for an input image is derived by aggregating the contributions of triggered rules through linear weights. By making decisions based on domain-invariant logical rules, CRL not only incorporates global interpretability, but also improves generalizability to unseen domains. Our main contributions can be summarized as follows:

- We propose CRL, a novel framework that learns Boolean logical rules from binary visual concepts to model concept correlations, while mitigating the issue of concept leakage.
- CRL not only delivers concept explanations for individual predictions but also extracts decision rules for the entire datasets, thereby unifying local and global interpretations.

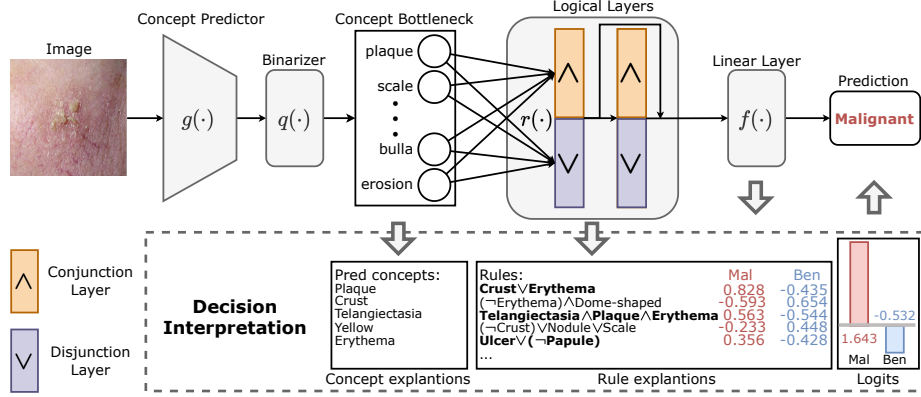


Fig. 1: The architecture of Concept Rule Learner (CRL). CRL is a composition of four functions ( $f \circ r \circ q \circ g$ ), where  $g(\cdot)$  represents the concept predictor,  $r(\cdot)$  denotes a stack of logical layers,  $q(\cdot)$  means discretization and  $f(\cdot)$  corresponds to a linear layer. The decision interpretation is presented in the dashed box, illustrating how concepts and rules contribute to the final prediction.

- We evaluate the effectiveness of the proposed method on two medical image classification tasks. The experimental results demonstrate that our approach could extract meaningful concept logical rules and exhibits superior generalizability to the unseen dataset.

## 2 Method

### 2.1 Model Architecture

Given an image  $\mathbf{x} \in \mathbb{R}^n$  annotated with a task label  $y \in \{1, \dots, L\}$  and  $K$  concept labels  $\mathbf{c} \in \{0, 1\}^K$ , our objective is to develop a **Concept Rule Learner** (CRL) capable of performing medical image classification through concept-driven logical rules. As illustrated in Fig. 1, the CRL framework is a **quadruple of functions** ( $g, q, r, f$ ), whose composition ( $f \circ r \circ q \circ g$ ) predicts the label based on the derived concepts and rules. **The first function**  $g: \mathbb{R}^n \rightarrow [0, 1]^K$ , referred to as the concept predictor, learns a mapping from the input image  $\mathbf{x}$  to a set of concept activations  $\hat{\mathbf{c}} = g(\mathbf{x}) \in [0, 1]^K$ , where  $\hat{c}_i$  represents the probability supporting the presence of the  $i$ -th concept. **The second function**  $q: [0, 1]^K \rightarrow \{0, 1\}^K$  serves as a binarizer that converts the concept activations into binary values. **The third function**  $r: \{0, 1\}^K \rightarrow \{0, 1\}^R$ , consists of a series of logical layers, mapping the activated concepts to a set of rule activations  $\mathbf{r} = r(q(\hat{\mathbf{c}})) \in \{0, 1\}^R$ , where  $R$  denotes the total number of extracted logical rules. A rule activation  $r_i = 1$  indicates that the  $i$ -th rule is activated for the input image  $\mathbf{x}$  and  $r_i = 0$  indicates otherwise. **The fourth function**  $f: \{0, 1\}^R \rightarrow \mathbb{R}^L$ , represents the final linear layer, which learns the mapping from the activated logical rules to the task logits  $\hat{\mathbf{y}} = f(\hat{\mathbf{r}}) \in \mathbb{R}^L$ .

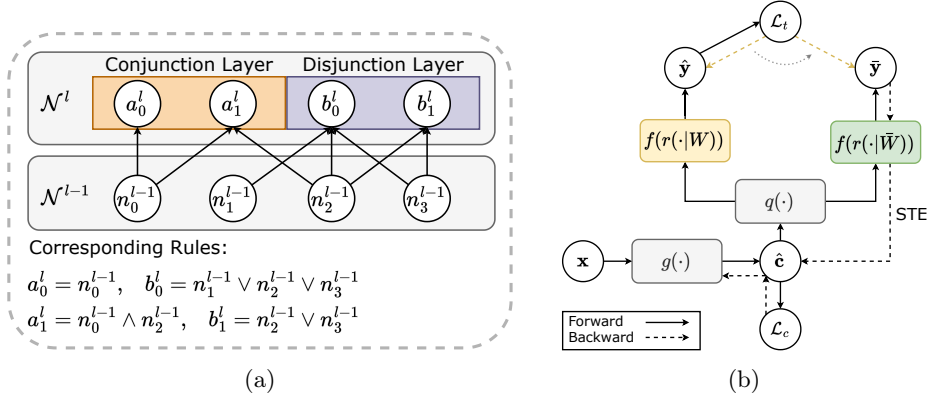


Fig. 2: (a) An example of two adjacent logical layers, where the directed arrows indicates the presence of the connections between nodes. The corresponding rules could be derived by analyzing the connections. (b) The computational graph of CRL. Arrows with solid lines represent forward pass while arrows with dashed lines represent backpropagation.

## 2.2 Logical Operation Modelling

To learn logical rules from concept activations, the function  $r(\cdot)$  adopts a series of logical layers to model the concept-based rules with Boolean logic. Let  $\mathcal{N}^l$  denote the  $l$ -th logical layer, where  $n_j^l$  represents the  $j$ -th node within the layer. The output of layer  $\mathcal{N}^l$  is a vector containing the values of all nodes, denoted as  $\mathbf{n}^l$ . As shown in Fig. 2a, each logical layer  $\mathcal{N}^l$  consists of a conjunction layer and a disjunction layer, denoted by  $\mathcal{A}^l$  and  $\mathcal{B}^l$ , respectively. The  $i$ -th node in the conjunction layer, denoted as  $a_i^l$ , represents the conjunction (AND operation) of nodes from the preceding layer connected to it. Conversely, the  $i$ -th node in the disjunction layer, denoted as  $b_i^l$ , encapsulates the disjunction (OR operation) of its connected predecessors. Formally, the nodes in the conjunction and disjunction layers are defined as:

$$a_i^{(l)} = \bigwedge_{W_{i,j}^{(l,0)}=1} n_j^{l-1}, \quad b_i^l = \bigvee_{W_{i,j}^{(l,1)}=1} n_j^{l-1}, \quad (1)$$

where  $W^{(l,0)}$  denotes the adjacency matrix of the conjunction layer  $\mathcal{A}^l$  and the previous layer  $\mathcal{N}^{l-1}$ , with  $W_{i,j}^{(l,0)} \in \{0, 1\}$ . Specifically,  $W_{i,j}^{(l,0)} = 1$  indicates the presence of an edge connecting  $a_i^l$  to  $n_j^{l-1}$ , while  $W_{i,j}^{(l,0)} = 0$  indicates the absence. Similarly,  $W^{(l,1)}$  is the adjacency matrix of the disjunction layer  $\mathcal{B}^l$  and  $\mathcal{N}^{l-1}$ . These adjacency matrices are treated as the weight matrices for the logical layers, analogous to weight matrices in neural networks. The output of the  $l$ -th layer is the concatenation of the outputs of the conjunction and disjunction layers, i.e.,  $\mathbf{n}^l = \mathbf{a}^l \oplus \mathbf{b}^l$ , where  $\mathbf{a}^l$  and  $\mathbf{b}^l$  are the outputs of  $\mathcal{A}^l$  and  $\mathcal{B}^l$  respectively.

Fig. 2a illustrates an example of two adjacent logical layers, where the directed arrows indicates the presence of the connections between nodes. By examining the weights  $W^{(l,0)}$  and  $W^{(l,1)}$ , we could derive the corresponding rules in both conjunctive and disjunctive normal forms.

### 2.3 Training Paradigm

While the logical layers are capable of expressing Boolean operations, their non-differentiable structure makes CRL challenging to optimize. To address this issue, we introduce continuous logical layers in training. These layers are differentiable and share the same parameters as the discrete counterparts, enabling end-to-end optimization while preserving the interpretability of the original design.

Let  $\bar{W}^{(l,0)}, \bar{W}^{(l,1)} \in [0, 1]$  denote the continuous weight matrices of conjunction and disjunction layers. To make Eq. (1) differentiable, we leverage the logical activation functions introduced by [19]:

$$Conj(\mathbf{n}, \bar{W}_i) = P\left(\prod_{j=1}^N F_c(n_j, \bar{W}_{i,j})\right), \quad Disj(\mathbf{n}, \bar{W}_i) = 1 - P\left(\prod_{j=1}^N F_d(n_j, \bar{W}_{i,j})\right),$$

where  $N$  is the number of nodes,  $F_c(n, w) = 1 - w(1 - n)$  and  $F_d(n, w) = 1 - n \cdot w$ . If  $\mathbf{n}$  and  $W_i$  are both binary vectors, then  $Conj(\mathbf{n}, W_i) = \bigwedge_{W_{i,j}=1} \mathbf{n}_j$  and  $Disj(\mathbf{n}, W_i) = \bigvee_{W_{i,j}=1} \mathbf{n}_j$ .  $P(x) = 1/(1 - \log x)$  is a projection function to prevent gradients vanishing. After using continuous weights and logical activation functions, the nodes in continuous logical layers are defined as follows:

$$\bar{a}_i^l = Conj(\bar{\mathbf{n}}^{l-1}, \bar{W}_i^{(l,0)}), \quad \bar{b}_i^l = Disj(\bar{\mathbf{n}}^{l-1}, \bar{W}_i^{(l,1)}).$$

By employing continuous logical layers, CRL facilitates end-to-end training. The computational graph of CRL is illustrated in Fig. 2b. As shown, the concept predictor generates concept activations  $\hat{\mathbf{c}} = g(\mathbf{x})$  from input images. The binarized concept activations are passed to both the discrete and continuous logical layers to generate the task predictions given by  $\hat{\mathbf{y}} = f(r(q(\hat{\mathbf{c}}))|W)$  and  $\bar{\mathbf{y}} = f(r(q(\hat{\mathbf{c}}))|\bar{W})$ . The objective function is defined as:

$$\mathcal{L} = \mathcal{L}_t(\hat{\mathbf{y}}, y) + \mathcal{L}_c(\hat{\mathbf{c}}, \mathbf{c}) + \lambda \|\bar{W}\|_2, \quad (2)$$

where  $\mathcal{L}_t$  is the cross-entropy loss,  $\mathcal{L}_c$  represents the mean cross-entropy loss across all training concepts and  $\lambda$  is the regularization hyperparameter that controls the complexity of the logical layers. For backpropagation, as indicated by the dotted arrow in Fig. 2b, the gradients  $\partial \mathcal{L}_t / \partial \hat{\mathbf{y}}$  are grafted onto the backward pass of  $f(r(\cdot|\bar{W}))$ . To address the non-differentiability of the binarizer, we utilize a Straight-Through Estimator (STE) [10], which approximates the gradients for the discretization.

## 2.4 Decision Interpretation

As illustrated in the dashed box of Fig. 1, CRL provides concept explanations for individual predictions as well as global logical rules for the entire dataset, thereby integrating both local and global interpretability. After training, the weights of logical layers  $r(\cdot)$  are analyzed to summarize  $R$  logical rules, where  $R$  depends on the number of nodes and the hyperparameter  $\lambda$  in Eq. (2). During inference, the concept predictor  $g(\cdot)$  first extracts concepts from images and matches them against the learned rules. Each rule is associated with a set of class-specific weights from the linear layer  $f(\cdot)$ . When a rule is triggered by a match (bold in Fig. 1), its corresponding weights are added to the overall class logits  $\hat{\mathbf{y}}$ . The final prediction is determined by the cumulative logits, which aggregate the contribution of all matched rules.

## 3 Experiments

### 3.1 Experimental Setup

**Tasks and datasets:** To evaluate the proposed method across different medical scenarios, we assess two tasks: skin disease diagnosis and white blood cell (WBC) classification. **Skin disease diagnosis:** We employ the Fitzpatrick17k (F17k) dataset [9] and the Diverse Dermatology Images (DDI) dataset [4], incorporating concept annotations from the SkinCon dataset [5]. The SkinCon dataset comprises 48 concepts annotated by board-certified dermatologists. Following the approach in [18], we focus our training and testing to images labeled as *Benign* (Ben) or *Malignant* (Mal). **WBC classification:** We utilize the PBC dataset [1], along with the concept annotations from the WBCAtt dataset [22]. The WBCAtt dataset contains 24 morphological attributes, and the classification includes five distinct classes.

**Implementation details:** We employ ResNet-34 pretrained on the ImageNet dataset, as the backbone network. The balance parameter between concept and task loss is set to 1. We use the AdamW optimizer with an initial learning rate of  $5 \times 10^{-5}$  and a weight decay of 0.01. The learning rate decays to zero following a cosine scheduler. Models are trained for 300 epochs with a batch size of 64. CRL is implemented with two logical layers, each comprising 256 nodes. The hyperparameter  $\lambda$  set to  $5 \times 10^{-6}$  to control the complexity of the rules. For evaluation, we report the average accuracy (ACC) and F1 score (F1) for both concept prediction and diagnosis tasks. For skin disease diagnosis, we perform 5-fold cross-validation. For WBC classification, we adopt the original data split from [22], conducting experiments with three different random seeds.

### 3.2 Results

**Model utility analysis:** To showcase the classification utility of the proposed method, we compare CRL with other concept-based methods on F17k and PBC

Table 1: Performance comparison on skin disease diagnosis and WBC classification tasks. **Bold** text indicates the best results, while underlined text denotes the second-best results. The symbol  $\star$  denotes methods employ binary concept values, while  $\dagger$  indicates methods with global interpretability.

Dataset	Method	Concept Metric		Diagnosis Metric	
		ACC(%)	F1(%)	ACC(%)	F1(%)
F17k	CBM [12]	91.41 $\pm$ 0.19	<b>59.44</b> $\pm$ 0.97	<u>76.37</u> $\pm$ 3.45	<u>76.30</u> $\pm$ 3.45
	align-CBM [18]	89.77 $\pm$ 0.52	58.98 $\pm$ 1.41	75.93 $\pm$ 2.39	75.83 $\pm$ 2.41
	CEM [7]	91.85 $\pm$ 0.55	<u>59.11</u> $\pm$ 2.06	76.26 $\pm$ 2.59	76.17 $\pm$ 2.52
	evi-CEM [8]	<u>92.04</u> $\pm$ 0.80	58.65 $\pm$ 1.71	<b>76.47</b> $\pm$ 1.69	<b>76.39</b> $\pm$ 1.66
	hard-CBM $\star$	86.36 $\pm$ 1.09	49.57 $\pm$ 0.48	73.51 $\pm$ 2.94	73.43 $\pm$ 2.92
	DCR $\dagger$ [2]	91.03 $\pm$ 0.80	49.40 $\pm$ 1.00	75.05 $\pm$ 2.12	74.97 $\pm$ 2.12
	<b>CRL</b> $\star\dagger$	<b>92.80</b> $\pm$ 0.47	52.39 $\pm$ 0.28	75.95 $\pm$ 3.09	75.90 $\pm$ 3.08
PBC	CBM [12]	<u>95.21</u> $\pm$ 0.10	<b>91.65</b> $\pm$ 0.15	98.93 $\pm$ 0.14	98.44 $\pm$ 0.22
	align-CBM [18]	95.01 $\pm$ 0.23	89.95 $\pm$ 0.59	99.14 $\pm$ 0.05	<u>99.25</u> $\pm$ 0.52
	CEM [8]	94.98 $\pm$ 0.34	<u>90.86</u> $\pm$ 0.69	<u>99.43</u> $\pm$ 0.12	99.23 $\pm$ 0.18
	evi-CEM [8]	94.44 $\pm$ 0.91	89.18 $\pm$ 2.09	<b>99.57</b> $\pm$ 0.03	<b>99.42</b> $\pm$ 0.03
	hard-CBM $\star$	64.53 $\pm$ 2.42	56.24 $\pm$ 2.74	98.22 $\pm$ 0.48	97.52 $\pm$ 0.73
	DCR $\dagger$ [2]	92.79 $\pm$ 0.43	82.92 $\pm$ 0.64	98.93 $\pm$ 0.32	98.47 $\pm$ 0.49
	<b>CRL</b> $\star\dagger$	<b>95.32</b> $\pm$ 0.26	90.51 $\pm$ 0.73	98.67 $\pm$ 0.25	98.03 $\pm$ 0.42

datasets. The comparative methods include CBM [12], align-CBM [18], CEM [7], evi-CEM [8] and DCR [2]. Among these, align-CBM integrates clinical knowledge to prioritize the most relevant, while evi-CEM employs evidential learning to model concept uncertainty. Hard-CBM, which serves as a baseline, only accepts binary concepts without logical layers. The comparison results are presented in Table 1. From the results, we observe that CRL achieves comparable predictive performance with other CBM variants, though it exhibits a performance trade-off compared to methods based on concept embeddings. Notably, when comparing CRL and hard-CBM, both of which utilize binary concepts, CRL significantly outperforms hard-CBM. This improvement can be attributed to the logical layers in CRL, which capture the concept correlations with logical operations.

**Model generalizability analysis:** To evaluate generalizability, we assess the out-of-domain (OOD) performance of the models on the unseen DDI dataset, using models trained on the source F17k dataset. As reported in Table 2, we can observe that concept models relying on soft concepts (probabilities or embeddings) exhibit relatively large performance drops, whereas methods employing binary concept values experience smaller declines. This suggests that binary concepts could help mitigate concept leakage. Notably, although both CRL and hard-CBM utilize binary concepts, CRL achieves a diagnostic ACC of 73.46, outperforming the second-best method by 9.21 while exhibiting the smallest performance drop. This improvement can be attributed to the domain-invariant

Table 2: Performance comparison on the unseen DDI dataset. **Bold** text indicates the best results, while underlined text denotes the second-best results. The symbol  $\star$  denotes methods employ binary concept values, while  $\dagger$  indicates methods with global interpretability.

Method	OOD Performance $\uparrow$				Performance Drop $\downarrow$			
	Concept Metric		Diagnosis Metric		Concept Metric		Diagnosis Metric	
	ACC(%)	F1(%)	ACC(%)	F1(%)	ACC(%)	F1(%)	ACC(%)	F1(%)
CBM [12]	90.59 $\pm$ 0.27	<u>53.09</u> $\pm$ 0.44	62.36 $\pm$ 4.05	57.49 $\pm$ 2.02	0.82	6.35	14.01	18.81
align-CBM [18]	89.34 $\pm$ 0.25	<u>51.47</u> $\pm$ 0.22	63.13 $\pm$ 4.01	59.66 $\pm$ 2.18	0.33	7.51	12.80	<u>16.17</u>
CEM [7]	<u>91.66</u> $\pm$ 0.26	<b>53.21</b> $\pm$ 0.57	62.39 $\pm$ 2.59	57.42 $\pm$ 2.01	<u>0.19</u>	5.90	13.87	18.75
evi-CEM [8]	91.35 $\pm$ 0.20	52.17 $\pm$ 0.64	<u>64.25</u> $\pm$ 2.04	58.82 $\pm$ 1.62	0.69	6.48	12.12	17.48
hard-CEM $\star$	84.63 $\pm$ 1.20	48.48 $\pm$ 0.28	63.77 $\pm$ 2.35	56.96 $\pm$ 1.52	1.73	<u>1.09</u>	<u>9.74</u>	16.47
DCR $\dagger$ [2]	90.97 $\pm$ 0.49	47.89 $\pm$ 0.72	62.23 $\pm$ 4.09	57.50 $\pm$ 2.35	<b>0.06</b>	1.51	12.82	17.47
<b>CRL<math>\star\dagger</math></b>	<b>92.14</b> $\pm$ 0.20	51.78 $\pm$ 0.19	<b>73.46</b> $\pm$ 2.36	<b>63.42</b> $\pm$ 1.27	0.66	<b>0.61</b>	<b>2.49</b>	<b>12.48</b>

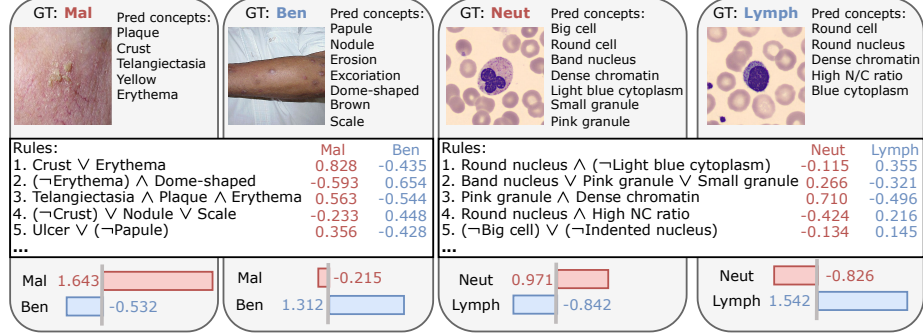


Fig. 3: An illustration of concept logical rules for both skin disease diagnosis task (left) and WBC classification (right). Note: We only present the rule weights and logits of *Neutrophil* (Neut) and *Lymphocyte* (Lymph) for WBC classification.

logical rules extracted by logical layers, which enhance robustness to distribution shifts.

**Model interpretability analysis:** To illustrate that CRL could generate meaningful logical rules, we present the rules obtained from both skin disease diagnosis and WBC classification tasks, as illustrated in Fig. 3. For skin disease diagnosis, we observe that concepts *Telangiectasia*, *Ulcer* and *Crust* are associated with *Malignant*, aligning with established clinical knowledge [3]. In case of WBC classification, the rules indicate that *Lymphocytes* are characterized by a *High NC ratio* and a *Round nucleus*, while *Neutrophils* typically exhibit *Pink granules* and *Dense chromatin*, consistent with clinical observations [22]. The case studies demonstrate that CRL can effectively extract concept-based logical rules that are clinically meaningful, offering both local concept explanations and global rule explanations for the entire medical dataset.

## 4 Conclusion

This paper introduces CRL, a framework for interpretable medical image classification that mitigates concept leakage and unifies local and global interpretability. By employing binary concepts and learnable logical layers, CRL effectively models concept correlations and extracts clinically meaningful decision rules. Experiments on two medical image classification tasks demonstrate that CRL achieves competitive performance and exhibits superior generalizability to unseen data.

**Acknowledgments.** This work was funded by the National Natural Science Foundation of China (grant No. 62372115) and the Shanghai Municipal Education Commission-Artificial Intelligence Initiative to Promote Research Paradigm Reform and Empower Disciplinary Advancement Plan (grant no. 24KXZNA13).

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Acevedo, A., Merino, A., Alf  rez, S.,   ngel Molina, Bold  , L., Rodellar, J.: A dataset of microscopic peripheral blood cell images for development of automatic recognition systems. *Data in Brief* **30**, 105474 (2020)
2. Barbiero, P., Ciravegna, G., Giannini, F., Zarlenga, M.E., Magister, L.C., Tonda, A., Lio, P., Precioso, F., Jamnik, M., Marra, G.: Interpretable neural-symbolic concept reasoning. In: *International Conference on Machine Learning* (2023)
3. Bologna, J., Schaffer, J., Cerroni, L.: *Dermatology*. Elsevier (2017)
4. Daneshjou, R., Vodrahalli, K., Novoa, R.A., Jenkins, M., Liang, W., Rotemberg, V., Ko, J., Swetter, S.M., Bailey, E.E., Gevaert, O., Zou, J., et al.: Disparities in dermatology ai performance on a diverse, curated clinical image set. *Science advances* **8**(31) (2022)
5. Daneshjou, R., Yuksekgonul, M., Cai, Z.R., Novoa, R.A., Zou, J.: SkinCon: A skin disease dataset densely annotated by domain experts for fine-grained debugging and analysis. In: *Neural Information Processing Systems* (2022)
6. Duan, X., Wang, X., Zhao, P., Shen, G., Zhu, W.: Deeplogic: Joint learning of neural perception and logical reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 4321–4334 (2023)
7. Espinosa Zarlenga, M., Barbiero, P., Ciravegna, G., Marra, G., Giannini, F., Dili-genti, M., Shams, Z., Precioso, F., Melacci, S., Weller, A., Li  , P., Jamnik, M.: Concept Embedding Models: Beyond the Accuracy-Explainability Trade-Off. In: *Advances in Neural Information Processing Systems*. vol. 35, pp. 21400–21413 (2022)
8. Gao, Y., Gao, Z., Gao, X., Liu, Y., Wang, B., Zhuang, X.: Evidential concept embedding models: Towards reliable concept explanations for skin disease diagnosis. In: *International Conference on Medical Image Computing and Computer Assisted Intervention* (2024)
9. Groh, M., Harris, C., Soenksen, L., Lau, F., Han, R., Kim, A., Koochek, A., Badri, O.: Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1820–1828 (2021)

10. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks. In: *Advances in Neural Information Processing Systems*. vol. 29 (2016)
11. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)
12. Koh, P.W., Nguyen, T., Tang, Y.S., Mussmann, S., Pierson, E., Kim, B., Liang, P.: Concept Bottleneck Models. In: *Proceedings of the 37th International Conference on Machine Learning*. pp. 5338–5348 (2020)
13. Li, C., Wong, C., Zhang, S., Usuyama, N., Liu, H., Yang, J., Naumann, T., Poon, H., Gao, J.: Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *arXiv preprint arXiv:2306.00890* (2023)
14. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**, 654 (2024)
15. Mahinpei, A., Clark, J., Lage, I., Doshi-Velez, F., Pan, W.: Promises and pitfalls of black-box concept learning models. *arXiv preprint arXiv:2106.13314* (2021)
16. Margeloiu, A., Ashman, M., Bhatt, U., Chen, Y., Jamnik, M., Weller, A.: Do concept bottleneck models learn as intended? In: *International Conference on Learning Representations workshop* (2021)
17. Oikarinen, T., Das, S., Nguyen, L.M., Weng, T.W.: Label-Free Concept Bottleneck Models. In: *International Conference on Learning Representations* (2023)
18. Pang, W., Ke, X., Tsutsui, S., Wen, B.: Integrating clinical knowledge into concept bottleneck models. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2024)
19. Payani, A., Fekri, F.: Learning algorithms via neural logic networks. *arXiv preprint arXiv:1904.01554* (2019)
20. Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* **1**(5), 206–215 (2019)
21. Schrouff, J., Baur, S., Hou, S., Mincu, D., Loreaux, E., Blanes, R., Wexler, J., Karthikesalingam, A., Kim, B.: Best of both worlds: local and global explanations with human-understandable concepts. *arXiv preprint arXiv:2106.08641* (2021)
22. Tsutsui, S., Pang, W., Wen, B.: Wbcatt: A white blood cell dataset annotated with detailed morphological attributes. In: *Advances in Neural Information Processing Systems (NeurIPS)*. (2023)
23. U.S. Food and Drug Administration: Transparency for machine learning-enabled medical devices: Guiding principles (2024)
24. van der Velden, B.H.M., Kuijf, H.J., Gilhuijs, K.G.A., Viergever, M.A.: Explainable artificial intelligence (XAI) in deep learning-based medical image analysis. *Medical Image Analysis* **79**, 102470 (2022)
25. Wang, Z., Zhang, W., Liu, N., Wang, J.: Learning interpretable rules for scalable data representation and classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(02), 1121–1133 (2024)
26. Yang, G., Song, L.: Learn to explain efficiently via neural logic inductive learning. In: *International Conference on Learning Representations* (2023)
27. Yang, Y., Panagopoulou, A., Zhou, S., Jin, D., Callison-Burch, C., Yatskar, M.: Language in a Bottle: Language Model Guided Concept Bottlenecks for Interpretable Image Classification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19187–19197 (2023)
28. Zhang, Y., Tiño, P., Leonardis, A., Tang, K.: A Survey on Neural Network Interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence* **5**(5), 726–742 (2021)