

Training-free Test-time Improvement for Explainable Medical Image Classification

Hangzhou He^{1,2,3,5}[0009-0009-3050-8773], Jiachen Tang^{4,5}[0009-0008-2295-6042],
Lei Zhu^{1,2,3}[0000-0003-0506-4268], Kaiwen Li^{1,2,3}[0000-0002-4988-2687], and
Yanye Lu^{1,2,3,6}[0000-0002-3063-8051]

¹ Department of Biomedical Engineering, Peking University, Beijing, China

² Institute of Medical Technology, Peking University, Beijing, China

³ National Biomedical Imaging Center, Peking University, Beijing, China

⁴ School of Computer Science and Technology, Beijing Institute of Technology,
Beijing, China

⁵ Equal Contribution

⁶ Corresponding Author

zhuang@stu.pku.edu.cn, yanye.lu@pku.edu.cn

Abstract. Deep learning-based medical image classification techniques are rapidly advancing in medical image analysis, making it crucial to develop accurate and trustworthy models that can be efficiently deployed across diverse clinical scenarios. Concept Bottleneck Models (CBMs), which first predict a set of explainable concepts from images and then perform classification based on these concepts, are increasingly being adopted for explainable medical image classification. However, the inherent explainability of CBMs introduces new challenges when deploying trained models to new environments. Variations in imaging protocols and staining methods may induce concept-level shifts, such as alterations in color distribution and scale. Furthermore, since CBM training requires explicit concept annotations, fine-tuning models solely with image-level labels could compromise concept prediction accuracy and faithfulness - a critical limitation given the high cost of acquiring expert-annotated concept labels in medical domains. To address these challenges, we propose a training-free confusion concept identification strategy. By leveraging minimal new data (e.g., 4 images per class) with only image-level labels, our approach enhances out-of-domain performance without sacrificing source domain accuracy through two key operations: masking misactivated confounding concepts and amplifying under-activated discriminative concepts. The efficacy of our method is validated on both skin and white blood cell images. Our code is available at: <https://github.com/riverback/TF-TTI-XMed>.

Keywords: Test-time improvement · Inherent explainability · Concept bottleneck model.

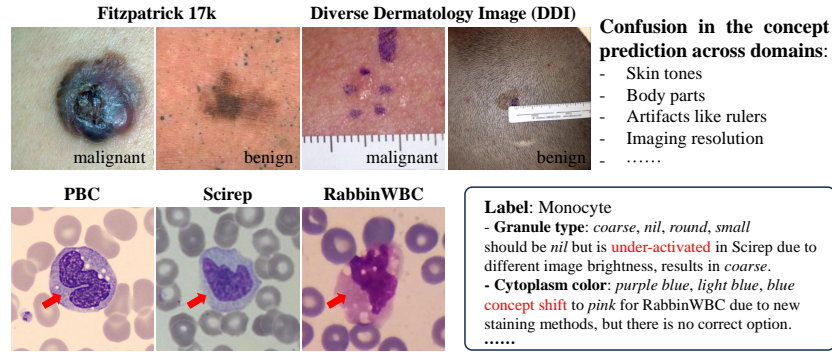


Fig. 1. Challenges when deploying CBMs to new environments: the concept features across different domains may vary intensively and this may result in under-activated discriminative concepts like the *nil granule type* for monocytes, or shifted concepts like the *pink cytoplasm color* in RabbinWBC (which are mostly blue tones in the PBC dataset). However, annotating fine-grained concept labels for the new data is costly due to the expertise, and directly fine-tuning with image-level labels may lead to overfitting and degrade source domain performance.

1 Introduction

The development and integration of deep learning technologies have greatly transformed the field of medical image analysis and achieved remarkable progress in multiple medical fields [2,23,32,35]. However, the particularity of medical applications poses additional challenges to the explainability of deep learning models, such as checking whether the models are reliable [9], fair [4], and robust [3]. However, for a long time, there has been a trade-off between model explainability and accuracy in the field of deep learning, and the accuracy of explainable models is often significantly lower than that of their black-box counterparts [11], which limits their application in real-world medical scenarios where accuracy is prioritized [26]. As an emerging self-explainable model architecture, the concept bottleneck model (CBM [17]) uses a set of interpretable concepts as the intermediate representation of images and then relies on an interpretable linear classification layer to achieve interpretable image classification. When an accurate and sufficient concept set is provided, CBMs can often maintain good explainability and achieve performance comparable to that of black-box models and thus have been widely applied to medical image classification tasks [6,15,25,33].

However, deploying CBMs to a new environment while maintaining explainability is challenging, e.g., a new medical institution or imaging device. Firstly, the concept features across domains may vary intensively (as shown in Fig. 1), which poses challenges for both fine-tuning in the target domain and accurate concept prediction, and the clinical knowledge required for annotating fine-grained medical concepts further exacerbates this drawback. Although some label-free methods have been proposed to construct CBMs by leveraging the capabilities of large language models or vision-language models [12,24,34], it is

difficult for us to assess whether the generated medical concepts are accurate and faithful [14]. Secondly, even if there is concept-level labeled data in the target domain, unless sufficient, direct fine-tuning of the model may lead to severe over-fitting and also degrade the model’s performance on the source domain [22]. Moreover, fine-tuning with only image labels may also undermine the faithfulness of CBMs, as the model may forget previously learned knowledge and fail to leverage the desired concepts. [21].

To address the above issues, we propose a training-free **confusion concept identification strategy** to improve the test-time performance of CBMs. By identifying and adjusting confusing concepts during the test phase, we can **mask misactivated confounding concepts** or **amplify under-activated discriminative concepts**. More importantly, our method does not harm the faithfulness of the concept prediction procedure and in-domain performance. Experimental results on skin and white blood cell images demonstrate that, with minimal data with only image labels (e.g., only 4 images per class), our method achieves performance comparable to fine-tuning, without sacrificing the source domain’s accuracy while keeping the explainability of concept bottleneck models.

2 Related Work

2.1 Concept Bottleneck Models for Medical Image Classification

The concept bottleneck model (CBM [17]) maps visual images to a space composed of predefined concepts, extracts intermediate concepts as the representation of the images, and uses an interpretable linear classifier for classification. Due to its inherent explainability, CBMs have been extensively utilized in medical image classification tasks, including white blood cell classification [25], skin disease classification [15], fundus disease classification [6], and the classification of thorax diseases [33]. Among these works, Pang *et al.* proposed to align CBMs with clinical knowledge during training to improve out-of-domain (OOD) performance, but needed additional clinical knowledge labels to indicate concept preference [25]. Yang *et al.* proposed Knowledge-enhanced Bottlenecks that used language models and PubMed as the knowledge resources for making CBMs less sensitive to domain shift [33]. Nevertheless, existing methods for improving the OOD performance of CBMs all require additional training or the integration of a vision-language model like CLIP to help concept prediction at test time [5], which can be challenging to quickly deploy in new environments.

Beyond these methods, CBM itself is also renowned for its test-time concept intervention capability, which permits users to modify activation values of specific concepts at the concept bottleneck [28]. However, when applied to OOD medical images, such interventions often rely on specialized medical knowledge and fail to consistently improve model performance across diverse datasets, thereby limiting their generalizability and practical utility. In response to this challenge, the Concept Bottleneck Memory Model (CB2M) [30] introduces a two-fold memory module designed to capture model errors during training and

validation phases, utilizing these memories to bolster test performance. Nonetheless, CB2M fails to adequately explore the issue of concept shift in OOD scenarios, which remains a critical concern. Another prominent approach for adapting pre-trained models to distribution shifts is test-time adaptation (TTA) [20]. TTA methods typically exploit unlabeled test-domain data or data streams to facilitate domain adaptation. However, the two-step inference process of CBM necessitates the preservation of the model’s concept prediction fidelity during the TTA phase. Regrettably, research on enhancing CBM’s test-time performance while maintaining fidelity remains unexplored.

3 Method

3.1 Preliminary: Concept Bottleneck Model

Formally, consider an image dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{c}_i, y_i)\}$ where \mathbf{x} denotes an image, $\mathbf{c} = \{c^1, c^2, \dots, c^L\}$ is L predefined concepts, and $y \in \{0, 1\}^{K \times 1}$ is the image label from K target classes. As shown in Fig. 2 (a), a CBM first predicts concepts using a concept predictor g and then predicts targets based on the concepts using a classifier f :

$$\hat{\mathbf{c}} = g(\mathbf{x}), \hat{y} = f(\hat{\mathbf{c}}) \quad (1)$$

In this work, we jointly train g and f to make the framework compatible with additional optimization strategies such as incorporating clinical knowledge [28].

$$\hat{f}, \hat{g} = \arg \min_{f, g} \sum_i^{|\mathcal{D}|} [\varphi \mathcal{L}_c(g(\mathbf{x}_i); \mathbf{c}_i) + \mathcal{L}_y(f(g(\mathbf{x}_i)); y_i)] \quad (2)$$

where \mathcal{L}_c and \mathcal{L}_y are loss functions for concept prediction and class prediction respectively, for some $\varphi > 0$.

3.2 Confusion Concept Identification Strategy

We first give some observations into the concept prediction procedure on OOD data. As shown in Fig. 2 (b), there are typically two types of confusion concepts: **1) under-activated discriminative concepts** that are crucial for predicting the target class but are less activated than other concepts because of feature distribution shifts caused by different image brightness, resolution and so on; and **2) misactivated confounding concepts** that are unable to be faithfully activated in the OOD dataset due to huge changes in feature representation, like the blue tones cytoplasm color is changed to pink in the RabbinWBC dataset. Based on these observations, we propose our **confusion concept identification strategy** which leverages the concept activation memory during the in-domain (ID) model development and then uses minimal image-level labeled data from the target domain to identify the two types of confusion concepts. As presented in Fig. 2 (c), we save the mean ($\bar{\mathbf{c}}_{ID}$) and the standard deviation

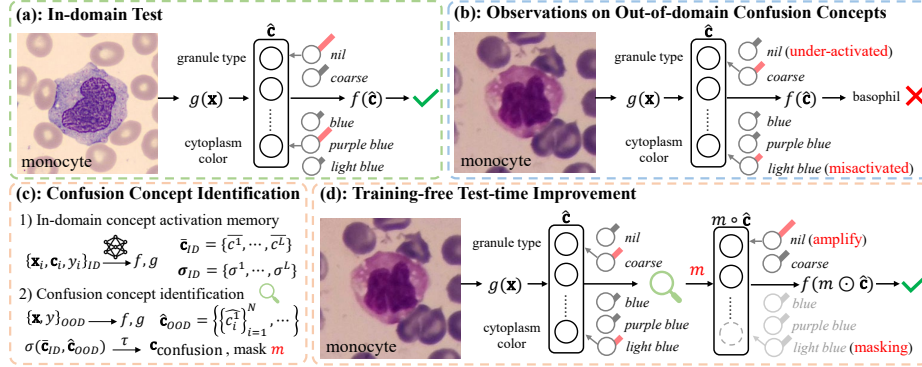


Fig. 2. Method overview: (a) a trained CBM can perform explainable classification with accurate concepts, (b) but fails on the out-of-domain (OOD) data because of two types of confusion concepts, under-activated discriminative concepts and misactivated confusing concepts. So we (c) propose a confusion concept identification strategy and (d) use the concept mask m to amplify or mask these concepts.

(σ_{ID}) for concept activation value on the validation set \mathcal{D}_{val} as the concept activation memory for each target class k :

$$\{\bar{c}_{ID}, \sigma_{ID}\}_k = \{(\bar{c}^1, \sigma^1), \dots, (\bar{c}^L, \sigma^L)\}_k, \quad k \in \{1, 2, \dots, K\} \quad (3)$$

where the subscript k denotes the k -th element of a set. For example, the concept activation memory of class k for c^1 is calculated as:

$$\bar{c}^1 = \frac{1}{|\mathcal{D}_{val}(k)|} \sum_{i, y_i=k} g(\mathbf{x}_i)[0], \quad \sigma^1 = \sqrt{\sum_{i, y_i=k} \frac{1}{|\mathcal{D}_{val}(k)|} (g(\mathbf{x}_i)[0] - \bar{c}^1)^2} \quad (4)$$

Given a set of image-level labeled OOD images $\mathcal{D}_{OOD} = \{\mathbf{x}_i, y_i\}_{i=1}^{N \times K}$ where N is the number of images for each class and is usually small, we iterate these data once to get the concept activation information in the target domain. The OOD concept activation information for class k is:

$$\hat{\mathbf{c}}_{OOD,k} = \left\{ \{\hat{c}_i^1\}_{i=1}^N, \dots, \{\hat{c}_i^L\}_{i=1}^N \right\}_k \quad (5)$$

then we can identify confusing concepts for class k by calculating the stability of concept activation (e.g., standard deviation) between ID and OOD samples:

$$\mathbf{c}_{confusion,k} = \bigcup_{i=1, \dots, L} \left\{ c^i \mid \sigma[\text{concat}(\bar{\mathbf{c}}_{ID,k}[i], \hat{\mathbf{c}}_{OOD,k}[i])] \geq \tau \cdot \sigma_{ID,k}[i] \right\} \quad (6)$$

where σ denotes the computation of the standard deviation, the factor $\tau > 1$ ensures the OOD concept activation variation is larger than that of ID. We use $\tau = 1.4$ for skin and $\tau = 1.1$ for WBC. A concept is deemed misactivated if

it confuses more than half of the classes. The remaining confusing concepts are treated as under-activated discriminative concepts if their in-domain activation probability for the corresponding class exceeds 0.5. Our intuition is that concepts with low variance in the source domain are likely class-discriminative; thus, large fluctuations in the test domain indicate unreliability. Conversely, concepts with high source-domain variance are likely non-discriminative and are naturally downweighted by Eq.(6), which compares against in-domain variance. If such noisy concepts are still activated by Eq.(6), they tend to confuse all classes and should therefore be masked.

3.3 Confusion Concepts Activation Manipulation

Upon acquiring the two categories of confusion concepts, we can adjust the activation levels of these concepts to enhance test-time performance. Towards this end, we put forward two straightforward yet efficacious operations: **(1) Masking Misactivated Confounding Concepts**, and **(2) Amplifying Under-Activated Discriminative Concepts**. As depicted in Fig. 2 (d), we initialize a concept mask m with a value of 1. We mask the misactivated confounding concepts by setting the corresponding values to zero. Meanwhile, to amplify the under-activated discriminative concepts, we multiply them by a factor greater than 1; for the sake of simplicity, we adopt 2 in this paper. During inference, we simply multiply the concept mask m element-wise with the predicted concepts $\hat{\mathbf{c}}$, and use the result as input to the final classifier f :

$$\hat{y}_{\text{new}} = f(m \odot \hat{\mathbf{c}}) = f(m \odot g(\mathbf{x})) \quad (7)$$

4 Experiment

Settings To evaluate the effectiveness of our method, we conduct experiments on two medical image classification tasks: dermatology and white blood cell (WBC) classification. For skin dataset, we train a dermatology classifier ($K = 2$) using concept annotations from the Fitzpatrick 17k dataset [10] and the SkinCon dataset [8]. We use the same concepts ($L = 22$) with [25] to train the concept predictor. The Fitzpatrick 17k dataset ($|\mathcal{D}| = 3,479$) is used as the in-domain dataset and the Diverse Dermatology Images (DDI) dataset [7] ($|\mathcal{D}| = 656$) serves as the out-of-domain dataset. The DDI dataset focuses on diverse skin tones and uncommon diseases, which introduces variability in skin tones, lighting conditions, and disease presentation. For the WBC dataset, we use the PBC dataset [1] combined with concept annotations ($L = 11$) from the WBCAtt dataset [31]. We train the models ($K = 5$) following the data split from [31]. The Scirep dataset [19] ($|\mathcal{D}| = 2,019$) and the RaabinWBC dataset [18] ($|\mathcal{D}| = 4,339$) are used as out-of-domain datasets. These WBC datasets include images captured under different staining conditions and imaging devices, resulting in different cell appearances and textures. Example Images are shown in Fig. 1.

We separately use VGG16 [29] and ResNet34 [13] pre-trained on ImageNet as g and use a linear classifier as f to build CBMs. For both tasks, we use an

Table 1. Classification accuracy and F1 score on skin images. N is the number of image-level labeled images for each class in the out-of-domain dataset.

CBM Backbone	Fitzpatrick 17k (In-domain)		DDI (Out-of-domain)	
	Accuracy	F1 Score	Accuracy	F1 Score
VGG16	75.28	75.28	65.72	58.72
+ \mathcal{F} (N=8)	61.40 \downarrow 13.88	57.09 \downarrow 17.38	69.65 \uparrow 3.93	56.58 \downarrow 2.14
+Ours (N=4)	74.53 \downarrow 0.75	74.36 \downarrow 0.92	70.28 \uparrow 4.56	60.46 \uparrow 1.74
+Ours (N=8)	74.42 \downarrow 0.86	74.25 \downarrow 1.03	70.44 \uparrow 4.72	60.75 \uparrow 2.03
VGG16+ \mathcal{K}	76.95	76.93	62.11	56.00
+ \mathcal{F} (N=8)	64.79 \downarrow 12.16	63.43 \downarrow 13.50	62.11	59.44 \uparrow 3.44
+Ours (N=4)	76.95	76.94 \uparrow 0.01	66.19 \uparrow 4.09	57.55 \uparrow 1.55
+Ours (N=8)	76.95	76.95 \uparrow 0.02	65.88 \uparrow 3.77	57.16 \uparrow 1.16
ResNet34	77.72	77.72	62.89	59.09
+ \mathcal{F} (N=8)	68.61 \downarrow 8.78	68.60 \downarrow 9.12	64.94 \uparrow 2.05	60.22 \uparrow 1.13
+Ours(N=4)	77.98 \uparrow 0.26	77.97 \uparrow 0.25	64.15 \uparrow 1.26	60.13 \uparrow 1.04
+Ours(N=8)	77.32 \downarrow 0.40	77.22 \downarrow 0.50	70.13 \uparrow 7.23	62.28 \uparrow 3.20
ResNet34+ \mathcal{K}	77.38	77.29	62.26	58.48
+ \mathcal{F} (N=8)	67.61 \downarrow 9.77	67.61 \downarrow 9.68	67.14 \uparrow 4.88	59.17 \uparrow 0.69
+Ours (N=4)	78.41 \uparrow 1.12	78.41 \uparrow 1.03	66.98 \uparrow 4.72	60.48 \uparrow 2.00
+Ours (N=8)	77.84 \uparrow 0.46	77.73 \uparrow 0.44	68.55 \uparrow 6.29	59.86 \uparrow 1.38


AdamW optimizer [16] with weight decay set to 0.01. The batch size is set to 64. The learning rate is set to $1\text{e-}4$. We use cross-entropy loss with label smoothing ($\phi=0.05$) and φ in Eq. 2 is set to 1. We follow [25] to train 30 epochs per run and also implement CBMs integrated with clinical knowledge (denoted as + \mathcal{K}). Except for the additional loss functions used in [25], other parameters are kept the same for both vanilla CBMs and the clinical knowledge integrated CBMs. For fine-tuning experiments (denoted as + \mathcal{F}), we use the same parameters and settings with in-domain settings. The same images (N for each class) are used for fine-tuning and our confusion concept identification strategy.

Results The main results for skin images are presented in Table 1 and those for WBC images are in Table 2. Without any fine-tuning or fine-grained concept labels, our method can enhance the classification performance of the CBM model on out-of-domain datasets. For the vanilla CBM with VGG16 as the backbone, with only 8 images from target domain ($N=4$ per class), our method achieves an accuracy improvement of 4.56 and an F1 score improvement of 1.74 on the DDI dataset (bold in Table 1), even surpassing out-of-domain performance when using twice the number of images for fine-tuning. More importantly, under nearly all experimental settings, our method can maintain or even improve the model’s in-domain performance. In contrast, fine-tuning may lead to a drastic decline in in-domain performance. For example, the accuracy drops from 99.77 to 62.96 on WBC classification, and the F1 score drops from 99.67 to 60.95 (bold in Table 2).

Table 2. Classification accuracy and F1 score on WBC images. N is the number of image-level labeled images for each class in the out-of-domain dataset.

CBM Backbone	In-domain Dataset		Out-of-domain Datasets			
	PBC ¹		Scirep		RaabinWBC	
	Accuracy	F1 Score	Accuracy	F1 Score	Accuracy	F1 Score
VGG16	99.71	99.61	84.10	70.85	75.06	50.14
+ \mathcal{F} (N=4)	97.42/96.16	96.96/97.09	88.01	75.55	86.89	78.34
+Ours (N=4)	99.71/99.52	99.61/99.41	84.25	71.45	76.05	51.64
+Ours (N=8)	99.32/99.71	99.45/99.74	86.43	72.18	78.43	55.09
VGG16+ \mathcal{K}	99.84	99.77	84.00	67.37	77.21	51.89
+ \mathcal{F} (N=4)	91.87/96.45	90.64/96.10	84.50	77.20	88.20	82.62
+Ours (N=4)	99.74/99.84	99.66/99.76	84.84	68.59	77.69	53.42
+Ours (N=8)	99.71/99.61	99.61/99.62	86.53	70.38	78.77	55.15
ResNet34	99.77	99.67	80.29	69.34	37.35	27.12
+ \mathcal{F} (N=4)	92.16/ 62.96	89.90/ 60.95	90.59	85.85	75.71	72.58
+Ours (N=4)	99.74/99.81	99.65/99.74	81.72	71.74	54.99	35.94
+Ours (N=8)	99.81/99.81	99.74/99.75	81.72	71.74	50.86	36.50
ResNet34+ \mathcal{K}	99.74	99.65	83.01	69.35	16.73	10.17
+ \mathcal{F} (N=4)	82.74/71.51	77.27/69.92	83.95	72.60	83.43	75.83
+Ours (N=4)	99.84/99.77	99.79/99.69	84.25	71.14	39.09	22.93
+Ours (N=8)	99.68/99.77	99.56/99.69	84.35	72.40	59.55	24.20

1: test separately for using Scirep or RabbinWBC as the out-of-domain dataset.

**Fig. 3.** Visualization of the inference procedure and focus areas of the model before and after applying our test-time improvement strategy.

Are Confusion Concepts Manipulated Correctly? We visualize the inference process and utilize Grad-CAM [27] to highlight the important features that the model relies on for predictions to see whether the concept manipulation is faithful. As presented in Fig. 3, the model becomes more focused on the target object after our confusion concept manipulation, indicating better faithfulness during inference, in addition to the improvement of classification accuracy.

5 Conclusion

In this study, we propose a novel training-free confusion concept identification strategy to enhance the adaptability of concept bottleneck models (CBMs) for out-of-domain data. Our approach requires only minimal image-level annotations, yet effectively improves test-time performance while maintaining in-

domain accuracy and model explainability. Both quantitative evaluations and visual analyses demonstrate the efficacy of our method, revealing the potential of explainability-guided test-time adaptation methods. Future research may focus on further reducing the labeling requirement and developing label-free methods for adopting explainable medical classification models.

Acknowledgments. This work was supported by the Natural Science Foundation of China (623B2001, 82371112), the Science Foundation of Peking University Cancer Hospital (JC202505), Natural Science Foundation of Beijing Municipality (Z210008) and the Clinical Medicine Plus X - Young Scholars Project of Peking University, the Fundamental Research Funds for the Central Universities (PKU2025PKULCXQ008).

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., 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. Acosta, J.N., Falcone, G.J., Rajpurkar, P., Topol, E.J.: Multimodal biomedical ai. *Nature Medicine* **28**(9), 1773–1784 (2022)
3. Balendran, A., Beji, C., Bouvier, F., Khalifa, O., Evgeniou, T., Ravaud, P., Porcher, R.: A scoping review of robustness concepts for machine learning in healthcare. *npj Digital Medicine* **8**(1), 38 (2025)
4. Chen, R.J., Wang, J.J., Williamson, D.F., Chen, T.Y., Lipkova, J., Lu, M.Y., Sahai, S., Mahmood, F.: Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nature biomedical engineering* **7**(6), 719–742 (2023)
5. Choi, J., Raghuram, J., Li, Y., Jha, S.: Adaptive concept bottleneck for foundation models under distribution shifts. In: *The Thirteenth International Conference on Learning Representations* (2025)
6. Chowdhury, T.F., Phan, V.M.H., Liao, K., To, M., Xie, Y., van den Hengel, A., Verjans, J.W., Liao, Z.: Adacbm: An adaptive concept bottleneck model for explainable and accurate diagnosis. In: *MICCAI* (10). *Lecture Notes in Computer Science*, vol. 15010, pp. 35–45. Springer (2024)
7. 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), eabq6147 (2022)
8. Daneshjou, R., Y  ksek  n  l, M., Cai, Z.R., Novoa, R.A., Zou, J.Y.: Skincon: A skin disease dataset densely annotated by domain experts for fine-grained debugging and analysis. In: *NeurIPS* (2022)
9. Fajtl, J., Welikala, R.A., Barman, S., Chambers, R., Bolter, L., Anderson, J., Olvera-Barrios, A., Shakespeare, R., Egan, C., Owen, C.G., et al.: Trustworthy evaluation of clinical ai for analysis of medical images in diverse populations. *NEJM AI* **1**(9), AIoa2400353 (2024)
10. 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: *CVPR Workshops*. pp. 1820–1828 (2021)

11. Gunning, D., Aha, D.: Darpa’s explainable artificial intelligence (xai) program. *AI magazine* **40**(2), 44–58 (2019)
12. He, H., Zhu, L., Zhang, X., Zeng, S., Chen, Q., Lu, Y.: V2c-cbm: Building concept bottlenecks with vision-to-concept tokenizer. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 39, pp. 3401–3409 (2025)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR*. pp. 770–778. IEEE Computer Society (2016)
14. Huang, Q., Song, J., Hu, J., Zhang, H., Wang, Y., Song, M.: On the concept trustworthiness in concept bottleneck models. In: *AAAI*. pp. 21161–21168. AAAI Press (2024)
15. Kim, I., Kim, J., Choi, J., Kim, H.J.: Concept bottleneck with visual concept filtering for explainable medical image classification. In: *ISIC/Care-AI/MedAGI/DeCaF@MICCAI*. *Lecture Notes in Computer Science*, vol. 14393, pp. 225–233. Springer (2023)
16. Kingma, D.P.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
17. Koh, P.W., Nguyen, T., Tang, Y.S., Mussmann, S., Pierson, E., Kim, B., Liang, P.: Concept bottleneck models. In: *ICML. Proceedings of Machine Learning Research*, vol. 119, pp. 5338–5348. PMLR (2020)
18. 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), 1123 (2022)
19. 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), 2160 (2023)
20. Liang, J., He, R., Tan, T.: A comprehensive survey on test-time adaptation under distribution shifts. *Int. J. Comput. Vis.* **133**(1), 31–64 (2025)
21. Mai, Z., Chowdhury, A., Zhang, P., Tu, C., Chen, H., Pahuja, V., Berger-Wolf, T.Y., Gao, S., Stewart, C.V., Su, Y., Chao, W.: Fine-tuning is fine, if calibrated. In: *NeurIPS* (2024)
22. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: *Psychology of learning and motivation*, vol. 24, pp. 109–165. Elsevier (1989)
23. Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J., Topol, E.J., Rajpurkar, P.: Foundation models for generalist medical artificial intelligence. *Nature* **616**(7956), 259–265 (2023)
24. Oikarinen, T.P., Das, S., Nguyen, L.M., Weng, T.: Label-free concept bottleneck models. In: *ICLR. OpenReview.net* (2023)
25. Pang, W., Ke, X., Tsutsui, S., Wen, B.: Integrating clinical knowledge into concept bottleneck models. In: *MICCAI (4). Lecture Notes in Computer Science*, vol. 15004, pp. 243–253. Springer (2024)
26. Sabuncu, M.R., Wang, A.Q., Nguyen, M.: Ethical use of artificial intelligence in medical diagnostics demands a focus on accuracy, not fairness. *NEJM AI* **2**(1), AIp2400672 (2025)
27. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. *Int. J. Comput. Vis.* **128**(2), 336–359 (2020)
28. Shin, S., Jo, Y., Ahn, S., Lee, N.: A closer look at the intervention procedure of concept bottleneck models. In: *ICML. Proceedings of Machine Learning Research*, vol. 202, pp. 31504–31520. PMLR (2023)

29. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
30. Steinmann, D., Stammer, W., Friedrich, F., Kersting, K.: Learning to intervene on concept bottlenecks. In: ICML. OpenReview.net (2024)
31. Tsutsui, S., Su, Z., Wen, B.: Benchmarking white blood cell classification under domain shift. In: ICASSP. pp. 1–5. IEEE (2023)
32. Varghese, C., Harrison, E.M., O’Grady, G., Topol, E.J.: Artificial intelligence in surgery. *Nature medicine* **30**(1), 1257–1268 (2024)
33. Yang, Y., Gandhi, M., Wang, Y., Wu, Y., Yao, M.S., Callison-Burch, C., Gee, J.C., Yatskar, M.: A textbook remedy for domain shifts: Knowledge priors for medical image analysis. In: NeurIPS (2024)
34. 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: CVPR. pp. 19187–19197. IEEE (2023)
35. Zhang, K., Yang, X., Wang, Y., Yu, Y., Huang, N., Li, G., Li, X., Wu, J.C., Yang, S.: Artificial intelligence in drug development. *Nature medicine* **31**(1), 45–59 (2025)