

MICA: Towards Explainable Skin Lesion Diagnosis via Multi-Level Image-Concept Alignment

Yequan Bie¹, Luyang Luo¹, Hao Chen^{1,2,3*}

¹Department of Computer Science and Engineering, Hong Kong University of Science and Technology

²Department of Chemical and Biological Engineering, Hong Kong University of Science and Technology

³HKUST Shenzhen-Hong Kong Collaborative Innovation Research Institute

ybie@connect.ust.hk, cseluyang@ust.hk, jhc@cse.ust.hk

Abstract

Black-box deep learning approaches have showcased significant potential in the realm of medical image analysis. However, the stringent trustworthiness requirements intrinsic to the medical field have catalyzed research into the utilization of Explainable Artificial Intelligence (XAI), with a particular focus on concept-based methods. Existing concept-based methods predominantly apply concept annotations from a single perspective (e.g., global level), neglecting the nuanced semantic relationships between sub-regions and concepts embedded within medical images. This leads to underutilization of the valuable medical information and may cause models to fall short in harmoniously balancing interpretability and performance when employing inherently interpretable architectures such as Concept Bottlenecks. To mitigate these shortcomings, we propose a multi-modal explainable disease diagnosis framework that meticulously aligns medical images and clinical-related concepts semantically at multiple strata, encompassing the image level, token level, and concept level. Moreover, our method allows for model intervention and offers both textual and visual explanations in terms of human-interpretable concepts. Experimental results on three skin image datasets demonstrate that our method, while preserving model interpretability, attains high performance and label efficiency for concept detection and disease diagnosis. The code is available at <https://github.com/Tommy-Bie/MICA>.

1 Introduction

Black-box deep learning methods have surfaced as powerful instruments in medical image analysis, offering significant potential to revolutionize healthcare diagnostics and treatments (Keremany et al. 2018; Esteva et al. 2017). These methods excel at handling the extensive and intricate data inherent to the medical field, rendering them suitable for many tasks (Litjens et al. 2017). Despite the encouraging performance, their end-to-end prediction nature leads to a lack of transparency, raising critical issues of trust and interpretability in high-stakes domains like healthcare (Rudin 2019).

The healthcare field, with its rigorous demands for trustworthiness, requires models that not only perform well but are also understandable and trustable by practitioners, which necessitates research into Explainable Artificial Intelligence

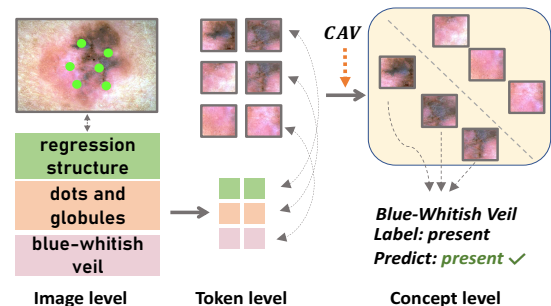


Figure 1: Our method learns image and concept semantic correspondences at the image, token, and concept levels.

(XAI). Within XAI, several approaches have been proposed to explain the neural networks using saliency map (Zhou et al. 2016; Selvaraju et al. 2017) highlighting the contribution of each pixel or region in the model’s prediction, while others leverage inspection of the learned features (Abbasi-Asl and Yu 2017), feature interactions (Tsang, Cheng, and Liu 2017), and influence functions (Koh and Liang 2017) to explain the models. However, the reliability of these post-hoc analyses, which offer explanations for a trained AI model, has been under considerable scrutiny recently. Some studies (Laugel et al. 2019; Rudin 2019) have shown that post-hoc explanation techniques often yield inconsistent results across different runs and are sensitive to slight changes in the input, making the post-hoc methods misleading as they could provide explanations that do not accurately reflect the model’s decision-making process.

Thus, ante-hoc explainable methods have garnered researchers’ interest, with a particular emphasis placed on concept-based methods. These methods aim to integrate interpretability into machine learning models by linking their predictions to human-understandable concepts (Koh et al. 2020; Yuksekgonul et al. 2023; Fang et al. 2020; Yan et al. 2023). For example, the Concept Bottleneck Model (CBM) (Koh et al. 2020) first predicts an intermediate set of pre-defined concepts and then uses these concepts to predict the final output. Yan et al. (2023) introduced a human-in-the-loop framework to eliminate confounding factors and improve model performance. These inherently interpretable methods offer concept-based explanations, which are gener-

*Corresponding author.

ally more understandable than post-hoc approaches.

However, concept-based methods are not devoid of limitations. A major challenge these methods face is that the model performance (e.g., classification accuracy) is sacrificed when designing an explainable architecture like Concept Bottlenecks, compared to the black-box methods. We argue this is caused by inefficient usage of valuable medical information, i.e., clinical-related concepts. Most existing methods apply concept annotations at a single level, e.g., Sarkar et al. (2022) only utilize concept labels to supervise the concept prediction results of the whole image, neglecting the intricate semantic connections between images' sub-regions and concepts. This narrow focus can limit the models' performance and interpretability, leading to unreliable concept detection and inaccurate diagnosis.

To address the mentioned challenges, we introduce a multi-modal explainable disease diagnosis framework to meticulously align medical images and clinical-related concepts semantically at multiple levels, encompassing the global image level, the regional token level, and the concept subspace level as shown in Figure 1. Specifically, image-level alignment encourages the model to learn the correspondences between images and concepts from a global perspective. Token-level alignment focuses on the similarity between sub-regions within images and concept tokens using an attention-based mechanism. Concept-level alignment leverages concept activation vectors (CAVs) (Kim et al. 2018) to project the concept-based attention-weighted image representations to a concept subspace and subsequently aligns the aggregated image representations with clinical concepts. It is noteworthy that since the utilized concepts are human-interpretable, we leverage the knowledge of a medical large language model (LLM) by employing it as a concept encoder to enable the model to comprehend the latent conceptual semantics. During disease diagnosis, our model detects the concepts before making the decision. In this manner, our method makes full use of concept-based medical semantics through multi-level image-concept joint learning and achieves better performance and interpretability.

We summarize our main contributions as follows: (1) We propose MICA, a novel explainable disease diagnosis framework that semantically aligns medical images and clinical concepts at three different levels, i.e., global image level, regional token level, and concept subspace level. (2) To the best of our knowledge, we are the first to encode dermoscopic concepts using medical LLM. (3) As an ante-hoc explainable framework, our method is capable of performing disease diagnosis and concept detection concurrently while offering both visual and textual explanations. (4) Experimental results on three skin lesion datasets show that our method achieves superior performance and label efficiency, benefiting from the high-quality semantic correlations between images and concepts learned within our framework.

2 Related Works

2.1 XAI & Concept-based Methods

With an increasing number of high-stake scenarios (e.g., healthcare, finance, law enforcement) requiring trustwor-

thiness, XAI has been gaining attraction. One general approach in XAI is post-hoc analysis, which aims to interpret a trained model by fitting explanations to the model outputs, such as LIME (Ribeiro, Singh, and Guestrin 2016), SHAP (Lundberg and Lee 2017), and SENN (Alvarez Melis and Jaakkola 2018). Particularly for CNN, many researchers focus on saliency visualization (Zhou et al. 2016; Selvaraju et al. 2017; Sundararajan, Taly, and Yan 2017) and activation maximization (Van Den Oord, Kalchbrenner, and Kavukcuoglu 2016; Yosinski et al. 2015; Nguyen et al. 2016). However, post-hoc methods, which typically provide explanations based on pixels, regions or features of input images, do not genuinely enable medical experts or patients to understand which specific symptoms contribute to the decision process. This premise has sparked researchers' interest in the exploration of concept-based methods that integrate high-level human-interpretable concepts into decision process. Several researchers work on automatically discovering the concepts (Lang et al. 2021; Yeh et al. 2020), which can reduce the need for concept annotations but may not be suitable for healthcare, since the semantic meanings of discovered concepts can be unclear and unreliable. Concept activation vectors (CAVs) like approaches (Kim et al. 2018; Lucieri et al. 2020; Patrício et al. 2023; Yan et al. 2023) train linear classifiers, e.g., SVM (Cortes and Vapnik 1995), to the model's features to verify whether the representations can separate the human-defined concept examples. The inherently interpretable Concept Bottleneck Model (Koh et al. 2020; Patrício et al. 2023; Rigotti et al. 2021) first predicts concepts, then uses the detected concepts to predict task labels. We argue CBM is an essential research direction in trustworthy medical image analysis, since it mimics the process wherein medical experts first assess symptoms before diagnosing diseases during clinical treatments.

2.2 Trustworthy Skin Disease Diagnosis

The diagnosis of skin diseases, especially skin cancer, has been a significant research area within the intersection of deep learning and healthcare. Many of explanation approaches for skin lesion diagnosis are based on saliency map (Young et al. 2019; Xiang and Wang 2019) and attention mechanisms (Barata, Celebi, and Marques 2021; Gu et al. 2020). However, considering the stringent demands for model decision interpretability in healthcare (Lipton 2017), some researchers have taken efforts in designing concept-based models upon the ABCD-rule (Nachbar et al. 1994) and the 7-point checklist (Argenziano et al. 1998), which are authoritative criteria established by dermatologists. For instance, Lucieri et al. (2020) predict dermoscopic concepts from a pre-trained network to explain its predictions using TCAV. Coppola et al. (2020) propose to predict dermoscopic features with information sharing between different subnetworks to increase interpretability through multi-task learning. Yan et al. (2023) discover and eliminate confounding concepts within the datasets using spectral relevance analysis (Lapuschkin et al. 2019). CBE (Patrício et al. 2023) uses an extra segmentation module to preprocess images and encourages the feature maps obtained by 1×1 convolutional kernels to learn the representations of each dermo-

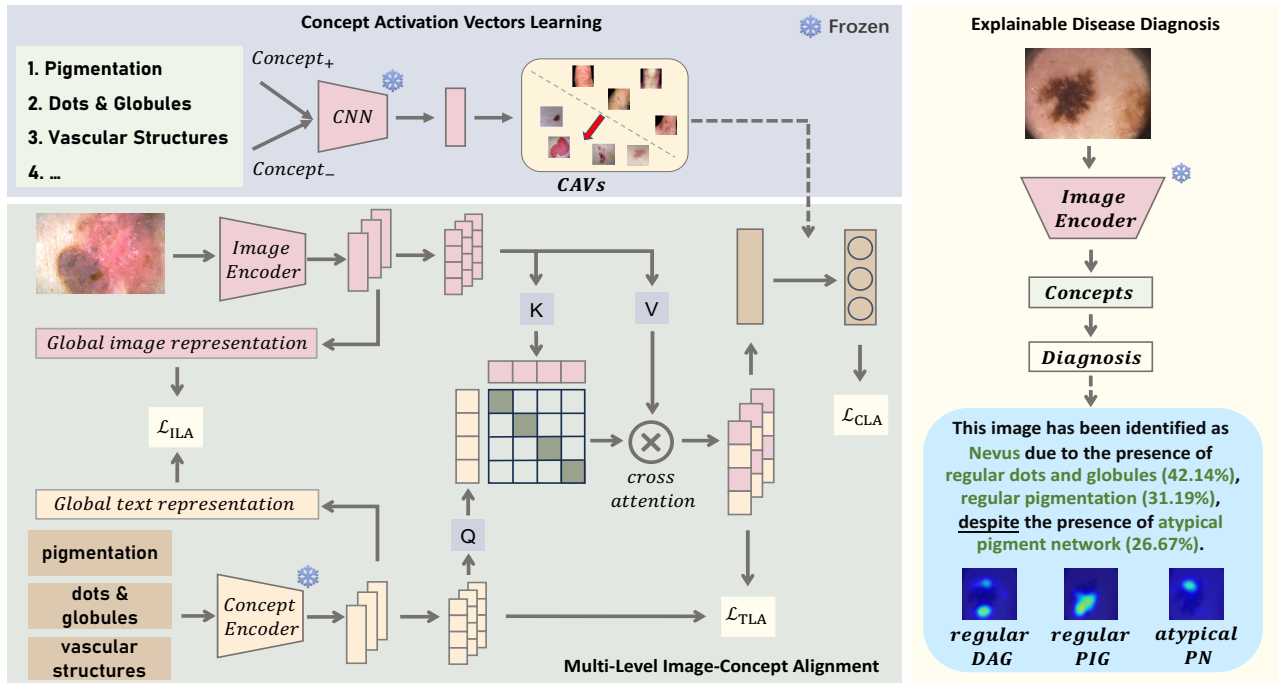


Figure 2: The overall pipeline of our proposed framework.

scopic concept, then diagnoses diseases in CBM architecture. However, most existing methods primarily utilize concept annotations from a single perspective and focus on the analysis of dermoscopic images. In contrast, our method semantically aligns skin images and concepts at multiple levels without using extra models to obtain disease masks, which can be applied to both dermoscopic images and raw, clinical images, e.g., SkinCon dataset (Daneshjou et al. 2022).

3 Method

3.1 Overall Framework

Figure 2 presents the overall architecture of our multi-level image-concept alignment framework for explainable disease diagnosis. Our method mainly consists of two stages: multimodal representation learning through image-concept alignment and explainable disease diagnosis. Specifically, in the first stage, we utilize a CNN-based image encoder and a large language model (LLM)-based concept encoder to extract semantic visual and textual features from the input medical images and corresponding clinical-related concepts. Then we align the images and concepts at three levels, i.e., image level, token level and concept level, directing the feature extractor to more effectively leverage the correspondences between images and concepts. To elaborate, we employ an image-level alignment module to maximize the similarity of the global representation between accurate image-concept pairs versus random pairs. Then, an attention-based token-level image-concept alignment module is proposed to cultivate fine-grained alignments between image sub-regions and concept word tokens. Moreover, to further refine the image and concept matches established by

the first two modules, we introduce a concept-level alignment module based on concept activation vectors (Kim et al. 2018), which maps the aggregated attention-weighted image representation onto the concept subspace and subsequently enhances the match with the concept ground truth. In the second stage, we add a single layer on top of the image encoder trained in the first stage to predict the concepts, and use the detected concepts to predict the diagnosis through a linear probe. Finally, concept-based explanations, including concept contributions and localization, are generated.

3.2 Multi-Level Image-Concept Alignment

Image-level Image-concept Alignment. To encourage the model to learn global correspondences between images and concepts, we employ an image-level alignment module. Specifically, given training set $D = \{(I_1, C_1), (I_2, C_2), \dots, (I_N, C_N)\}$, where (I_i, C_i) denotes the i -th image-concept pairs, and N denotes the number of training samples, we adopt a CNN-based image encoder $E_I : I \rightarrow V$, which takes an image I_i as input and outputs the global average pooling result of the last layer’s feature map, i.e., the visual representation V_i . To extract the semantic representations of concepts, we utilize a medical LLM-based concept encoder $E_C : C \rightarrow \{t, T\}$, which outputs the encoded concept text tokens t_i and the aggregated representation T_i given the concept C_i . The same dimensional image representation F_I and concept representation A_I are obtained using two projection layers that transform V and T into latent space embeddings, respectively. Then, the similarity $s_{ij} = (F_{I_i})^T A_{I_j}$ between image representation F_{I_i} and concept representation A_{I_j} is calculated by cosine similarity. To maximize the alignment between the correct pairs

of images and concept versus random pairs, we follow previous work (Zhang et al. 2022) to maximize the posterior probability of the image representation given its corresponding concept representation using contrastive loss function:

$$L_I^{(v|t)} = \sum_{i=1}^N -\log \frac{\exp(s_{ii}/\tau_1)}{\sum_{j=1}^N \exp(s_{ij}/\tau_1)}, \quad (1)$$

where τ_1 is a scaling temperature parameter. Considering the mutual relationship between the image and concept pairs, we also maximize the posterior probability of the concept given its corresponding image by minimizing the symmetrical contrastive loss function:

$$L_I^{(t|v)} = \sum_{i=1}^N -\log \frac{\exp(s_{ii}/\tau_1)}{\sum_{j=1}^N \exp(s_{ji}/\tau_1)} \quad (2)$$

The overall objective function of the image-level image-concept alignment module \mathcal{L}_{ILA} is the average of the two symmetrical losses.

Token-level Image-concept Alignment. In medical imaging, the primary focus is often on specific, small regions within an image. Diagnosis typically hinges on the symptoms observable within these particular image regions, which is different from the natural image domain. Therefore, considering that different clinical concepts may correspond to distinct sub-regions of a medical image, we propose the attention-based token-level image-concept alignment module. This module aims to address the limitation of image-level alignment and additionally explore the correlation between the image features of sub-regions and the respective concept tokens. Given the image encoder E_I , we extract the region-level visual feature maps from the intermediate convolution layer and vectorize to get the features of each image sub-region. A projection layer is also applied to the features to get the low-dimensional embeddings F_T . Similarly, the concept token features t extracted by concept encoder E_C are projected into token-level text embeddings A_T . We then calculate the concept-based attention-weighted image representation g_i using the cross attention between the region-level visual embeddings F_{T_i} and the token-level concept embedding A_{T_i} . Specifically, we regard A_T as queries, F_T as keys and values, then the attention weight α_{ij} is obtained by calculating the following formulation:

$$\alpha_{ij} = \frac{\exp((F_{T_i})^T A_{T_j}/\tau_2)}{\sum_{k=1}^M \exp((F_{T_i})^T A_{T_k}/\tau_2)}, \quad (3)$$

where τ_2 is a scaling temperature parameter. Then the concept-based attention weighted image representation $g_i = \sum_{j=0}^M \alpha_{ij} F_{T_j}$ is the weighted sum of sub-region features.

In order to align image and concept representations at the token level, we aggregate the similarity between all W concept features and their corresponding attention-weighted image representations using the token-wise matching function:

$$G(g_i, A_{T_i}) = \log \left(\sum_{i=1}^W \exp(\langle g_i, A_{T_i} \rangle / \tau_3) \right)^{\tau_3} \quad (4)$$

where τ_3 is a scaling temperature parameter. Then, similar to the image-level alignment module, we define the token-level image-concept alignment contrastive losses as:

$$L_T^{(v|t)} = \sum_{i=1}^N -\log \frac{\exp(G(g_i, A_{T_i})/\tau_2)}{\sum_{j=1}^N \exp(G(g_i, A_{T_j})/\tau_2)},$$

$$L_T^{(t|v)} = \sum_{i=1}^N -\log \frac{\exp(G(g_i, A_{T_i})/\tau_2)}{\sum_{j=1}^N \exp(G(g_j, A_{T_i})/\tau_2)} \quad (5)$$

The overall objective of the token-level alignment module \mathcal{L}_{TLA} is the average of the two symmetrical losses.

Concept-level Image-concept Alignment. The ILA and TLA modules encode concepts based on the knowledge of the LLM, yet they come with certain limitations: (1) For the sake of efficiency, the parameters of the LLM are fixed in our framework; (2) We cannot guarantee that all knowledge derived from the LLM is entirely accurate. To alleviate these issues, we design a concept-level image-concept alignment module to fully exploit concept annotations directly sourced from the labels for more effective learning of the correlation and to further refine the alignment of attention-weighted image representations and concepts. To enhance the matching in the concept subspace, we initially make use of concept activation vectors (CAVs) to learn concept representations.

Specifically, given concepts $C = \{c_1, c_2, \dots, c_{N_c}\}$, where c_i denotes the i -th concept (e.g., *Blue-Whitish Veil*), N_c denotes the number of concepts, we first split the concept samples $S^c = \{P^c, N^c\}$ into positive concept examples $P^c = \{P_i^c\}_{i=1}^{N_p}$ and negative concept examples $N^c = \{N_i^c\}_{i=1}^{N_n}$, where P_i^c and N_i^c are the CNN features of the images that contain or not contain the given concept c , respectively. N_p and N_n denote the number of positive and negative examples, respectively. We train an SVM to obtain the classification boundary, which separates features in P^c and N^c . We learn the corresponding CAV \mathbf{b} with weights ω^c and bias ϕ^c , defined as the vector normal to the boundary, which satisfies $(\omega^c)^T P_i^c + \phi^c > 0$ and $(\omega^c)^T N_i^c + \phi^c < 0$ for all examples. Given the concept-based attention weighted image representation g_i and the CAV $\mathbf{b} \in \mathbb{R}^{d \times N_c}$, where d is the dimension of concept subspace, we project g_i onto the concept subspace spanned by concept vectors using b :

$$h_i = \frac{\langle g_i, \mathbf{b} \rangle}{\|\mathbf{b}\|^2} \mathbf{b}, \quad (6)$$

where h_i is the projected concept embedding, which can also be regarded as concept scores. Then, cross-entropy loss is applied to estimate the discrepancy between the concept scores and the concept ground truth:

$$\mathcal{L}_{CLA} = - \sum_{i=1}^N (C_i \log(h_i) + (1 - C_i) \log(1 - h_i)) \quad (7)$$

Overall Objective In the first stage of our method, we introduce Multi-Level Image-Concept Alignment modules to make full use of the concept annotations and jointly learn the correspondences between images and concepts. The overall training objective of the first stage is represented as:

$$\mathcal{L} = \lambda_1 * \mathcal{L}_{ILA} + \lambda_2 * \mathcal{L}_{TLA} + \lambda_3 * \mathcal{L}_{CLA}, \quad (8)$$

where λ_1 , λ_2 , and λ_3 are hyperparameters.

METHOD	Derm7pt			PH ²			SkinCon		
	AUC (%)	ACC (%)	F1 (%)	AUC (%)	ACC (%)	F1 (%)	AUC (%)	ACC (%)	F1 (%)
Sarkar et al. (2022)	76.22 _{2.06}	73.89 _{1.47}	66.81 _{1.23}	79.33 _{0.62}	88.00 _{3.26}	79.66 _{2.11}	68.21 _{1.44}	71.14 _{1.21}	71.32 _{1.38}
PCBM (2023)	72.96 _{2.19}	76.98 _{1.39}	71.04 _{1.15}	78.33 _{1.17}	89.33 _{1.89}	81.49 _{2.57}	68.94 _{1.59}	71.04 _{1.13}	70.47 _{0.75}
PCBM-h (2023)	83.27 _{1.14}	79.89 _{0.89}	74.48 _{1.37}	92.32 _{1.47}	90.67 _{1.89}	83.30 _{2.55}	69.53 _{1.67}	72.28 _{1.39}	72.28 _{1.29}
CBE (2023)	76.60 _{0.35}	<u>83.75</u> _{0.26}	<u>78.13</u> _{0.44}	97.50 _{0.00}	<u>96.00</u> _{0.00}	93.89 _{0.00}	72.75 _{1.15}	73.75 _{1.10}	73.56 _{1.31}
MICA (w bot)	<u>84.11</u> _{1.10}	82.20 _{1.31}	78.08 _{1.22}	<u>97.66</u> _{1.24}	<u>96.00</u> _{3.26}	<u>94.40</u> _{1.48}	<u>75.89</u> _{1.11}	<u>74.29</u> _{1.09}	<u>74.74</u> _{1.21}
MICA (w/o bot)	85.59 _{1.11}	83.94 _{0.99}	79.38 _{1.34}	98.18 _{1.43}	98.67 _{1.89}	95.34 _{1.17}	75.92 _{1.13}	75.63 _{1.07}	75.43 _{1.24}

Table 1: Quantitative comparison on disease diagnosis with the concept-based state-of-the-arts. The performance is reported as mean_{std} of three random runs. Our method is highlighted, where *w bot* and *w/o bot* represent with or without concept bottleneck, respectively. The best and second-best results are highlighted in bold and underlined, respectively.

3.3 Explainable Disease Diagnosis

As illustrated in Figure 2, the second stage of our method is to perform an explainable disease diagnosis. In order to mimic the process wherein medical experts first assess symptoms before making a final diagnosis during clinical treatments, we propose an explainable decision module with a concept bottleneck. Given the image encoder, which learned the image-concept association in the first stage, we initially utilize it to predict the clinical concepts in the images, i.e., detecting the presence, absence, and types of various symptoms within the images. Then, we apply a linear predictor that maps the concept subspace to the disease prediction based on the detected concepts. It is worth noting that the linear predictor is highly interpretable because its decision is based on the detected clinical concepts, which is consistent with human medical experts. In addition, the weight matrix of the linear predictor denotes the importance of each concept to the final decision. Moreover, a human medical expert can easily edit the predictor to get a more reliable diagnosis when observing a wrong or counter-intuitive phenomenon. Specifically, given the image representation v_i , of an input image I_i encoded by the freezing image encoder E_c , we first detect the concepts and get the concept scores through an FC layer f_c , then the disease classification is performed based on the detected concepts through a linear layer f_d . We jointly train both the concept detection and disease classification layer through the following objective function:

$$\hat{f}_c, \hat{f}_d = \arg \min_{f_c, f_d} \sum_{i=1}^N [CE(f_c(v_i), C_i) + \beta CE(f_d(f_c(v_i)), y_i)], \quad (9)$$

where $CE(\cdot)$ is the cross-entropy loss, y_i is the diagnosis label of the i -th image, and β is the hyperparameter to balance the concept detection and disease diagnosis.

Discussion What happens when we directly perform disease diagnosis using the trained image encoder without a concept bottleneck? Ideally, we would like to achieve higher performance while retaining the explainability of the concept-based methods. Thus, we simply apply a classification head on top of the image encoder to predict diagnosis labels.

Since we only use the concept annotations when training the image encoder, it can be regarded as our model leveraging the concept knowledge to perform disease diagnosis, which is still an emulation of the *image* \rightarrow *concept* \rightarrow *diagnosis* decision-making process.

4 Experiments

4.1 Experimental Setup

Datasets: *Derm7pt* (Kawahara et al. 2018) is a dermoscopic image dataset contains 1011 images with clinical concepts for melanoma skin lesions in dermatology. We consider all 7 dermoscopic concepts. Following Patrício et al. (2023), we filter the dataset to obtain 827 images of *Nevus* and *Melanoma* classes. Only the dermoscopic images are used. The specific names of concepts can be found in section 4.2. *PH²* (Mendonça et al. 2013) contains a total of 200 dermoscopic images of melanocytic lesions, including 80 common nevi, 80 atypical nevi, and 40 melanomas. We consider 5 concepts that are also included in the *Derm7pt* dataset. Following previous work (Patrício et al. 2023), we combine the *Common Nevi* and *Atypical Nevi* classes of the *PH²* dataset into one global class label called *Nevus*. *SkinCon* (Daneshjou et al. 2022) is a skin disease dataset with 3230 images densely annotated by experts for fine-grained model debugging and analysis. We choose 22 concepts that have at least 50 images representing the concept from the *F17k* (Groh et al. 2021) part. The classification categories are *malignant*, *benign* and *non-neoplastic*. The dataset is split into training set, validation set and test set according to the proportion of 70%, 15% and 15%, respectively.

Compared Approaches: **Sarkar et al.** (2022) design an ante-hoc model where the output of the concept encoder is passed to a decoder that reconstructs the image, encouraging the model to capture the semantic features of the input image. **PCBM(-h)** (Yuksekonul et al. 2023) allows transferring concepts from other datasets and designs a residual modeling step to preserve performance of CBM. **CBE** (Patrício et al. 2023) proposes a coherence loss to improve the visual coherence of concept activations. **CAV** (Lucieri et al. 2020) uses concept activation vectors to perform a detailed concept analysis for skin tumor classification.

Dataset	Method	AUC (%)	ACC (%)	F1 (%)
Derm7pt	CAV (2020)	73.8	71.2	61.5
	CBE (2023)	72.2	74.1	71.0
	MICA (Ours)	78.6	76.0	72.6
PH ²	CAV (2020)	71.2	67.2	66.4
	CBE (2023)	81.3	71.6	70.0
	MICA (Ours)	83.6	75.2	68.4
SkinCon	CAV (2020)	76.5	86.4	60.2
	CBE (2023)	79.3	89.0	62.1
	MICA (Ours)	82.6	91.7	63.8

Table 2: Quantitative results in concept detection.

Dataset	Method	PN	DaG	STR	RS	BWV	PIG	VS	Avg.
Derm7pt	CAV	72.7	69.1	74.3	65.5	74.6	68.6	73.8	71.2
	CBE	78.5	64.5	77.0	71.0	81.1	75.4	73.7	74.1
	MICA	74.4	79.1	71.3	72.8	84.4	68.8	81.3	76.0
PH ²	CAV	68.0	60.0	56.0	76.0	76.0	N/A	N/A	67.2
	CBE	64.0	58.0	80.0	80.0	76.0	N/A	N/A	71.6
	MICA	60.0	60.0	80.0	88.0	88.0	N/A	N/A	75.2

Table 3: Comparison of concept detection accuracy (%). PN, DaG, etc. denote each clinical concept of the corresponding dataset and Avg. presents the average values of the previous columns. The best results are highlighted in bold.

Implementation Details: Our framework uses ResNet-50 (He et al. 2016) as the image encoder. For concept encoder, we use a BERT encoder (Devlin et al. 2018) with the ClinicalBERT weights (Alsentzer et al. 2019). During training in the first stage (i.e., multi-level alignment), we train the image encoder using only concept labels. In the second stage (i.e., disease diagnosis), the parameters of the image encoder are fixed, and we only train the classification heads. We adopt Adam (Kingma and Ba 2014) optimizer with learning rate of $5e-5$ in the first stage and $1e-4$ in the second stage. For the hyperparameter selection, we use grid search and set $\tau_1 = 0.25$, $\tau_2 = 0.2$, $\tau_3 = 0.1$. We set $\beta = 1$ for *Derm7pt* and *SkinCon* dataset, and set $\beta = 0.5$ for *PH²* dataset.

4.2 Experimental Results

To demonstrate that our method’s competitive performance on disease diagnosis and concept detection, we first compare with other state-of-the-art concept-based approaches on various datasets. Then we conduct comprehensive ablation experiments to validate the effectiveness of each module designed in our method. Finally, we evaluate the explainability of our method using multiple XAI metrics.

Disease Diagnosis. In Table 1, we report the classification comparison results of our method under three metrics (AUROC, Accuracy and F1 Score) on the considered datasets. MICA outperforms other methods in overall performance especially AUC. As discussed in section 3.3, the results of MICA without concept bottleneck are also presented.

ILA	TLA	CLA	Derm7pt		PH ²		SkinCon	
			AUC _D	AUC _C	AUC _D	AUC _C	AUC _D	AUC _C
✗	✗	✗	71.2	63.1	88.3	70.6	67.4	70.7
✓	✗	✗	81.1	75.5	95.0	80.5	72.9	78.5
✓	✓	✗	82.8	77.3	95.6	84.5	74.7	81.5
✓	✗	✓	83.3	78.0	96.4	82.7	75.1	81.8
✗	✓	✓	82.9	77.7	96.1	82.9	75.0	82.1
✓	✓	✓	84.1	78.6	97.7	83.6	75.9	82.6

Table 4: Ablation study of MICA on disease diagnosis (AUC_D [%]) and concept detection (AUC_C [%]). ILA, TLA, and CLA represent the image-level, token-level, and concept-level alignment modules, respectively.

Concept Detection. Table 2 shows the quantitative results of clinical concept detection. MICA outperforms other methods at most metrics in considered datasets. We also report the test classification accuracy of each concept in *Derm7pt* and *PH²* dataset in Table 3. These two datasets have five common dermoscopic concepts from Seven-point Checklist (Argenziano et al. 1998), where PN stands for “Pigment Network”, DaG stands for “Dots and Globules”, STR stands for “Streaks”, RS stands for “Regression Structures”, BWV stands for “Blue-Whitish Veil”, PIG stands for “Pigmentation”, VS stands for “Vascular Structures”. Test accuracy of each concept reported is the average of the fine-grained labels that belong within each criteria. For example, PN (Pigment Network) denotes the mean accuracy of subclasses “Typical Pigment Network” and “Atypical Pigment Network”. The “avg” column reports the mean value of the test accuracy of all concepts.

Ablation Study. As shown in Table 4, we observe that our framework can benefit from all three alignment modules, including ILA, TLA and CLA. The ablation results show that without any one of the three modules, the performance of both disease diagnosis and concept detection may suffer. The last configuration of Table 4 demonstrates that our method achieves the best overall performance with all three designed image-concept alignment components.

4.3 Analysis of Explainability

In this section, we evaluate and analyze the explainability of our method. Specifically, inspired by prior research (Hsiao et al. 2021; Guidotti et al. 2018; Rigotti et al. 2021; Jin et al. 2023), we outline and evaluate our framework on multiple essential metrics for XAI techniques, including *faithfulness*, *plausibility (understandability)* and *efficiency*.

Faithfulness. *Faithfulness* requires that the explanation should be highly faithful with the designed model mechanism and thus reflects the model decision process (Lakkaraju et al. 2019). In this paper, we employ test-time intervention on concepts to assess faithfulness. During inference time, we first predict the concepts and obtain the corresponding concept scores, then we intervene on a concept by changing its value. The diagnosis labels are predicted based on the

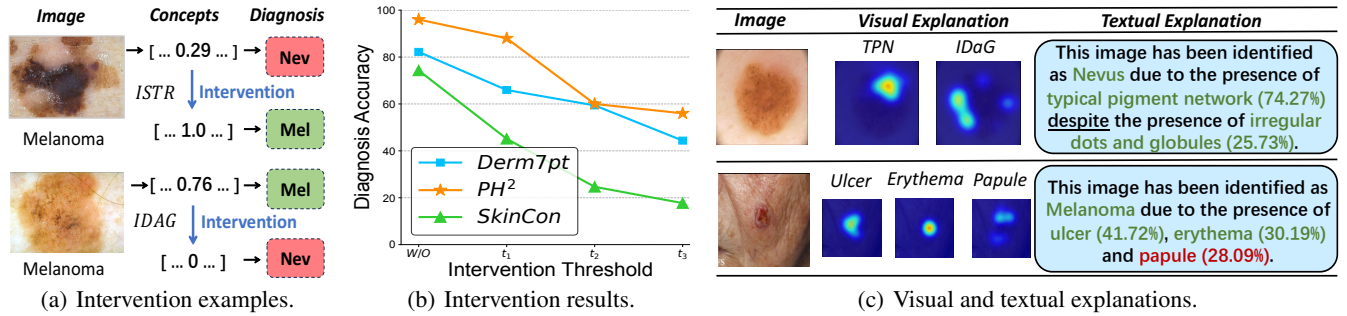


Figure 3: Illustration of our model’s *faithfulness*, *plausibility* and *understandability*. (a)(b) Test time concept-intervention examples and results. (c) Examples of visual and textual explanations provided by our method given skin images from different datasets. Correct prediction results are marked in green, while red highlights incorrect predictions.

concepts after intervention. Figure 3(a) shows two cases of test-time intervention. In the first case, we replace the concept score with the ground-truth label of *irregular streaks* and subsequently the diagnosis result changes from *Nevus* to *Melanoma*, which amends the diagnosis prediction. In the second case of Figure 3(a), we set the correct predicted value of concept *irregular dots and globules* to 0 and the model decision changes from *Melanoma* to *Nevus*, which is consistent with the dermatologists’ findings (Argenziano et al. 1998; Kawahara et al. 2018). We also report test-time intervention results in Figure 3(b), which reflects the change of diagnosis accuracy when we zero those above the threshold values. The accuracy decreases when the threshold values decrease demonstrates that the predicted concepts are faithfully explaining the model’s decisions.

Plausibility & Understandability. *Plausibility* refers to how believable or likely the given explanation seems, given what human-being know about the world or the domain of the problem (Carvalho, Pereira, and Cardoso 2019; Guidotti et al. 2018), while *understandability* refers to how easily a human user can comprehend the provided explanation without requiring technical knowledge (Jin et al. 2023). In this paper, our model achieves both *plausibility* and *understandability* by providing concept-based visual and textual explanations. Figure 3(c) shows the examples of explanations in detail. Given an input image, our method predicts the diagnosis labels with each predicted concept’s localization and contribution (in %). The concept localization is generated by visualizing the token-level correspondences with the images. As for the concept contribution, we leverage the softmax result of the last linear layer, which multiplies the concept logits and the corresponding weights. Our framework aggregates the prediction results to get textual explanations. It is worth noting that a “despite” (underlined) is used in case of negative class fluence to signalize contradiction.

Efficiency. High label efficiency can allow the explainable model to be practically implemented in real-world applications without using extra data or annotations. In this paper, we assess our model’s label efficiency by using different proportions of training data in the second stage. As shown in Table 5, we report the disease diagnosis results using 100%,

Dataset	Derm7pt		PH ²		SkinCon	
Num#	34 / 173 / 346		15 / 75 / 150		225 / 1126 / 2252	
PCT (%)	10 / 50 / 100		10 / 50 / 100		10 / 50 / 100	
Proportion	AUC	ACC	AUC	ACC	AUC	ACC
10%	77.9	76.9	65.7	81.3	70.1	71.5
50%	83.2	81.0	96.5	93.3	73.5	73.2
100%	84.1	82.2	97.7	96.0	75.9	74.3

Table 5: Disease diagnosis performance based on different portion of training data. PCT represents percentage.

50%, and 10% of training data. For *Derm7pt* and *SkinCon* datasets, it can be observed that the diagnosis performance of the model does not exhibit significant decline when only 50% or 10% of the diagnosis labels are used. The obvious decrease for *PH²* dataset with 10% training data may be because only 15 labels are used to train the classifier. Therefore, our method can achieve competitive diagnosis results only using a small proportion of diagnosis labels, signifying that our method encourages the model to learn the correspondences between medical images and clinical-related concepts, thus facilitating disease diagnosis effectively. This experimental observation is faithfully consistent with the doctors’ diagnostic process wherein medical experts make a diagnosis decision based on the detected symptoms.

5 Conclusion

In this paper, we propose MICA, a multi-modal explainable concept-based framework for skin disease diagnosis, which semantically aligns medical images and clinical-related concepts at multiple levels. By thoroughly learning the correspondences between images and concepts at the global image level, regional token level, and concept subspace level, our method outperforms other concept-based models while preserving inherent interpretability and offering both visual and textual explanations. Extensive experiments and explainability analysis conducted on skin image datasets demonstrate that our method simultaneously achieves superior performance, label efficiency, and interpretability.

Acknowledgments

This work was supported by the Hong Kong Innovation and Technology Fund (Project No. ITS/028/21FP), Shenzhen Science and Technology Innovation Committee Fund (Project No. SGDX20210823103201011), and the Project of Hetao Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone (HZQB-KCZYB-2020083).

References

- Abbasi-Asl, R.; and Yu, B. 2017. Structural compression of convolutional neural networks. *arXiv preprint arXiv:1705.07356*.
- Alsentzer, E.; Murphy, J. R.; Boag, W.; Weng, W.-H.; Jin, D.; Naumann, T.; and McDermott, M. 2019. Publicly available clinical BERT embeddings. *arXiv preprint arXiv:1904.03323*.
- Alvarez Melis, D.; and Jaakkola, T. 2018. Towards robust interpretability with self-explaining neural networks. *Advances in neural information processing systems*, 31.
- Argenziano, G.; Fabbrocini, G.; Carli, P.; De Giorgi, V.; Sammarco, E.; and Delfino, M. 1998. Epiluminescence microscopy for the diagnosis of doubtful melanocytic skin lesions: comparison of the ABCD rule of dermatoscopy and a new 7-point checklist based on pattern analysis. *Archives of dermatology*, 134(12): 1563–1570.
- Barata, C.; Celebi, M. E.; and Marques, J. S. 2021. Explainable skin lesion diagnosis using taxonomies. *Pattern Recognition*, 110: 107413.
- Carvalho, D. V.; Pereira, E. M.; and Cardoso, J. S. 2019. Machine learning interpretability: A survey on methods and metrics. *Electronics*, 8(8): 832.
- Coppola, D.; Lee, H. K.; and Guan, C. 2020. Interpreting mechanisms of prediction for skin cancer diagnosis using multi-task learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 734–735.
- Cortes, C.; and Vapnik, V. 1995. Support-vector networks. *Machine learning*, 20: 273–297.
- Daneshjou, R.; Yuksekgonul, M.; Cai, Z. R.; Novoa, R.; and Zou, J. Y. 2022. Skincon: A skin disease dataset densely annotated by domain experts for fine-grained debugging and analysis. *Advances in Neural Information Processing Systems*, 35: 18157–18167.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Esteva, A.; Kuprel, B.; Novoa, R. A.; Ko, J.; Swetter, S. M.; Blau, H. M.; and Thrun, S. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *nature*, 542(7639): 115–118.
- Fang, Z.; Kuang, K.; Lin, Y.; Wu, F.; and Yao, Y.-F. 2020. Concept-based explanation for fine-grained images and its application in infectious keratitis classification. In *Proceedings of the 28th ACM international conference on Multimedia*, 700–708.
- Groh, M.; Harris, C.; Soenksen, L.; Lau, F.; Han, R.; Kim, A.; Koochek, A.; and Badri, O. 2021. 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*, 1820–1828.
- Gu, R.; Wang, G.; Song, T.; Huang, R.; Aertsen, M.; Deprest, J.; Ourselin, S.; Vercauteren, T.; and Zhang, S. 2020. CA-Net: Comprehensive attention convolutional neural networks for explainable medical image segmentation. *IEEE transactions on medical imaging*, 40(2): 699–711.
- Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Gianotti, F.; and Pedreschi, D. 2018. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5): 1–42.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hsiao, J. H.-w.; Ngai, H. H. T.; Qiu, L.; Yang, Y.; and Cao, C. C. 2021. Roadmap of designing cognitive metrics for explainable artificial intelligence (XAI). *arXiv preprint arXiv:2108.01737*.
- Jin, W.; Li, X.; Fatehi, M.; and Hamarneh, G. 2023. Guidelines and evaluation of clinical explainable AI in medical image analysis. *Medical Image Analysis*, 84: 102684.
- Kawahara, J.; Daneshvar, S.; Argenziano, G.; and Hamarneh, G. 2018. Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *IEEE journal of biomedical and health informatics*, 23(2): 538–546.
- Kermany, D. S.; Goldbaum, M.; Cai, W.; Valentim, C. C.; Liang, H.; Baxter, S. L.; McKeown, A.; Yang, G.; Wu, X.; Yan, F.; et al. 2018. Identifying medical diagnoses and treatable diseases by image-based deep learning. *cell*, 172(5): 1122–1131.
- Kim, B.; Wattenberg, M.; Gilmer, J.; Cai, C.; Wexler, J.; Viegas, F.; et al. 2018. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, 2668–2677. PMLR.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Koh, P. W.; and Liang, P. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*, 1885–1894. PMLR.
- Koh, P. W.; Nguyen, T.; Tang, Y. S.; Mussmann, S.; Pierson, E.; Kim, B.; and Liang, P. 2020. Concept bottleneck models. In *International conference on machine learning*, 5338–5348. PMLR.
- Lakkaraju, H.; Kamar, E.; Caruana, R.; and Leskovec, J. 2019. Faithful and customizable explanations of black box models. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 131–138.
- Lang, O.; Gandselman, Y.; Yarom, M.; Wald, Y.; Elidan, G.; Hassidim, A.; Freeman, W. T.; Isola, P.; Globerson, A.; Irani,

- M.; et al. 2021. Explaining in style: Training a gan to explain a classifier in stylespace. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 693–702.
- Lapuschkin, S.; Waldchen, S.; Binder, A.; Montavon, G.; Samek, W.; and Muller, K.-R. 2019. Unmasking Clever Hans predictors and assessing what machines really learn. *Nature communications*, 10(1): 1096.
- Laugel, T.; Lesot, M.-J.; Marsala, C.; Renard, X.; and Detyniecki, M. 2019. The dangers of post-hoc interpretability: Unjustified counterfactual explanations. *arXiv preprint arXiv:1907.09294*.
- Lipton, Z. C. 2017. The doctor just won’t accept that! *arXiv preprint arXiv:1711.08037*.
- Litjens, G.; Kooi, T.; Bejnordi, B. E.; Setio, A. A. A.; Ciompi, F.; Ghafoorian, M.; Van Der Laak, J. A.; Van Ginneken, B.; and Sanchez, C. I. 2017. A survey on deep learning in medical image analysis. *Medical image analysis*, 42: 60–88.
- Lucieri, A.; Bajwa, M. N.; Braun, S. A.; Malik, M. I.; Dengel, A.; and Ahmed, S. 2020. On interpretability of deep learning based skin lesion classifiers using concept activation vectors. In *2020 international joint conference on neural networks (IJCNN)*, 1–10. IEEE.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Mendonça, T.; Ferreira, P. M.; Marques, J. S.; Marcal, A. R.; and Rozeira, J. 2013. PH 2-A dermoscopic image database for research and benchmarking. In *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, 5437–5440. IEEE.
- Nachbar, F.; Stolz, W.; Merkle, T.; Cognetta, A. B.; Vogt, T.; Landthaler, M.; Bilek, P.; Braun-Falco, O.; and Plewig, G. 1994. The ABCD rule of dermatoscopy: high prospective value in the diagnosis of doubtful melanocytic skin lesions. *Journal of the American Academy of Dermatology*, 30(4): 551–559.
- Nguyen, A.; Dosovitskiy, A.; Yosinski, J.; Brox, T.; and Clune, J. 2016. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. *Advances in neural information processing systems*, 29.
- Patricio, C.; Neves, J. a. C.; Teixeira, L. F.; et al. 2023. Coherent Concept-Based Explanations in Medical Image and Its Application to Skin Lesion Diagnosis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 3798–3807.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. ” Why should i trust you?” Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.
- Rigotti, M.; Miksovic, C.; Giurgiu, I.; Gschwind, T.; and Scotton, P. 2021. Attention-based interpretability with concept transformers. In *International Conference on Learning Representations*.
- Rudin, C. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature machine intelligence*, 1(5): 206–215.
- Sarkar, A.; Vijaykeerthy, D.; Sarkar, A.; and Balasubramanian, V. N. 2022. A Framework for Learning Ante-Hoc Explainable Models via Concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10286–10295.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. In *International conference on machine learning*, 3319–3328. PMLR.
- Tsang, M.; Cheng, D.; and Liu, Y. 2017. Detecting statistical interactions from neural network weights. *arXiv preprint arXiv:1705.04977*.
- Van Den Oord, A.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. In *International conference on machine learning*, 1747–1756. PMLR.
- Xiang, A.; and Wang, F. 2019. Towards interpretable skin lesion classification with deep learning models. In *AMIA annual symposium proceedings*, volume 2019, 1246. American Medical Informatics Association.
- Yan, S.; Yu, Z.; Zhang, X.; Mahapatra, D.; Chandra, S. S.; Janda, M.; Soyer, P.; and Ge, Z. 2023. Towards Trustable Skin Cancer Diagnosis via Rewriting Model’s Decision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11568–11577.
- Yeh, C.-K.; Kim, B.; Arik, S.; Li, C.-L.; Pfister, T.; and Ravikumar, P. 2020. On completeness-aware concept-based explanations in deep neural networks. *Advances in neural information processing systems*, 33: 20554–20565.
- Yosinski, J.; Clune, J.; Nguyen, A.; Fuchs, T.; and Lipson, H. 2015. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*.
- Young, K.; Booth, G.; Simpson, B.; Dutton, R.; and Shrapnel, S. 2019. Deep neural network or dermatologist? In *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support: Second International Workshop, iMIMIC 2019, and 9th International Workshop, ML-CDS 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 17, 2019, Proceedings 9*, 48–55. Springer.
- Yuksekgonul, M.; Wang, M.; Zou, J.; et al. 2023. Post-hoc Concept Bottleneck Models. In *The Eleventh International Conference on Learning Representations*.
- Zhang, Y.; Jiang, H.; Miura, Y.; Manning, C. D.; and Langlotz, C. P. 2022. Contrastive learning of medical visual representations from paired images and text. In *Machine Learning for Healthcare Conference*, 2–25. PMLR.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.