



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Integrating Clinical Knowledge into Concept Bottleneck Models

Winnie Pang, Xueyi Ke, Satoshi Tsutsui, and Bihan Wen*

Nanyang Technological University, Singapore
winnie001@e.ntu.edu.sg, bihan.wen@ntu.edu.sg

Abstract. Concept bottleneck models (CBMs), which predict human-interpretable concepts (e.g., nucleus shapes in cell images) before predicting the final output (e.g., cell type), provide insights into the decision-making processes of the model. However, training CBMs solely in a data-driven manner can introduce undesirable biases, which may compromise prediction performance, especially when the trained models are evaluated on out-of-domain images (e.g., those acquired using different devices). To mitigate this challenge, we propose integrating clinical knowledge to refine CBMs, better aligning them with clinicians’ decision-making processes. Specifically, we guide the model to prioritize the concepts that clinicians also prioritize. We validate our approach on two datasets of medical images: white blood cell and skin images. Empirical validation demonstrates that incorporating medical guidance enhances the model’s classification performance on unseen datasets with varying preparation methods, thereby increasing its real-world applicability. Our code is available at https://github.com/PangWinnie0219/align_concept_cbm.

Keywords: Explainable-AI · Concept bottleneck models

1 Introduction

The integration of deep learning in medical image analysis has greatly improved prediction performance [9,4,16]. However, the inherent “black-box” nature of these deep models often presents challenges in understanding the decision-making processes, thus limiting their acceptance in clinical settings. The Concept Bottleneck Model (CBM) [10] was proposed to address this issue by establishing causal relationships between interpretable concepts and final class predictions, thereby enhancing the explainability of the decision-making process. For example, in the classification of white blood cells (WBC), CBM facilitates the interpretation of the model’s decisions in terms of morphological concepts such as nucleus shape and granule color. Furthermore, CBM has been applied to other medical tasks [14,3,21].

However, models trained solely on data often inadvertently capture biases present in the training data, leading to diminished performance, especially under domain shifts from training data, such as changes in staining conditions or

* Bihan Wen is the corresponding author.

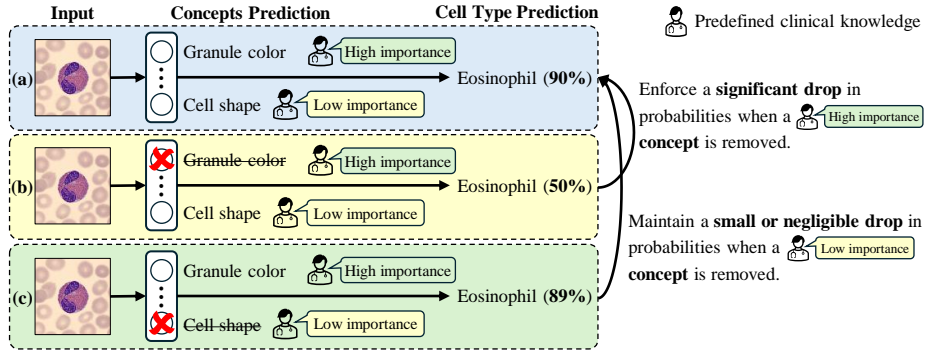


Fig. 1. We propose a method to guide concept bottleneck models (CBMs) using knowledge aligned with clinicians’ perspectives. **(a)**: CBMs predict interpretable concepts (e.g., granule color, cell shape, etc.) and then make a final prediction (e.g., eosinophil) based on them. During training, models usually do not consider the clinical importance of the concepts. Therefore, granule color and cell shape are treated equally despite granule color being a much more important factor for predicting eosinophils. **(b)**: To incorporate clinical knowledge, we enforce the CBM to exhibit a significant drop in cell type prediction probabilities when a clinically important concept is removed from the prediction. For instance, the predicted eosinophil probability should be lower when granule color, a key factor in recognizing eosinophil, is missing. **(c)**: Conversely, the cell type prediction probabilities should experience a negligible drop when a less clinically important concept is removed from the prediction. For instance, the eosinophil probability should not be affected much when cell shape, which is irrelevant to recognizing eosinophil, is missing.

microscopic devices [20]. In contrast to deep learning models, clinical experts can generally achieve reliable classification performance because they are equipped with domain knowledge. When classifying WBC types based on morphological concepts, for example, experts focus on concepts that differentiate between different cell types [2,18]. While all concepts are clinically relevant to WBCs, certain ones hold more relevance to specific cell types. Experts concentrate on these critical concepts to make decisions. This knowledge helps human experts perform better in recognizing WBCs. Motivated by the above observation, we hypothesize that CBMs can similarly benefit from domain knowledge if appropriately integrated. Then, *how to effectively incorporate domain knowledge into CBMs?* To address this, we present an approach that encourages CBMs to mirror the decision-making process of clinicians based on clinical concepts.

In this paper, we **propose a novel approach to guide the CBM to prioritize the concepts utilized by the clinical experts in decision-making**, as illustrated in Fig. 1. We first identify the importance of each concept utilized by clinical experts in cell type classification. Subsequently, we utilize a perturbation method to determine the importance of each concept from the model in predicting the cell type. We define a concept as *important* if its removal from the prediction results in a substantial variance in prediction probabilities. This

approach allows us to determine the importance of all concepts in the model prediction process. Armed with this knowledge, we align the model’s perception of concept importance with that of the clinical experts. This is achieved by enforcing a significant drop in prediction probabilities when a clinically important concept is removed from the prediction, thereby ensuring that the model makes its predictions based on concepts aligned with expert prioritization. We conduct experiments on two sets of datasets containing different medical images—WBC images and skin images—and demonstrate that the proposed method can enhance the performance of CBM through the integration of clinical knowledge.

1.1 Related Work

Although there are other CBM variants introduced to enhance classification performance or reduce annotation effort [23,15,7], to the best of our knowledge, no existing work incorporates clinical knowledge into CBMs. However, prior work has explored integrating clinical knowledge into black-box models. For instance, Yin et al. [22] proposed a regularization module to guide attention maps towards clinically interpretable features like nuclei and fat droplets in histological pattern analysis. Similarly, Zhou et al. [24] designed a selection module that mimics dermatologists’ focus on lesion features such as plaque and scale from noisy backgrounds. In addition, Manh et al. [13] introduced a multi-attribute attention network to guide the model in learning clinically relevant concepts, such as calcification and nodule shape, for predicting thyroid nodule malignancy.

While the utilization of clinical concepts itself is not a new idea, our novelty lies in the following two aspects. Firstly, while prior work utilize clinical knowledge to guide black-box models, our work focuses on interpretability and thus devises techniques to integrate clinical knowledge specifically into CBMs. Secondly, while prior work defines clinical knowledge in a class-agnostic manner (e.g., nuclei should always be prioritized for any classes), our method enables the definition of more fine-grained knowledge per class. For instance, our method can assign priority to granule color when predicting eosinophils, while emphasizing cytoplasm texture and nucleus shape when predicting monocytes, considering the absence of granules in these cells. Based on these two key differences (focus on CBMs and fine-grained knowledge representations), prior work based on attention [22,24,13] cannot be used as our baselines without substantial modifications because adding these attention mechanisms to CBMs for utilizing fine-grained knowledge is not trivial.

2 Method

Our work is built upon the CBM [10], which is designed to make final class predictions based on human-interpretable concepts. Formally, the model is abstracted as $x \mapsto \hat{c} \mapsto \hat{y}$, trained on n data points of $\{(\text{image } x^{(i)}, \text{concepts } c^{(i)}, \text{target } y^{(i)})\}_{i=1}^n$, where \hat{c} and \hat{y} denote the predicted c and y , respectively. One-hot vector $y \in \{0, 1\}^{K \times 1}$ is used to represent K classes. Moreover, to represent L concepts, we

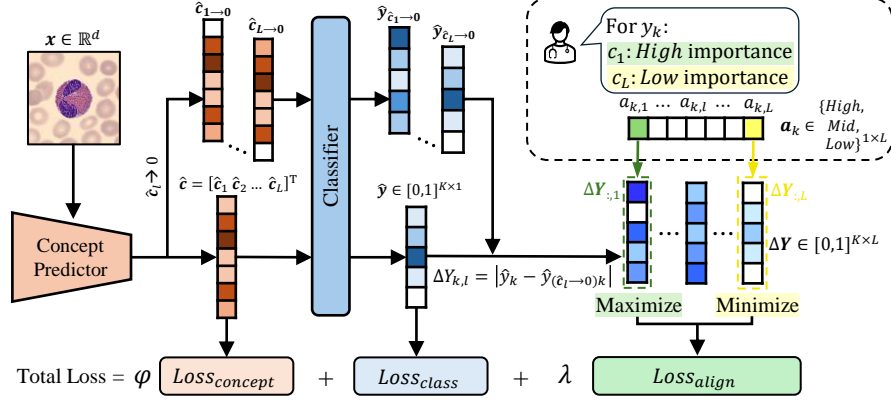


Fig. 2. Integration of Clinical Knowledge into CBM: The lower path illustrates the original CBM, where the final class prediction (\hat{y}) is based on the prediction of concepts (\hat{c}). The concept importance learned from the model (ΔY) is obtained through computing the absolute difference between the prediction probabilities when all concepts are used for prediction (\hat{y}) and when the concept is removed from prediction ($\hat{y}_{\hat{c}_l \rightarrow 0}$). We align the ΔY with the concept importance from the clinician’s perspective through $Loss_{align}$. Specifically, we enforce the model to maximize the $\Delta Y_{:,l}$ for the concept that is considered as *High* important by the clinicians, while we constrain the model to minimize the $\Delta Y_{:,l}$ for the concept that is considered *Low* important by the clinicians.

formulate c as a concatenation of probability vectors, $c = [c_1 c_2 \cdots c_L]^\top$ with corresponding prediction $\hat{c} = [\hat{c}_1 \hat{c}_2 \cdots \hat{c}_L]^\top$. The l -th concept $c_l \in \{0, 1\}^{N_l}$ represents the ground truth value out of N_l possible values. For instance, the concept of *cytoplasm-color* is the 5th concept and has the three possible values {light-blue, purple-blue, blue}, where its ground truth of *blue* is represented as $c_5 = [0, 0, 1]^\top$, while its prediction can be $\hat{c}_5 = [0.1, 0.1, 0.8]^\top$. When we *remove* this concept from the input of the classifier ($\hat{c}_5 \rightarrow 0$), we fill zeros like $[0, 0, 0]^\top$, and denote the resulting class prediction after the removal as $\hat{y}_{\hat{c}_5 \rightarrow 0}$. The architecture of the proposed method is illustrated in Fig. 2.

2.1 Align Model’s Concept Importance with Clinical Knowledge

Clinical Knowledge Representation: We rank L concepts based on a three-tier scale of importance: *High*, *Mid*, and *Low*, reflecting domain knowledge. Specifically, we assign a concept importance vector $\alpha_k \in \{High, Mid, Low\}^{1 \times L}$ per final prediction class $k \in \{1, \dots, K\}$. For example, pathologists often identify eosinophils based on the red color of granules [18]. In this case, the concept *granule-color* is marked as *High* for the class of eosinophils. If an eosinophil is assigned $k = 2$ and *granule-color* is indexed at $l = 7$, then $\alpha_{2,7} = High$.

Alignment: Our aim is to align the importance of knowledge-based concepts, as specified above, with the importance perceived by the model. To quantify

the model’s perception, we calculate the impact of each concept on the final class prediction using a perturbation method. Specifically, we first compute the class probabilities when l -th concept is removed from the prediction $\hat{y}_{\hat{c}_l \rightarrow 0}$, for $l = 1, \dots, L$. Then, we calculate the absolute difference between \hat{y} (the original predicted probability before removing a concept) and $\hat{y}_{\hat{c}_l \rightarrow 0}$, denoted as:

$$\Delta Y_{k,l} := |\hat{y}_k - \hat{y}_{(\hat{c}_l \rightarrow 0)k}| \in [0, 1]$$

where \hat{y}_k is the predicted probability for class k when all concepts are used for prediction, and $\hat{y}_{(\hat{c}_l \rightarrow 0)k}$ is the predicted probability for class k when l -th concept is removed from the input to the class prediction layer.

We consider concepts whose removal significantly alters the class prediction probability as important according to the model. Our objective is to align these important concepts with those ranked as *High* by clinical experts. To accomplish this, for concepts ranked as *High* importance by experts, each denoted by l_H , we employ the L1 loss to maximize the $\Delta Y_{:,l_H} \in [0, 1]^{K \times 1}$:

$$\mathcal{L}_{\text{Align(High)}} = \sum_{l_H} \sum_k |1 - \Delta Y_{k,l_H}|$$

Conversely, denoting a *Low* importance concept as l_L , we seek to minimize $\Delta Y_{:,l_L}$ for the concepts ranked as *Low* importance by experts:

$$\mathcal{L}_{\text{Align(Low)}} = \sum_{l_L} \sum_k |\Delta Y_{k,l_L}|$$

Our loss function is then formulated as:

$$\mathcal{L} = \varphi \mathcal{L}_c + \mathcal{L}_y + \lambda (\mathcal{L}_{\text{Align(Low)}} + \mathcal{L}_{\text{Align(High)}})$$

where \mathcal{L}_c and \mathcal{L}_y represent the cross-entropy loss functions for concept and class prediction, respectively. The hyperparameters φ and λ control the weights assigned to \mathcal{L}_c and $\mathcal{L}_{\text{Align}}$, respectively.

3 Experiments

To evaluate the effectiveness of our approach to integrate clinical knowledge into CBMs, we assess its performance on two types of medical image datasets: WBC and skin. We evaluate the classification performance on both in-domain and out-of-domain datasets.

(1) WBC datasets: We used the PBC [1] dataset with concept annotations from the WBCAtt dataset [19], following their train/val/test splits [19]. Subsequently, we trained the concept predictor ($L = 11$) and classifier ($K = 5$) for WBC classification. Performance evaluation was conducted on the PBC dataset’s testing set for in-domain data and on two other out-of-domain datasets: Scirep [12] ($n = 2019$) and RaabinWBC [11] ($n = 4339$). These out-of-domain datasets contain WBC images under various staining conditions and using different imaging devices. Fig. 3 (a) displays example images. **Collection of clinical**




Fig. 3. (a) Examples for each cell type from the in-domain dataset (PBC), as well as the out-of-domain datasets (Scirep and RaabinWBC). (b) Skin images for *Benign* and *Malignant* classes from Fitzpatrick 17k (in-domain) and DDI (out-of-domain) datasets.

knowledge: We consulted a pathologist to rank the importance of each concept from her perspective when classifying the WBC images. The pathologist was given 50 images (10 images per WBC class) to rank the importance level on a three-tier scale: *High*, *Mid*, *Low*. We averaged the importance level of each concept for each WBC class, which is used as our clinical knowledge, as shown in Supp. Fig 6(a).

(2) Skin datasets: We train a skin disease classification model ($K = 2$) using images from the Fitzpatrick 17k dataset [8] ($n = 910$) with concept annotations sourced from the SkinCon dataset [6]. Following the methodology outlined in [6], we include only the 22 concepts with at least 50 images ($L = 22$). Subsequently, we evaluate the model’s performance on two datasets: the Fitzpatrick 17k images without concept annotations serve as the in-domain dataset ($n = 3479$), and the Diverse Dermatology Images (DDI) dataset [5] ($n = 656$) serves as the out-of-domain dataset. Since DDI images contain *Benign* and *Malignant* classes only, we utilize the Fitzpatrick 17k images with these two classes for training and testing. Example of images from these datasets are shown in Fig. 3(b). **Collection of clinical knowledge:** We extracted the clinical knowledge (available within the code) from the discussion in the original paper [6].

3.1 Results

Evaluation Metrics: Due to the class imbalance in the datasets, we use the macro F1 score, which calculates the harmonic mean of precision and recall, instead of plain accuracy.

WBC Classification: We compared the performance of the CBM with (*Ours*) and without (*Baseline*) the integration of clinical knowledge across various backbone architectures.¹ As shown in Table 1, on the in-domain dataset,

¹ Refer to Supp. Table 4 and Supp. Table 5 for more details on the classifiers architectures and the implementation details, respectively.

Table 1. F1 score (mean \pm 95% CI) of the WBC type classification on cross-dataset testing, with (*Ours*) and without (*Baseline*) the integration of clinical knowledge across various backbone architectures. Generally, incorporating clinical knowledge with the model can enhance performance on out-of-domain dataset evaluations. VGG16 and ViT-B/16 are the backbones of concept predictors, while Linear, MLP(20), MLP(128) refers to the type of classifier with their details available in Supp. Table 4.

CBM Backbone+Classifier	In-domain Dataset		Out-of-domain Datasets			
	PBC		Scirep		RaabinWBC	
	Baseline	Ours	Baseline	Ours	Baseline	Ours
VGG16+Linear	99.68 \pm 0.06	99.63 \pm 0.04	68.99 \pm 3.35	70.22 \pm 0.28	51.01 \pm 2.40	54.45 \pm 5.59
VGG16+MLP(20)	99.69 \pm 0.04	99.73 \pm 0.02	66.91 \pm 3.95	70.02 \pm 6.48	52.06 \pm 5.83	53.29 \pm 4.77
VGG16+MLP(128)	99.78 \pm 0.04	99.67 \pm 0.06	66.31 \pm 3.73	70.50 \pm 2.94	54.45 \pm 3.14	58.40 \pm 5.55
ViT-B/16+Linear	99.61 \pm 0.12	99.69 \pm 0.07	75.24 \pm 5.88	77.04 \pm 1.56	41.19 \pm 10.68	47.27 \pm 6.46
ViT-B/16+MLP(20)	99.42 \pm 0.41	99.67 \pm 0.06	71.54 \pm 2.22	75.45 \pm 1.75	40.03 \pm 9.05	48.13 \pm 5.17
ViT-B/16+MLP(128)	99.72 \pm 0.02	99.63 \pm 0.04	76.39 \pm 4.54	80.15 \pm 3.19	43.64 \pm 6.67	49.75 \pm 2.39

the CBM consistently performs well across all backbones, achieving an average F1-score exceeding 99%. However, when evaluated on the out-of-domain datasets, a noticeable decline in performance is observed, especially on Raabin-WBC. The decline suggests that the model’s performance is significantly affected when tested on images prepared with different staining conditions and imaging devices (Fig. 3(a)). Notably, integrating clinical knowledge into the CBM results in an improvement in out-of-domain classification performance across all employed backbones. This improvement underscores the efficacy of guiding the model to base its predictions on concepts prioritized by domain experts, thereby enhancing its adaptability and robustness in real-world scenarios. The qualitative results are shown in Supp. Fig 5. We also demonstrated that the concept importance learned by our model is better aligned with clinical importance in Supp. Fig 6. Additionally, we investigated the sensitivity of λ on these datasets, as illustrated in Supp. Fig 4.

Skin Disease Classification: We also validated our approach to integrate clinical knowledge into CBM on datasets with skin images¹, as shown in Table 2 and Supp. Fig 5. Similar observations were found for the skin datasets, as integrating CBM with clinical knowledge improves the classification performance on out-of-domain datasets. For in-domain validation, while CBM with VGG16 backbones showed improved performance using our approach, CBM with ViT-B/16 and MLP classifiers experienced a slight decrease in performance (<0.3%). This could be due to the constraint of enforcing the model to learn in a clinical approach, limiting its ability to capture specific patterns within the in-domain dataset. Additionally, there was a drop in performance with CBM using MLP classifiers compared to linear ones, although a similar drop is observed in both our model and baselines. This may be due to overfitting in binary classification.

Table 2. Evaluation on datasets with skin images using F1 score (mean \pm 95% CI). The results indicate that training the model with clinical knowledge (*Ours*) enhances its classification performance, particularly on the out-of-domain dataset.

CBM Backbone+Classifier	Fitzpatrick 17k (In-domain)		DDI (Out-of-domain)	
	Baseline	Ours	Baseline	Ours
VGG16+Linear	76.58 \pm 0.80	77.11 \pm 0.23	58.87 \pm 0.39	60.13 \pm 0.85
VGG16+MLP(20)	74.28 \pm 1.74	75.64 \pm 0.91	52.33 \pm 6.01	54.78 \pm 7.10
VGG16+MLP(128)	73.78 \pm 1.49	75.04 \pm 0.99	51.60 \pm 7.28	56.01 \pm 1.06
ViT-B/16+Linear	76.88 \pm 1.03	76.97 \pm 0.86	57.78 \pm 2.54	59.46 \pm 1.58
ViT-B/16+MLP(20)	78.45 \pm 1.85	78.17 \pm 0.69	53.44 \pm 3.44	57.11 \pm 4.97
ViT-B/16+MLP(128)	77.89 \pm 1.25	77.70 \pm 1.03	55.88 \pm 3.52	59.04 \pm 3.72

Table 3. Comparison of WBC classification F1 score (mean \pm 95% CI) on out-of-domain testing, with integration (*Ours*), with random integration (*Random*), and without the integration of clinical knowledge (*Baseline*).

Scirep			RaabinWBC		
Baseline	Random	Ours	Baseline	Random	Ours
66.31 \pm 3.73	67.36 \pm 0.30	70.50 \pm 2.94	54.45 \pm 3.14	52.68 \pm 3.92	58.40 \pm 5.55

Is knowledge working as expected? A part of our method randomly removes some input concepts from the classifier, which might act as regularization like Dropout [17]. To ensure that the performance improvement comes from our knowledge integration, we investigate whether random interference, as opposed to systematic integration of knowledge, also improves the model’s performance from the baseline or not. The results are shown in Table 3. Remarkably, when clinical knowledge is integrated, the F1-Score for out-of-domain testing achieves a notable improvement compared to the model without the integration of knowledge and with random integration, highlighting the potential benefits of incorporating clinical insights into the classification process.

4 Conclusion

Our findings underscore the considerable potential of enhancing the performance of CBMs with clinical knowledge, especially when testing on out-of-domain datasets. The CBM guided by clinical knowledge exhibits better performance when dealing with diverse sources of images, such as WBC images sourced from varied staining conditions and microscopy techniques. Unlike purely data-driven CBMs that rely solely on features extracted from the training set, our clinically guided model incorporates insights accumulated over years of experience. This alignment with clinical knowledge enhances the model’s adaptability across varied image sources, thereby boosting its robustness and predictive accuracy.

Acknowledgments. This research was carried out at the Rapid-Rich Object Search (ROSE) Lab, Nanyang Technological University, Singapore. The research is supported in part by the NTU-PKU Joint Research Institute, a collaboration between Nanyang Technological University and Peking University, sponsored by a donation from the Ng Teng Fong Charitable Foundation. This work is also supported in part by Sysmex Corporation, Japan. We thank Mari Kono and Joao Nunes for their helpful discussions. The computational work for this article was partially performed on the resources of the National Supercomputing Centre (NSCC), Singapore (NSCC project ID: 12003885).

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., Alférez, S., Molina, Á., Boldú, L., Rodellar, J.: A dataset of microscopic peripheral blood cell images for development of automatic recognition systems. *Data in Brief* **30** (2020)
2. Al-Dulaimi, K.A.K., Banks, J., Chandran, V., Tomeo-Reyes, I., Nguyen Thanh, K.: Classification of white blood cell types from microscope images: Techniques and challenges. *Microscopy Science: Last Approaches on Educational Programs and Applied Research* (2018)
3. Amgad, M., Hodge, J.M., Elsebaie, M.A., Bodelon, C., Puvanesarajah, S., Gutman, D.A., Siziopikou, K.P., Goldstein, J.A., Gaudet, M.M., Teras, L.R., et al.: A population-level digital histologic biomarker for enhanced prognosis of invasive breast cancer. *Nature Medicine* **30**(1) (2024)
4. Ardila, D., Kiraly, A.P., Bharadwaj, S., Choi, B., Reicher, J.J., Peng, L., Tse, D., Etemadi, M., Ye, W., Corrado, G., et al.: End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nature Medicine* **25**(6) (2019)
5. Daneshjou, R., Vodrahalli, K., Novoa, R.A., Jenkins, M., Liang, W., Rotemberg, V., Ko, J., Swetter, S.M., Bailey, E.E., Gevaert, O., et al.: Disparities in dermatology ai performance on a diverse, curated clinical image set. *Science Advances* **8**(31) (2022)
6. Daneshjou, R., Yuksekogonul, M., Cai, Z.R., Novoa, R., Zou, J.Y.: Skincon: A skin disease dataset densely annotated by domain experts for fine-grained debugging and analysis. *Advances in Neural Information Processing Systems* (2022)
7. Espinosa Zarlenga, M., Barbiero, P., Ciravegna, G., Marra, G., Giannini, F., Dili-genti, M., Shams, Z., Precioso, F., Melacci, S., Weller, A., Lio, P., Jamnik, M.: Concept embedding models: Beyond the accuracy-explainability trade-off. *Advances in Neural Information Processing Systems* **35** (2022)
8. 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: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021)
9. Han, Z., Wei, B., Zheng, Y., Yin, Y., Li, K., Li, S.: Breast cancer multi-classification from histopathological images with structured deep learning model. *Scientific Reports* **7**(1) (2017)
10. Koh, P.W., Nguyen, T., Tang, Y.S., Mussmann, S., Pierson, E., Kim, B., Liang, P.: Concept bottleneck models. In: *International Conference on Machine Learning* (2020)

11. Kouzehkhanan, Z.M., Saghari, S., Tavakoli, S., Rostami, P., Abaszadeh, M., Mirzadeh, F., Satlsar, E.S., Gheidishahran, M., Gorgi, F., Mohammadi, S., et al.: A large dataset of white blood cells containing cell locations and types, along with segmented nuclei and cytoplasm. *Scientific Reports* **12**(1) (2022)
12. Li, M., Lin, C., Ge, P., Li, L., Song, S., Zhang, H., Lu, L., Liu, X., Zheng, F., Zhang, S., et al.: A deep learning model for detection of leukocytes under various interference factors. *Scientific Reports* **13**(1) (2023)
13. Manh, V.T., Zhou, J., Jia, X., Lin, Z., Xu, W., Mei, Z., Dong, Y., Yang, X., Huang, R., Ni, D.: Multi-attribute attention network for interpretable diagnosis of thyroid nodules in ultrasound images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **69**(9) (2022)
14. Marcinkevičs, R., Wolfertstetter, P.R., Klimiene, U., Chin-Cheong, K., Paschke, A., Zerres, J., Denzinger, M., Niederberger, D., Wellmann, S., Ozkan, E., et al.: Interpretable and intervenable ultrasonography-based machine learning models for pediatric appendicitis. *Medical Image Analysis* **91** (2024)
15. Oikarinen, T., Das, S., Nguyen, L.M., Weng, T.W.: Label-free concept bottleneck models. In: *International Conference on Learning Representations* (2023)
16. Ran, A.R., Tham, C.C., Chan, P.P., Cheng, C.Y., Tham, Y.C., Rim, T.H., Cheung, C.Y.: Deep learning in glaucoma with optical coherence tomography: A review. *Eye* **35**(1) (2021)
17. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* **15**(1) (2014)
18. Standring, S., Ellis, H., Healy, J., Johnson, D., Williams, A., Collins, P., Wigley, C.: *Gray's anatomy: The anatomical basis of clinical practice*. *American Journal of Neuroradiology* **26**(10) (2005)
19. Tsutsui, S., Pang, W., Wen, B.: Wbcatt: A white blood cell dataset annotated with detailed morphological attributes. *Advances in Neural Information Processing Systems* (2023)
20. Tsutsui, S., Su, Z., Wen, B.: Benchmarking white blood cell classification under domain shift. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE (2023)
21. Yan, S., Yu, Z., Zhang, X., Mahapatra, D., Chandra, S.S., Janda, M., Soyer, P., Ge, Z.: Towards trustable skin cancer diagnosis via rewriting model's decision. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023)
22. Yin, C., Liu, S., Shao, R., Yuen, P.C.: Focusing on clinically interpretable features: Selective attention regularization for liver biopsy image classification. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer (2021)
23. Yuksekgonul, M., Wang, M., Zou, J.: Post-hoc concept bottleneck models. In: *International Conference on Learning Representations* (2023)
24. Zhou, Y.J., Liu, W., Gao, Y., Xu, J., Lu, L., Duan, Y., Cheng, H., Jin, N., Man, X., Zhao, S., et al.: A novel multi-task model imitating dermatologists for accurate differential diagnosis of skin diseases in clinical images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer (2023)