



Research Article

A two-step concept-based approach for enhanced interpretability and trust in skin lesion diagnosis

Cristiano Patrício^{a,c,*}, Luís F. Teixeira^{b,c}, João C. Neves^a^a Universidade da Beira Interior and NOVA LINCS, Portugal^b Faculdade de Engenharia da Universidade do Porto, Portugal^c INESC TEC, Portugal

ARTICLE INFO

Keywords:

Concept bottleneck models
Vision-language models
Interpretability
Skin cancer
Dermoscopy

ABSTRACT

The main challenges hindering the adoption of deep learning-based systems in clinical settings are the scarcity of annotated data and the lack of interpretability and trust in these systems. Concept Bottleneck Models (CBMs) offer inherent interpretability by constraining the final disease prediction on a set of human-understandable concepts. However, this inherent interpretability comes at the cost of greater annotation burden. Additionally, adding new concepts requires retraining the entire system. In this work, we introduce a novel two-step methodology that addresses both of these challenges. By simulating the two stages of a CBM, we utilize a pretrained Vision Language Model (VLM) to automatically predict clinical concepts, and an off-the-shelf Large Language Model (LLM) to generate disease diagnoses grounded on the predicted concepts. Furthermore, our approach supports test-time human intervention, enabling corrections to predicted concepts, which improves final diagnoses and enhances transparency in decision-making. We validate our approach on three skin lesion datasets, demonstrating that it outperforms traditional CBMs and state-of-the-art explainable methods, all without requiring any training and utilizing only a few annotated examples. The code is available at <https://github.com/CristianoPatrício/2-step-concept-based-skin-diagnosis>.

1. Introduction

Recent advancements in automated computer-aided diagnosis systems for detecting diseases from medical images have demonstrated a significant improvement in performance, largely driven by the superior capabilities of deep learning models. This shift has considerably boosted the accuracy of these systems in delivering precise diagnoses across various medical imaging tasks, including skin lesion analysis, sometimes achieving results comparable to those of dermatologists [1,2]. Despite these advancements, state-of-the-art deep learning-based systems in dermatology operate as “black-boxes”, failing to provide explainable rationale behind the decision-making process. This lack of transparency remains a major obstacle to their widespread adoption and integration into clinical practice [3].

To address this issue, explainable models have been developed to enhance the transparency and reliability of medical AI systems. Among these, Concept Bottleneck Models (CBMs) [4] have gained considerable attention in medical imaging analysis [5–7], as they base their decisions on the presence or absence of human-interpretable concepts, closely resembling the way clinicians analyze medical images. Studies have

also shown that concept-based explanations are generally preferred by humans over other methods, such as heatmaps or example-based explanations [8]. However, despite their growing popularity, CBMs face two key challenges: i) they require detailed annotations of human-understandable concepts, a process that is both time-consuming and demands domain expertise; and ii) the model must be retrained whenever new concepts are introduced. While recent studies [9–12] have attempted to reduce the reliance on concept annotations by leveraging Large Language Models (LLMs) to generate candidate concepts, these approaches still require retraining when new target classes or concepts are added to the system.

In this work, we address the limitations of CBMs by proposing a novel two-step approach that provides concept-based explanations and generates disease diagnoses grounded in predicted concepts, all without the need for additional training. In the first stage, we utilize pretrained Vision-Language Models (VLMs) to predict the presence of a set of clinical concepts. In the second stage, these predicted concepts are incorporated into a custom designed prompt, which is used to query an LLM for a final diagnosis class to ensure the diagnosis is based on the identified concepts. Fig. 1 illustrates our methodology compared to

* Corresponding author.

E-mail address: cristiano.patricio@ubi.pt (C. Patrício).




Fig. 1. Overview of the proposed framework. The linear classifier layer (left) is replaced by a pretrained Large Language Model (LLM) (right), which grounds its responses on clinical concepts predicted by a pretrained Vision-Language Model (VLM). This approach is training-free and not restricted by predefined labels, allowing the LLM to generate diverse diagnostic possibilities for different diseases.

the traditional approach where a linear layer is trained to predict the final target class. Similar to CBMs, our approach allows for test-time intervention, enabling clinicians to correct concept predictions at inference time. However, since the predicted concepts by VLM are presented as plain text for the final diagnosis, this enables more intuitive human intervention compared to adjusting numerical values in the bottleneck layer, as in traditional CBM test-time interventions. Additionally, leveraging a VLM to predict concepts offers the advantage of incorporating both visual and textual information during model inference. In the other hand, unlike CBMs, our approach does not require training to provide the final diagnosis class and can be easily adapted to incorporate new concepts.

We validate our approach on three well-known skin lesion datasets, demonstrating that it outperforms traditional CBMs and state-of-the-art explainable methods, all without requiring any training and using only a few annotated examples within the prompt. Our contributions are as follows:

- We introduce a novel two-step approach for concept-based explainability in skin lesion diagnosis without requiring additional training;
- We reduce the annotation burden of CBMs by leveraging zero-shot capabilities of pretrained VLMs to automatically predict clinical concepts;
- We provide the final diagnosis class using tailored prompts and leveraging the few-shot capabilities of LLMs;
- Our approach enables human test-time intervention, allowing corrections to any mispredicted concepts;
- We demonstrate superior interpretability and performance compared to traditional CBMs on three public benchmarks.

The remainder of the paper is organized as follows: Section 2 reviews related work in the context of our study. Section 3 introduces the proposed two-step methodology. Section 4 outlines the experimental setup and implementation details. Section 5 presents key results and analysis. A general discussion is provided in Section 6. Finally, Section 7 and Section 8 discuss the limitations and conclude the paper, respectively.

2. Related work

2.1. Vision-language models

Vision-language models (VLMs) aim to learn joint representations from image-text pairs using contrastive learning. While models like CLIP [13] have demonstrated impressive zero-shot capabilities across a wide range of tasks, their application in the biomedical domain remains challenging due to distribution shifts and the domain specific vocabulary. To address these challenges, several medical-specific VLMs have been developed, such as BiomedCLIP [14], PubMedCLIP [15], and Med-CLIP [16], which are pretrained on large-scale image-text pairs from sources like PubMed articles and radiology datasets. However, despite these advancements, such models still underperform when compared to task-specific models [17] and lack transparency and interpretability in their decision-making processes. Our novel methodology leverages the zero-shot capabilities of medical-domain VLMs to automatically annotate concepts, forming the first step of our approach.

2.2. Explainable models for skin lesion diagnosis

Concept Bottleneck Models (CBMs) [4] first predict a set of human-interpretable concepts and then use these concepts to determine the final class label using a linear layer, thus offering transparency in the decision-making process. However, these models require extensive manual concept annotations and depend on a linear layer for the final diagnosis prediction. To address the need for annotations, several studies [9–12] have explored the use of Large Language Models (LLMs) to generate candidate concepts for target classes. Other methods [17–19] align visual features with pre-defined clinically relevant concepts using vision-language models (VLMs).

Despite these improvements, most models still rely on a linear classifier to predict the final diagnostic label, whether based on concepts or visual features. Our approach overcomes this by prompting an LLM to directly predict the diagnosis using a tailored prompt that incorporates the concepts extracted by a pretrained VLM. This eliminates the fixed label constraint, improving scalability and removing the need for retraining when new diagnostic categories or concepts are introduced.

Instruction

You're a doctor, make a good choice based on the question and options. You need to answer the letter of the option without further explanations.

Demonstrations

Consider the following examples:

1. The color is highly variable, often with multiple colors (black, brown, red, white, blue), the shape is irregular, the border is often blurry and irregular, the dermoscopic patterns are atypical pigment network, irregular streaks, blue-whitish veil, irregular, the texture is a raised or ulcerated surface, the symmetry is asymmetrical, the elevation is flat to raised. Thus the diagnosis is **melanoma**.
 2. The color is uniformly tan, brown, or black, the shape is round, the border is sharp and well-defined, the dermoscopic patterns are regular pigment network, symmetric dots and globules, the texture is smooth, the symmetry is symmetrical, the elevation is raised with possible central ulceration. Thus the diagnosis is **nevus**.
- (...)

Question

What is the type of skin lesion that is associated with the following dermoscopic concepts ?

The color is uniformly tan, brown, or black, the shape is round, the border is sharp and well-defined, the dermoscopic patterns are regular pigment network, symmetric dots and globules, the texture is smooth, the symmetry is symmetrical, the elevation is raised with possible central ulceration.

Options

- A. Nevus
- B. Melanoma

Answer:

Fig. 2. Example prompt with $K = 2$ demonstration examples. Few-shot prompting is expected to improve performance as the number of demonstration examples increases.

3. Methodology

Given the task of predicting a target disease $y \in \mathbb{R}$ from input $x \in \mathbb{R}^d$, let $\mathcal{D} = \{(x^{(i)}, y^{(i)}, c^{(i)})\}_{i=1}^n$ represent a batch of training samples, where $c \in \mathbb{R}^l$ is a vector of l clinical concepts. CBMs first map the input x into a set of interpretable concepts c (the “bottleneck”) by learning a function $g : \mathbb{R}^d \rightarrow \mathbb{R}^l$ ($x \rightarrow c$), and use these concepts to predict the target y through $f : \mathbb{R}^l \rightarrow \mathbb{R}$ ($c \rightarrow y$). As a result, the final prediction $\hat{y} = f(g(x))$ is entirely based on the predicted concepts $\hat{c} = g(x)$. To estimate the presence of clinical concepts from an input image, we utilize a pretrained VLM and calculate the cosine similarity between the image features and a set of clinical concepts in the concept set $C = \{c^{(0)}, \dots, c^{(i)}\}$ (see Section 3.1). Next, we prompt a pretrained LLM to generate the final diagnosis class grounded on the predicted concepts (detailed in Section 3.2). Fig. 1 provides an overview of our methodology, which is further described in the following subsections.

3.1. Concept prediction

The first stage focuses on predicting the presence of dermoscopic concepts from a given input image, representing the “bottleneck” of CBM: $x \rightarrow c$. Following previous work [17], we adopt a pretrained VLM and determine the presence of a dermoscopic concept c in the input image by assessing the similarity between the image feature embedding, $\mathcal{V}(x)$, and the feature embedding of each concept c . Formally, the similarity scores are given by:

$$\text{scores} = \text{sim}(\mathcal{V}(x), \mathcal{T}(c^{(i)})), \quad (1)$$

where $\text{sim}(\cdot)$ is a similarity metric (e.g., cosine similarity), $\mathcal{V}(\cdot)$ is the visual encoder, and $\mathcal{T}(\cdot)$ is the text encoder.

3.2. Disease classification

In the second stage, we prompt a Large Language Model (LLM) to generate the final disease diagnosis grounded on the predicted dermoscopic concepts, corresponding to the classification phase of CBM: $c \rightarrow y$. Using the concept scores generated in the first stage (Equation (1)), we binarize them using a threshold and map them to their respective concept names. These concepts are then incorporated into the designed prompt. An example of this prompt is provided on the right side of Fig. 1.

This method ensures that the diagnosis is based on the dermoscopic concepts, improving the interpretability and transparency of the model's output, and eliminates the need for training a linear layer, allowing for more flexible and varied diagnostic output formats.

3.3. Few-shot prompting

Few-shot prompting, also known as In-Context Learning (ICL), is a prompt engineering technique where the model is given with a few task demonstrations during inference within the prompt as conditioning [20]. This method has been shown to significantly boost the performance of LLMs [21]. The process involves presenting the model with K demonstration examples, containing the context and ground truth answers, followed by a final example (query), where the model is expected to generate the answer following the template of the given demonstrations. The key advantage of this approach is that it requires no updates to the model's weights. An example of such a prompt is shown in Fig. 2. While the demonstrations could be selected randomly, Zhou et al. [22] have shown that this approach can lead to decreased model performance compared to the zero-shot scenario. To address this, following the method in [23], we adopt a selection strategy where demonstrations are chosen based on the similarity of the query image x to a bank of precomputed image features from training set. Specifically, given the feature vector $\mathcal{V}(x)$ of the query image, we select the N most similar image feature vectors from the training set $\mathcal{D}$. Let S denote the set of indices corresponding to the N most similar feature vectors:

$$S = \text{argsort}(\{\text{sim}(\mathcal{V}(x), f_i)\}_{i=0}^n)[:, N], \quad (2)$$

where $\text{sim}(\cdot)$ represents a similarity metric (e.g., cosine similarity), f_i the image feature from the training set, and $\mathcal{V}(\cdot)$ is an arbitrary visual encoder. Once the indices S of the similar images are identified, we retrieve their corresponding ground truth concepts and class labels from the training set. These are used to create demonstration examples, which are then appended to the prompt, as illustrated in Fig. 2.

4. Experiments

In this section, we outline our experimental setup, including details of the adopted models, datasets, and evaluation metrics, as well as the implementation details.

Table 1

Dataset statistics. Numbers between rounded brackets represent the # of Melanoma examples in the split. MEL is Melanoma and NEV is Nevus.

Dataset	Classes	Train size	Validation size	Test size	% (MEL/NEV)
PH ² [27]	2	160 (28 to 34)	-	40 (6 to 12)	25/75
Derm7pt [28]	2	346 (90)	161 (61)	320 (101)	30/70
HAM10000 [29]	2	2,646 (752)	599 (172)	200 (100)	30/70

4.1. Setup

Models. In our experiments, we evaluate the performance of our approach against the standard CBM [4] and several pretrained Vision-Language Models (VLMs), including the generic CLIP [13], the medical-domain BiomedCLIP [14], and three dermatology-specific VLMs: MONET [18], CBI-VLM [17], and ExplICD [19]. For the Large Language Models (LLMs), we compare the generic Mistral-7B-Instruct [24], the medical-domain MMed-LLama3-8B-EnIns [25], and conduct additional experiments with GPT-4 [26] in the top-performing setting.

Datasets. We conduct experiments using three publicly available dermoscopic datasets: PH² [27], Derm7pt [28], and HAM10000 [29]. The choice of these datasets was determined by the availability of clinical concept annotations. The PH² dataset consists of dermoscopic images of melanocytic lesions, including “melanoma” and two types of “nevus”, which we merged and treated as a single “nevus” category. For this dataset, we applied 5-fold cross-validation due to its small size. Derm7pt contains over 2,000 clinical and dermoscopic images, which we filtered to include only images of “nevus” and “melanoma” classes. HAM10000 contains 10,015 dermoscopic images of various skin lesions across different body locations. Following [30], we used all biopsy-verified “melanoma” and “nevus” images from HAM10000, resulting in a set of $n = 3,445$ images. Detailed statistics, including the train/validation/test splits for each dataset, are shown in Table 1. All reported results are based on the entire test set for each dataset. For Derm7pt, we followed the official split partition, and for HAM10000, we used the partitioning described in [30].

Evaluation metrics.

We adopt Balanced Accuracy (BAcc) as the arithmetic mean of Sensitivity (TPR) and Specificity (TNR), as this metric is tailored to deal with imbalanced datasets, which is the case of the considered datasets (see % (MEL/NEV) column of Table 1).

4.2. Implementation details

For experiments with concept-based models, we reproduce the results using the official implementation of CBM [4]. For pretrained VLMs, we use the checkpoints from the HuggingFace library for CLIP, BiomedCLIP and MONET. Additionally, we employ the official implementations of CBI-VLM and ExplICD, as provided in [17] and [19], respectively. Checkpoints for Mistral-7B-Instruct and MMed-LLama3-8B-EnIns are also sourced from the HuggingFace library.

For the concept prediction task, we consider the concept ontology of MONET as target set of concepts for experiments with CLIP, BiomedCLIP, MONET and CBI-VLM. For ExplICD, we utilize its own concept ontology, as the model was fine-tuned with a pre-defined set of concepts, either gathered by LLMs or medical experts. A summary of the ontologies used for each model is provided in Table 2. It is important to note that we only rely on annotated ground-truth concepts for each dataset when performing test-time interventions.

For the selection strategy used to perform ICL, the image similarity is determined using the cosine similarity between image features, which are extracted from the visual encoder of ExplICD. Further details, documentation and a demo of our approach are available at <https://github.com/CristianoPatricio/2-step-concept-based-skin-diagnosis>. All experiments are run on a single A40 GPU with 48 GB of VRAM.

5. Results and analysis

This section presents the key results of our two-step approach for skin lesion classification across three dermoscopic datasets. Specifically, we compare the performance of our approach (**VLM + LLM**) against two baseline approaches:

- **CBM:** We use the original CBM, which operates in two stages: i) first, a CNN is used to predict concepts from the input image ($x \rightarrow c$), and ii) a linear model then uses these predicted concepts to infer the final target ($c \rightarrow y$). This approach requires concept annotations and training in both stages.
- **VLM + Linear Classifier:** We use a pretrained VLM to predict the concepts from input images (the “bottleneck” part), followed by training a linear classifier on the predicted concepts to determine the class label. Training is only needed for the $c \rightarrow y$ phase.

Our two-step approach, **VLM + LLM**, leverages a pretrained VLM to predict concepts from an input image (Step 1) and prompts an off-the-shelf LLM to generate the final diagnosis (Step 2) based on the predicted concepts.

5.1. Disease classification performance

Table 3 presents the disease classification performance of models in terms of Balanced Accuracy (BAcc), Sensitivity, and Specificity across different datasets in a zero-shot setting.

Zero-shot classification performance. Using our two-step methodology (VLM + LLM), CLIP-like models such as CLIP and BiomedCLIP generally exhibit poor zero-shot performance, especially in terms of sensitivity, as they frequently predict the benign class (nevus). In contrast, Vision-Language Models (VLMs) fine-tuned on dermoscopic images perform significantly better, with MONET and ExplICD showing the highest accuracy. Although CBI-VLM was also fine-tuned on dermoscopic data, it underperformed compared to MONET and ExplICD. This may be due to its reliance on image-text alignment only in the final layers of CLIP, which were trained on generic data, while keeping the remaining weights frozen. When combining ExplICD with MMed for PH2 and HAM10000, and ExplICD with Mistral for Derm7pt, we increase BAcc over traditional CBM: 8% on PH2, 1.5% on Derm7pt, and 5% on HAM10000. These gains can be attributed to ExplICD’s use of a pretrained vision-language model and its incorporation of textual diagnostic criteria as knowledge anchors. These anchors are aligned with predefined visual concept tokens through contrastive loss, resulting in more fine-grained alignment than traditional image-text alignment in VLMs. We also evaluate the combination of ExplICD + GPT-4 due to the strong performance of ExplICD with MMed and Mistral. However, no substantial improvements were observed compared to using open-source LLMs like MMed and Mistral.

Few-shot classification performance. Fig. 3 shows that the results for PH2 and Derm7pt can be further improved by incorporating additional context into the prompt through the inclusion of n demonstration examples, selected using our retrieval strategy.

As shown in Fig. 3, for PH2, we achieve a score of 83.56% and 85.05% when using Mistral and MMed, respectively, with only one

Table 2

Concept ontology used in each model and dataset. Model-specific concepts were collected by either medical experts or LLM and therefore automatically annotated. Dataset-specific concepts are the ground-truth annotations.

Model-specific	ExpLICD
	Color: highly variable, often with multiple colors (black, brown, red, white, blue), uniformly tan, brown, or black
	Shape: irregular, round
	Border: often blurry and irregular; sharp and well-defined
	Dermoscopic Patterns: atypical pigment network, irregular streaks, blue-whitish veil, irregular, regular pigment network, symmetric dots and globules
Dataset-specific	Texture: a raised or ulcerated surface; smooth
	Symmetry: asymmetrical; symmetrical
	Elevation: flat to raised
	MONET
	Asymmetry, Irregular, Black, Blue, White, Brown, Erosion, Multiple Colors, Tiny, Regular
Dataset-specific	HAM10000
	thick reticular or branched lines; black dots or globules in the periphery of the lesion; white lines or white structureless area; eccentrically located structureless area; grey patterns; polymorphous vessels; pseudopods or radial lines at the lesion margin that do not occupy the entire lesional circumference; asymmetric combination of multiple patterns or colors in the absence of other melanoma criteria; melanoma simulator
	PH² and Derm7pt
	typical pigment network, atypical pigment network, irregular streaks, regular streaks, regular dots and globules, irregular dots and globules, blue-whitish veil, regression structures

Table 3

Disease classification performance across various datasets. Reported results in %. The best results are highlighted in **bold**, and the second-best results are underlined.

Model	Trainable	Dataset									
		PH ²			Derm7pt			HAM10000			
		Sens.	Spec.	BAcc	Sens.	Spec.	BAcc	Sens.	Spec.	BAcc	
Image to Concepts to Label Classification [x → c → y]											
	CBM [4]	✓	42.04	98.06	70.14	69.31	89.50	77.56	51.00	94.00	70.83
VLM + Linear Classifier	CLIP [13] + Linear Classifier	✓	59.29	96.32	77.80	37.62	91.32	64.47	21.00	94.00	57.50
	BiomedCLIP [14] + Linear Classifier	✓	22.78	98.18	60.49	6.93	99.09	53.01	0.00	100.0	50.00
	MONET [18] + Linear Classifier	✓	71.43	96.79	<u>84.11</u>	53.47	92.69	73.08	20.00	92.00	56.00
	CBI-VLM [17] + Linear Classifier	✓	24.60	98.22	61.41	50.50	89.95	70.22	56.00	92.00	74.00
	ExpLICD [19] + Linear Classifier	✓	58.49	96.47	77.48	55.45	95.43	75.44	60.00	87.00	73.50
Ours (VLM + LLM)	CLIP [13]										
	+ MMed-Llama3-8B-EnIns	✗	88.17	3.71	45.94	59.41	14.61	37.01	87.00	2.00	44.50
	+ Mistral-7B-Instruct	✗	8.33	96.90	52.62	19.80	88.58	54.19	11.00	94.00	55.00
	BiomedCLIP [14]										
	+ MMed-Llama3-8B-EnIns	✗	21.75	87.38	54.56	21.78	70.78	46.28	21.00	70.00	45.50
	+ Mistral-7B-Instruct	✗	5.00	93.59	49.29	9.90	79.00	44.45	10.00	78.00	44.00
	MONET [18]										
	+ MMed-Llama3-8B-EnIns	✗	100.0	28.67	64.33	89.11	31.96	60.54	77.00	45.00	61.00
	+ Mistral-7B-Instruct	✗	73.17	61.30	67.24	73.27	62.56	67.91	40.00	67.00	53.50
	CBI-VLM [17]										
	+ MMed-Llama3-8B-EnIns	✗	6.11	96.27	51.19	6.93	97.26	52.10	77.00	11.00	44.00
	+ Mistral-7B-Instruct	✗	86.67	0.59	43.63	99.01	11.87	55.44	100.0	0.00	50.00
	ExpLICD [19]										
	+ MMed-Llama3-8B-EnIns	✗	58.49	97.65	78.07	65.35	91.78	78.56	66.00	86.00	76.00
	+ Mistral-7B-Instruct	✗	72.06	82.82	77.44	67.33	90.87	<u>79.10</u>	72.00	79.00	<u>75.50</u>
	+ GPT-4	✗	72.78	93.96	83.37	67.33	90.87	<u>79.10</u>	77.00	72.00	<u>74.50</u>
★ (ExpLICD + LLM w/ few-shot)	✗	79.44	90.65	85.05	67.33	92.24	79.78	76.00	74.00	75.00	




Fig. 3. Few-shot disease classification performance (in BACC %) across different n -shot settings. Each bar corresponds to an n -shot scenario ($n = 0, 1, 2, 4, 8$).

Table 4

Prompts used for zero-shot image classification. The {class} is replaced by “Nevus” or “Melanoma” in our experiments.

Model	Template
BiomedCLIP	This is a dermoscopic image of {class}
CLIP	This is a dermoscopic image of {class}
MONET	This is skin image of {class}
CBI-VLM	This is dermatoscopy of {class}
ExpLICD	This is a dermoscopic image of {class}

demonstration example, significantly outperforming the zero-shot performance. For Derm7pt, the zero-shot performance improves with $K = 4$ demonstrations in Mistral and with $K = 8$ demonstration examples in MMed. Although no BAcc improvement was observed for HAM10000 using ICL, we achieved increased sensitivity with two demonstration examples in MMed and one in Mistral.

5.2. Ablation study on disease classification

To further validate the effectiveness of our method, we conduct ablation studies on the disease classification strategy. Specifically, we explore direct label prediction from the input image ($x \rightarrow y$) by discarding the intermediate concept prediction step. Instead, we compare image features to short text prompt describing the class, detailed in Table 4.

As shown in Table 5, directly comparing image features with a text class prompt yields lower performance than our proposed method (last row of Table 5), which first predict the clinical concepts and then uses that information for class label prediction. Nonetheless, CBI-VLM achieves the highest performance across most datasets, benefiting from fine-tuning that aligns image and label features within the same class. These results emphasize the value of an intermediate concept prediction step, enhancing both performance and interpretability.

5.3. Test-time intervention

Correcting mispredicted concepts during inference time is especially valuable in high-stakes fields like medicine. By addressing incorrect concepts, clinicians can help the LLM correcting its final diagnosis. To simulate a scenario where a clinician intervenes to correct erroneous concepts, we replace all predicted concepts with their ground-truth values and evaluate the LLM’s classification performance. We conducted experiments on all datasets using the top-performing CLIP-like VLM (MONET) for comparison.

As presented in Table 6, intervening on the predicted concepts leads to a significant improvement in Balanced Accuracy across all datasets, with MMed demonstrating the greatest average improvement. These findings underscore the critical role of robust concept detection and emphasize the value of human supervision when applying such methods in clinical practice.

5.4. Diagnostic interpretability

An appealing ability of the proposed methodology is its ability to provide clinical concept-based explanations alongside a diagnosis grounded in those concepts. Fig. 4 presents two examples of correctly diagnosed skin images along with their concept-based explanations, as well as one example with an incorrect diagnosis. The concept ontology defined in ExpLICD, which align with the ABCDEs of melanoma [31], allows for further examination of the skin image, especially when a predicted concept appears inconsistent with the diagnosis (e.g. the symmetry concept in the bottom row of Fig. 4). This two-step approach ensures both interpretability and transparency in the model’s decision-making.

Fig. 5 presents the most frequently predicted dermoscopic concepts for each target class in Derm7pt dataset. For Nevus, the predominant predictions include uniform color, well-defined borders, regular

pigment network, smooth texture, and symmetry, which are key features typically linked to benign lesions like Nevus. In contrast, for the melanoma class, concepts such as multiple colors, irregular borders, atypical pigment network, ulcerated surface, and asymmetry are the most frequently predicted, and commonly associated with melanoma diagnosis. This analysis offers a transparent and interpretable framework that can be easily understood and validated by medical experts.

6. General discussion

The effect of few-shot prompting in the overall performance. Our proposed two-step approach consistently outperforms all baseline methods across datasets when using ICL, demonstrating the utility of incorporating additional context into prompts with a few demonstration examples to facilitate LLM reasoning. This improvement was particularly noticeable in the PH2 and Derm7pt datasets. However, no performance gains were observed for HAM10000, which can be attributed to two key factors:

- **Incorrect prediction of concepts:** Our empirical analysis shows that predicted concepts have a significant impact on LLM responses. While demonstration examples can guide the LLM to provide correct diagnoses, situations may arise where incorrectly predicted concepts, such as predicting features commonly associated with melanomas when the ground truth label is nevus, lead the LLM to produce incorrect disease diagnoses, as its response is grounded in the predicted concepts.
- **Inaccuracy in demonstration examples:** Another observed issue is when demonstration examples contain inaccurate associations between concepts and ground truth labels. Since the concepts are automatically predicted by the VLM, there is no guarantee that demonstration examples from the training set accurately represent the correct associations. This mismatch could mislead the LLM into providing incorrect diagnoses.

Overall, these issues were frequently observed in HAM10000 dataset, suggesting that it contains more challenging images, which may contribute to a higher False Positive rate. Further refinement of concept prediction and selection of high-quality demonstration examples could mitigate these limitations.

Transparency and interpretability in model responses. Overall, the results presented in Section 5 highlight the effectiveness of our two-step approach in enabling pretrained VLMs to predict clinical concepts from images and use those concepts to ground the final diagnosis. This enhances both interpretability and transparency in the model’s responses, all without requiring additional training and using only a few annotated examples. Notably, our approach outperforms traditional CBMs and state-of-the-art XAI methods across all datasets.

Human interaction through test-time intervention. Our two-step methodology allows for human intervention in the clinical concepts predicted by the VLM. Since the concepts are presented in plain text to the user, this offers a significant advantage over traditional CBMs, which require adjusting numerical values in the bottleneck layer. Furthermore, our results in Section 5.3 underscore the importance of this test-time intervention in enhancing overall performance.

Overall advantages of the proposed two-step approach. When comparing our work with studies [9–12] using similar techniques, our two-step approach stands out for multiple reasons:

- **Training-free methodology:** We use pretrained models with medical knowledge that require only two inference steps (taking less than 3 seconds) to generate an interpretable final diagnosis. In contrast, existing CBMs and supervised methods require extensive training and are annotation-dependent;

Table 5

Ablation study on zero-shot classification performance. The best results are highlighted in **bold**, and the second-best results are underlined.

Model	Dataset								
	PH ²			Derm7pt			HAM10000		
	Sens.	Spec.	BAcc	Sens.	Spec.	BAcc	Sens.	Spec.	BAcc
Image to Label [x → y]									
CLIP [13]	100.0	0.00	50.00	93.07	7.31	50.19	91.00	4.00	47.50
BiomedCLIP [14]	18.89	94.11	56.50	69.31	62.56	65.93	45.00	68.00	56.50
MONET [18]	78.25	88.48	<u>83.36</u>	55.45	88.13	71.79	86.00	44.00	65.00
CBI-VLM [17]	40.40	74.92	77.46	82.18	73.06	<u>77.62</u>	32.00	99.00	<u>65.50</u>
ExpLICD [19]	38.02	58.84	48.43	46.53	53.88	50.21	54.00	52.00	53.00
★ (ExpLICD + LLM w/ few-shot)	79.44	90.65	85.05	67.33	92.24	79.78	76.00	74.00	75.00

Table 6

Test-time intervention results. BAcc is Balanced Accuracy. Reported results in %. The best results are highlighted in **bold**.

Model	Intervention	Dataset		
		PH ²	Derm7pt	HAM10000
		BAcc	BAcc	BAcc
Concepts to Label [c → y]				
MMed-Llama3-8B-EnIns [25]	✗	64.33	60.54	61.00
MMed-Llama3-8B-EnIns [25]	✓	75.27 (+10.94)	81.09 (+20.55)	80.00 (+19.00)
Mistral-7B-Instruct	✗	67.24	67.91	53.50
Mistral-7B-Instruct	✓	73.43 (+6.19)	72.56 (+4.65)	58.00 (+4.50)


Image	Disease Diagnosis	Concept-based Explanation
	Actual Melanoma Predicted Melanoma ✓	The color is highly variable, often with multiple colors (black, brown, red, white, blue), the shape is irregular, the border is often blurry and irregular, the dermoscopic patterns are atypical pigment network, irregular streaks, blue-whitish veil, irregular, the texture is a raised or ulcerated surface, the symmetry is asymmetrical, the elevation is flat to raised. Thus the diagnosis is melanoma.
	Actual Nevus Predicted Nevus ✓	The color is uniformly tan, brown, or black, the shape is round to irregular, the border is sharp and well-defined, the dermoscopic patterns are regular pigment network, symmetric dots and globules, the texture is smooth, the symmetry is symmetrical, the elevation is flat to raised. Thus the diagnosis is nevus.
	Actual Melanoma Predicted Nevus ✗	The color is uniformly tan, brown, or black, the shape is round to irregular, the border is sharp and well-defined, the dermoscopic patterns are atypical pigment network, irregular streaks, blue-whitish veil, irregular, the texture is smooth, the symmetry is symmetrical, the elevation is flat to raised. Thus the diagnosis is Nevus.

Fig. 4. Examples of skin images diagnosed using our approach. The concept-based explanations alongside the predicted disease label allow for further inspection of the skin image, especially when a concept prediction appears inconsistent with the diagnosis.

- **Intuitive test-time intervention:** We enable intuitive human intervention to correct any predicted concept by editing plain text;
- **Flexible incorporation of concepts and class labels:** Additional concepts or class labels can be included without the need to re-train the entire architecture. In contrast, CBMs and supervised task-specific models are limited to a predefined set of class labels.

7. Limitations

Although AI-based methods have shown promising results in diagnosing skin lesions, their real-world effectiveness in clinical settings remains largely unexplored. Many studies focus on controlled datasets and scenarios that lack diversity in both diseases and skin tones [32]. Therefore, there is a need for the acquisition, annotation, and curation

of diverse data, alongside close collaboration between clinicians and AI researchers, to enhance transparency, explainability, and the accuracy of AI integration into clinical practice [33].

8. Conclusions and future work

We present a novel two-step methodology that utilizes off-the-shelf Vision-Language Models (VLMs) to predict clinical concepts and a Large Language Model (LLM) to generate disease diagnoses based on those concepts. This approach enhances transparency and explainability in decision-making, which is crucial for high-stakes applications like medical use cases. Our evaluation on three skin lesion datasets demonstrates that our method outperforms traditional Concept Bottleneck Models (CBMs), achieving a relative improvement of 21% on PH2, 3% on




Fig. 5. Predicted dermoscopic concepts for each target class in Derm7pt. The width of the lines indicates the frequency with which each concept is predicted for its respective class. This visualization highlights the key features associated with Nevus and melanoma, providing insights into the most relevant characteristics for each diagnosis.

Derm7pt, and 6% on HAM10000. Additionally, it surpasses state-of-the-art explainable approaches without requiring additional training and uses only a few demonstration examples, thereby addressing the two primary limitations of CBMs. Furthermore, our approach easily incorporates new concepts without retraining and supports test-time interventions. While our methodology was evaluated only on skin image datasets, it is generalizable and can be applied to other domains. In future work, we plan to explore the integration of visual explanations to further enhance trust and promote the adoption of automated diagnosis systems in clinical settings.

CRedit authorship contribution statement

Cristiano Patrício: Writing – original draft, Methodology, Investigation, Conceptualization, Data curation, Software, Validation. **Luís F. Teixeira:** Writing – review & editing, Supervision, Conceptualization, Methodology. **João C. Neves:** Writing – review & editing, Supervision, Conceptualization, Methodology.

Declaration of competing interest

We declare that the authors have no conflicts of interest related to this work.

Acknowledgements

This work was funded by the Portuguese Foundation for Science and Technology (FCT) under the PhD grant “2022.11566.BD”, and supported by NOVA LINGS (UIDB/04516/2020) with the financial support of FCT/IP.

References

- [1] Codella NC, Nguyen Q-B, Pankanti S, Gutman DA, Helba B, Halpern AC, et al. Deep learning ensembles for melanoma recognition in dermoscopy images. *IBM J Res Dev* 2017;61(4/5): 5–1.
- [2] Esteve A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 2017;542(7639):115–8.
- [3] Rotemberg V, Halpern A, Dusza S, Codella NC. The role of public challenges and data sets towards algorithm development, trust, and use in clinical practice. In: *Seminars in cutaneous medicine and surgery*, vol. 38. 2019. p. E38–42.
- [4] Koh PW, Nguyen T, Tang YS, Musmann S, Pierson E, Kim B, et al. Concept bottleneck models. In: *Proceedings of the international conference on machine learning (ICML)*; 2020. p. 5338–48.
- [5] Fang Z, Kuang K, Lin Y, Wu F, Yao Y-F. Concept-based explanation for fine-grained images and its application in infectious keratitis classification. In: *Proceedings of the ACM international conference on multimedia*; 2020. p. 700–8.
- [6] Lucieri A, Bajwa MN, Braun SA, Malik MI, Dengel A, Ahmed S. On interpretability of deep learning based skin lesion classifiers using concept activation vectors. In: *Proceedings of the international joint conference on neural networks (IJCNN)*; 2020. p. 1–10.
- [7] Patrício C, Neves JC, Teixeira LF. Coherent concept-based explanations in medical image and its application to skin lesion diagnosis. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops (CVPRW)*; 2023. p. 3799–808.
- [8] Ramaswamy VV, Kim SS, Fong R, Russakovsky O. Overlooked factors in concept-based explanations: dataset choice, concept learnability, and human capability. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*; 2023. p. 10932–41.
- [9] Yang Y, Panagopoulou A, Zhou S, Jin D, Callison-Burch C, Yatskar M. Language in a bottle: language model guided concept bottlenecks for interpretable image classification. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*; 2023. p. 19187–97.
- [10] Oikarinen T, Das S, Nguyen LM, Weng T-W. Label-free concept bottleneck models. Available from: [arXiv:2304.06129](https://arxiv.org/abs/2304.06129), 2023.
- [11] Menon S, Vondrick C. Visual classification via description from large language models. Available from: [arXiv:2210.07183](https://arxiv.org/abs/2210.07183), 2022.
- [12] Yan A, Wang Y, Zhong Y, He Z, Karypis P, Wang Z, et al. Robust and interpretable medical image classifiers via concept bottleneck models. Available from: [arXiv:2310.03182](https://arxiv.org/abs/2310.03182), 2023.
- [13] Radford A, Kim JW, Hallacy C, Ramesh A, Goh G, Agarwal S, et al. Learning transferable visual models from natural language supervision. In: *ICML*; 2021. p. 8748–63.
- [14] Zhang S, Xu Y, Usuyama N, Xu H, Bagga J, Tinn R, et al. BiomedCLIP: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. Available from: [arXiv:2303.00915](https://arxiv.org/abs/2303.00915), 2023.
- [15] Eslami S, Meinel C, De Melo G. PubMedCLIP: how much does CLIP benefit visual question answering in the medical domain? In: *Findings of the association for computational linguistics: EACL* 2023; 2023. p. 1151–63.
- [16] Wang Z, Wu Z, Agarwal D, Sun J. MedCLIP: contrastive learning from unpaired medical images and text. In: *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*; 2022. p. 3876–87.
- [17] Patrício C, Teixeira LF, Neves JC. Towards concept-based interpretability of skin lesion diagnosis using vision-language models. In: *2024 IEEE international symposium on biomedical imaging (ISBI)*. IEEE; 2024. p. 1–5.
- [18] Kim C, Gadgil SU, DeGrave AJ, Omiye JA, Cai ZR, Daneshjou R, et al. Transparent medical image ai via an image-text foundation model grounded in medical literature. *Nat Med* 2024. <https://doi.org/10.1038/s41591-024-02887-x>.
- [19] Gao Y, Gu D, Zhou M, Metaxas D. Aligning human knowledge with visual concepts towards explainable medical image classification. Available from: [arXiv:2406.05596](https://arxiv.org/abs/2406.05596), 2024.
- [20] Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I, et al. Language models are unsupervised multitask learners. *OpenAI blog* 2019;1(8):9.
- [21] Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. In: *Advances in neural information processing systems*, vol. 33. 2020. p. 1877–901.
- [22] Zhou G, Han Z, Chen S, Huang B, Zhu L, Khan S, et al. Adapting large multimodal models to distribution shifts: the role of in-context learning. [arXiv:2405.12217](https://arxiv.org/abs/2405.12217), 2024.
- [23] Yang Z, Gan Z, Wang J, Hu X, Lu Y, Liu Z, et al. An empirical study of GPT-3 for few-shot knowledge-based VQA. In: *Proceedings of the AAAI conference on artificial intelligence*, vol. 36. 2022. p. 3081–9.

- [24] Jiang AQ, Sablayrolles A, Mensch A, Bamford C, Chaplot DS, Casas Ddl, et al. Mistral 7b. Available from: [arXiv:2310.06825](https://arxiv.org/abs/2310.06825), 2023.
- [25] Qiu P, Wu C, Zhang X, Lin W, Wang H, Zhang Y, et al. Towards building multilingual language model for medicine. Available from: [arXiv:2402.13963](https://arxiv.org/abs/2402.13963), 2024.
- [26] OpenAI, Gpt-4 technical report, Available from: [arXiv:2303.08774](https://arxiv.org/abs/2303.08774), 2023. <https://arxiv.org/abs/2303.08774>.
- [27] Mendonça T, Ferreira PM, Marques JS, Marcal ARS, Rozeira J. PH2 - a dermoscopic image database for research and benchmarking. In: Proceedings of the international conference of the IEEE engineering in medicine and biology society (EMBC); 2013. p. 5437–40.
- [28] Kawahara J, Daneshvar S, Argenziano G, Hamarneh G. Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *IEEE J Biomed Health Inform* 2019;23(2):538–46.
- [29] Tschandl P, Rosendahl C, Kittler H. The ham10000 dataset, a large collection of multi-source dermoscopic images of common pigmented skin lesions. *Sci Data* 2018;5(1):1–9.
- [30] Chanda T, Hauser K, Hobelsberger S, Bucher T-C, Garcia CN, Wies C, et al. Dermatologist-like explainable ai enhances trust and confidence in diagnosing melanoma. *Nat Commun* 2024;15(1):524.
- [31] Rigel DS, Friedman RJ, Kopf AW, Polsky D. ABCDE—an evolving concept in the early detection of melanoma. *Arch Dermatol* 2005;141(8):1032–4.
- [32] Hsiao Y-P. The rise of artificial intelligence in dermatology: perspectives and potentials. *Dermatolog Sin* 2024;42(2):77–9.
- [33] Phung M, Muralidharan V, Rotemberg V, Novoa RA, Chiou AS, Sadée CY, et al. Best practices for clinical skin image acquisition in translational artificial intelligence research. *J Invest Dermatol* 2023;143(7):1127–32.