



CBVLM: Training-free explainable concept-based Large Vision Language Models for medical image classification

Cristiano Patrício ^{a,c,d,*}, Isabel Rio-Torto ^{b,c}, Jaime S. Cardoso ^c, Luís F. Teixeira ^c, João C. Neves ^{a,d}

^a Universidade da Beira Interior, Covilhã, Portugal

^b Departamento de Ciência de Computadores, Faculdade de Ciências, Universidade do Porto, Rua Campo Alegre s/n, 4169-007 Porto, Portugal

^c INESC TEC, Faculdade de Engenharia, Universidade do Porto, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal

^d NOVA LINCS, Portugal

ARTICLE INFO

Keywords:

Large Vision-Language Models
Multimodal Large Language Models
In-context learning
Medical imaging
Explainability
Concept-based explanations
Concept Bottleneck Models

ABSTRACT

The main challenges limiting the adoption of deep learning-based solutions in medical workflows are the availability of annotated data and the lack of interpretability of such systems. Concept Bottleneck Models (CBMs) tackle the latter by constraining the model output on a set of predefined and human-interpretable concepts. However, the increased interpretability achieved through these concept-based explanations implies a higher annotation burden. Moreover, if a new concept needs to be added, the whole system needs to be retrained. Inspired by the remarkable performance shown by Large Vision-Language Models (LVLMs) in few-shot settings, we propose a simple, yet effective, methodology, *CBVLM*, which tackles both of the aforementioned challenges. First, for each concept, we prompt the LVLm to answer if the concept is present in the input image. Then, we ask the LVLm to classify the image based on the previous concept predictions. Moreover, in both stages, we incorporate a retrieval module responsible for selecting the best examples for in-context learning. By grounding the final diagnosis on the predicted concepts, we ensure explainability, and by leveraging the few-shot capabilities of LVLMs, we drastically lower the annotation cost. We validate our approach with extensive experiments across four medical datasets and twelve LVLMs (both generic and medical) and show that *CBVLM* consistently outperforms CBMs and task-specific supervised methods without requiring any training and using just a few annotated examples. More information on our project page: <https://cristianopatricio.github.io/CBVLM/>.

1. Introduction

The field of medical imaging analysis has witnessed significant advancements through the use of supervised deep learning methods. However, their performance greatly depends on the availability and quality of training samples, which are particularly costly to obtain, as clinician expertise is required. The recent rise of Large Vision-Language Models (LVLMs) has also made its way into the medical field, of which Med-LVLMs such as [1–4] are prominent examples, where pre-trained LVLMs are fine-tuned using medical data. However, despite the improved performance, these Med-LVLMs require even more annotated data. Consequently, a different research direction has emerged, exploring the potential of applying LVLMs to specific downstream tasks with minimal or no supervision, by leveraging their few-shot capabilities [5–7], particularly the use of in-context learning (ICL) for

improving model performance through a short number of examples in the prompt. In addition to the need for annotated data, another major challenge with deep learning-based medical imaging is the lack of interpretability, which is of the utmost importance in high-stakes decision-making scenarios [8]. One way of introducing interpretability into such frameworks is via concept-based explanations. Concept Bottleneck Models (CBMs) provide this type of explanations by introducing a layer, the “bottleneck”, that predicts the presence or absence of a predefined set of human-interpretable concepts. The final classification prediction is then based on these concepts. Nevertheless, CBMs still require training and annotated data; in fact, the amount of annotated data actually increases since concept annotations are needed on top of the classification labels. We propose combining the best of both worlds: the interpretability of CBMs and the few-shot abilities of LVLMs to

* Corresponding author at: Universidade da Beira Interior, Covilhã, Portugal.

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

¹ Equal contribution.




Fig. 1. Comparison between the proposed methodology (*CBVLM*) with existing methods. Unlike traditional approaches, our methodology grounds the final diagnosis on a set of clinical concepts predicted by the LVLM itself (cf. *Concept Detector*), thus increasing classification performance, while being interpretable.

decrease the annotation cost. Our methodology, *CBVLM*, encompasses a two-stage approach, where, first, the LVLM is prompted to predict the existence/absence of a set of concepts in an image, and, in a second stage, the predicted concepts are incorporated into the prompt such that the LVLM provides a final diagnosis grounded on these predicted concepts. Fig. 1 represents this methodology compared to the common approach where an LVLM is simply prompted to perform classification. Moreover, contrary to CBMs, which require retraining the entire model to incorporate a new concept, *CBVLM* can easily incorporate additional concepts, requiring only two inference steps. We validate *CBVLM* on four medical datasets encompassing different imaging modalities and using twelve LVLMs, and show that *CBVLM* achieves comparable performance or even surpasses task-specific supervised methods, doing so without training, with few annotated examples, and without sacrificing interpretability. Although, individually, the techniques used in *CBVLM* (e.g., ICL) are not novel, this is, to the best of our knowledge, the first work combining them in a novel methodology exploring the capability of LVLMs to provide concept-based explanations for medical image classification.

2. Related work

Large vision-language models in the medical domain. The emergent capabilities of Large Language Models (LLMs) have led to the development of LLMs able to process visual inputs, i.e. LVLMs. These models can reason about images in a zero- or few-shot manner (e.g., LLaVA [9], GPT-4V [10]), allowing them to tackle downstream tasks without fine-tuning. General-purpose LVLMs have recently been adapted for the medical domain through training on extensive medical datasets, coining the term Med-LVLMs. The generalization capability of these models has prompted researchers to assess their performance in multiple medical data. Royer et al. [11] recently released an open-source toolkit for fair and reproducible evaluation of Med-LVLMs in tasks like classification and report generation. Han et al. [5] evaluated the performance of a general LVLM on the NEJM Image Challenge, as well as the impact of fine-tuning the LVLM. In [12], Med-LVLMs were evaluated in different specialties and compared with medical experts, with results comparable to the performance of junior doctors. In [7], the authors benchmarked the performance of GPT-4V against fully supervised image classifiers and found that their performance is comparable. In a different line of

research, Xia et al. [13] evaluated the trustworthiness and fairness of Med-LVLMs and found that models often display factual inaccuracies and can be unfair across different demographic groups. *CBVLM* distinguishes itself from previous studies as it is not limited to prompting the LVLM to provide a diagnosis, but it is, to the best of our knowledge, the first to generate concept-based explanations on which the LVLM response is grounded prior to diagnosis prediction, thus enhancing the performance and transparency of the decision-making process.

Concept bottleneck models. CBMs [14] are based on an encoder–decoder paradigm, where the encoder predicts a set of human-specified concepts from the input image, while the decoder leverages these predicted concepts to generate the final predictions. Due to their inherent interpretability, these models have attracted particular interest in medical image classification. Several CBMs have been proposed for a wide variety of medical tasks (e.g., melanoma diagnosis [15–17], knee osteoarthritis [14] and chest X-ray classification [18]). Nevertheless, CBMs have well-known limitations: (i) they require concept annotations (image-level or class-level), (ii) they tend to have lower performance compared to traditional (black-box) supervised methods, (iii) they struggle to predict concepts accurately, and (iv) incorporating new concepts implies retraining the whole CBM. These challenges are particularly significant in medical imaging, where data availability is typically limited due to the expertise required for annotations. Consequently, LVLMs represent a promising alternative to traditional CBMs, as they are able to generalize to unseen domains with minimal or no supervision. Moreover, incorporating additional concepts becomes trivial and does not require fine-tuning, as happened with traditional CBMs.

3. Methodology

Given the task of predicting a target $y \in \mathbb{R}$ from input $x \in \mathbb{R}^d$, let $D = \{(x^{(i)}, y^{(i)}, c^{(i)})\}_{i=1}^n$ be a batch of training samples where $c \in \mathbb{R}^{\{0,1\}^l}$ is a vector of l concepts. Traditional (black-box) models learn a function $h : \mathbb{R}^d \rightarrow \mathbb{R}$ that directly predicts the target $y \in \mathbb{R}$ from the input $x \in \mathbb{R}^d$. However, CBMs first map the input x into a set of interpretable concepts c (“bottleneck”) by learning $g : \mathbb{R}^d \rightarrow \mathbb{R}^l$, and use these concepts to predict the target y using $f : \mathbb{R}^l \rightarrow \mathbb{R}$. Therefore, the prediction $\hat{y} = f(g(x))$ is entirely based on the predicted concepts $\hat{c} = g(x)$.




Fig. 2. Overview of CBVLM. Our methodology is organized into two key stages: (1) The *Concept Detection* stage, where the LVLM predicts the individual presence of each predefined clinical concept in the query image. This is achieved using a custom prompt that supports both zero- and few-shot settings. In the latter, we include a set of demonstration examples (middle block of *Prompt Construction*) chosen by the *Retrieval Module*, responsible for selecting the N most similar examples to the input image. To evaluate the LVLM answer, we employ an *Evaluation Block* which first tries to extract the desired LVLM response using a rule-based formulation. If this fails, we adopt an auxiliary LLM to extract the desired response. (2) In the *Disease Diagnosis* stage, the final diagnosis is generated by the LVLM based on the clinical concepts predicted in the first stage, which are directly incorporated in the query (highlighted in yellow). This approach ensures that the diagnosis is grounded on the identified clinical concepts, enhancing the interpretability and transparency of the LVLM's response. In this second stage, the *Retrieval Module* is also used to select the N most similar demonstrations.

In order to generate concept-based explanations with LVLMs, we follow the CBM paradigm and introduce a two-stage methodology, as outlined in Fig. 2 and detailed in the following subsections.

3.1. Concept detection

The first stage involves predicting the l concepts present in a given input image. For each predefined clinical concept, we start by providing the LVLM with a short description of that concept and then prompt it to answer if that concept is shown in the image (see Fig. 2 on the left). Thus, concept prediction is given by $p_{LVLM}(t_{ans}|t_I^1, \dots, t_I^m, t_P^1, \dots, t_P^k)$, where t_{ans} corresponds to the first generated token, $t_I^1, \dots, t_I^m$ correspond to the input image tokens, and $t_P^1, \dots, t_P^k$ correspond to the tokens of the prompt (description of the concept and the question). The structure of the (zero-shot) prompt for the prediction of a single concept is as follows:

```
{A brief description of the {concept}}.
In the <image>, the {concept} is:
A) present
B) absent
Choose one option. Do not provide additional information. Answer:
```

The descriptions for each clinical concept, as presented in the *Instruction* section shown in Fig. 2, were generated by ChatGPT [19] with the following prompt: "According to published literature in {dermatology, radiology, retinal lesions}, which phrases best describe a {skin, X-ray, eye fundus} image containing {concept}?"

Examples of prompts for different concepts related to the classification of diabetic retinopathy are shown in Fig. 2 on the left. The full prompts for all datasets can be found in the Supplementary Material.

3.2. Disease diagnosis

The second stage involves prompting the LVLM to answer if a certain disease (i.e. class) is present in the image. The concepts predicted in the previous stage are appended to the question, such that when providing an answer, the LVLM not only considers the image and the question, but also takes into account the concepts. Thus, the disease diagnosis is given by $p_{LVLM}(t_{ans}|t_I^1, \dots, t_I^m, t_P^1, \dots, t_P^k, t_C^1, \dots, t_C^l)$,

where $t_C^1, \dots, t_C^l$ are the tokens representing the concepts. This approach ensures that the diagnosis is grounded on the identified clinical concepts, enhancing the interpretability and transparency of the model's response. The structure of the (zero-shot) prompt for diagnosis classification is as follows:

```
{A brief description of {concept 0}}.
(...)
{A brief description of {concept l - 1}}.
What is the diagnosis that is associated with the following concepts: {predicted
concepts}
Options:
A) class 0
B) class 1
(...)
Choose one option. Do not provide additional information. Answer:
```

An example of the prompt designed for the classification of diabetic retinopathy is shown in Fig. 2 on the right. The full prompts for all datasets can be found in the Supplementary Material.

3.3. Few-shot prompting

A key factor contributing to the appeal of LLMs, and consequently LVLMs, is their ability to perform ICL. ICL, or few-shot learning/prompting, is a prompt engineering technique in which the model is provided with examples, known as demonstrations, consisting of: i) questions, and (ii) their corresponding answers. It has been shown to significantly improve the performance of both LLMs [20] and LVLMs [21]. Therefore, our methodology incorporates ICL in both stages, as shown in Fig. 2.

3.4. Demonstration example retrieval

Although the demonstrations for ICL could be chosen randomly, Zhou et al. [22] have shown that this selection policy can decrease the model performance when compared to the 0-shot scenario (when no examples are included in the prompt). Thus, following other works [21, 23–25], we introduce a Retrieval Module to select the demonstrations based on their similarity to the query image x_0 , a method known as Retrieval-based In-Context Examples Selection (RICES). Given the feature vector $\phi(x_0)$ of the query image obtained from an arbitrary

vision encoder, we select the N most similar image feature vectors from the training set $\mathcal{D}$ as the few-shot demonstrations to include in the prompt. For each different type of prompt, the question is repeated for every demonstration. The ground-truth answer to each demonstration question is also provided; depending on the stage of our methodology, this can be the absence/presence of a given concept (1st stage), or the disease diagnosis, as well as the absence/presence of all concepts (2nd stage).

3.5. Answer extraction

As LVLMs produce their outputs in natural language (i.e. it is open-ended), it is necessary to process the generated text to compare it with the ground-truth concepts or the ground-truth class labels. To simplify this process and limit the space of possible responses, all questions are posed as multiple choice and the LVLM is instructed to answer by choosing one of the options and not providing additional information (see the example prompts in Fig. 2).

One way of extracting the LVLM answer would be to select the option with the higher first token log probability, i.e. if token B is the one with the highest probability, then this would mean that the LVLM had chosen option B. However, [26] have shown that there is a mismatch between the first token probabilities and the actual LVLM answer. Thus, we use a Regular Expression (RegEx) to directly find the pattern “option_letter”, either in lower or upper case. We choose to provide the options as “option_letter” + “)” instead of “option_letter” + “.”, as it would be less probable to find the pattern “option_letter)” in the middle of the LVLM response. If none of the predefined options is found in the LVLM answer (e.g., when the LVLM answers the question but in a more verbose fashion), we follow the approach of previous works [3] and employ an auxiliary LLM² to extract the information from the LVLM answer and match it to the predefined list of options. If both approaches fail (e.g., when the LVLM response is just a description of the input image without any actual answer to the question), we mark that instance as unknown and report separately the percentage of instances for which this occurred.

4. Experimental setup

In this section, we outline our experimental setup with information about the adopted LVLMs, datasets, and evaluation metrics, as well as the implementation details.

4.1. Datasets

We conduct experiments on four publicly available datasets, encompassing dermatology (**Derm7pt** [27] and **SkinCon** [28]), radiology (**CORDA** [29]) and eye fundus imaging (**DDR** [30]) modalities. The selection of these datasets is constrained by the availability of datasets with annotated clinical features, hereinafter referred to as “concepts”. Unless otherwise stated, we use the official train-val-test splits of the datasets.

For the **Derm7pt** [27] dataset, with a total of 827 images, we consider only dermoscopic images of the “nevus” and “melanoma” classes, jointly with the 5 annotated dermoscopic attributes, namely Absent/Typical/Atypical Pigment Network, Absent/Regular/Irregular Streaks, Absent/Regular/Irregular Dots and Globules, Blue-Whitish Veils, and Regression Structures. The **SkinCon** [28] dataset contains 3,230 images from the Fitzpatrick 17k dataset [31] annotated with 48 clinical concepts. Following previous works, we use the binary labels denoting skin malignancy for the target task and randomly split the dataset into train/val/test with the proportion of 70%, 15%, and 15%, respectively. Additionally, only 22 out of the 48 concepts are used,

Table 1

Detailed statistics for the datasets used.

Dataset	Image Type & Task	Classes	Concepts	Train	Test
Derm7pt [27]	Dermoscopic Melanoma classification	2	5	346	320
SkinCon [28]	Dermoscopic Malignancy classification	2	22	2582	554
CORDA [29]	X-ray Covid-19 classification	2	11	967	392
DDR [30]	Fundus DR classification	5	4	817	404

as the other 26 concepts appear in less than 50 images. From the **CORDA** [29] dataset, a COVID-19 diagnosis dataset, we use its 1601 chest X-rays. Inspired by the work of Barbano et al. [32], we use a pretrained model on the CheXpert [33] dataset to annotate the samples with a binary label to denote the presence or absence of 14 radiological observations. Of those 14 concepts, we exclude the “Fracture” and “Support Devices” concepts, as these are not directly related to the presence or absence of the target class. The “No Finding” concept is implicitly included in the remaining 11 concepts given that if all 11 concepts are 0 it represents a healthy case without any findings. The **DDR** [30] dataset comprises 13,673 fundus images divided into six Diabetic Retinopathy (DR) levels. In our experiments, we focus on the five-class classification task for DR grading, discarding the images belonging to the ungradable class and considering a subset of 757 images annotated with 4 types of lesions correlated with DR, which we use as concepts. We augment this subset with 743 “No DR” images, resulting in a dataset of 1500 samples.

Table 1 provides a detailed summary of the statistics for each of the considered datasets. All results (cf. Section 5) are reported on the entire test set of each dataset.

4.2. LVLMs

In our experiments, we assess the performance of open-source LVLMs that can process both text and multiple images. Specifically, we focus on models trained on general-domain data, such as Open-Flamingo [23], Idefics3 [34], VILA [35], LLaVA-OneVision [36], Qwen2-VL [37], MiniCPM-V [38], InternVL 2.5 [39] and mPLUG-Owl3 [40]. Additionally, we consider models specialized for the medical domain, including Med-Flamingo [1], LLaVA-Med [2], CheXagent [3] and SkinGPT-4 [4]. We use their corresponding checkpoints from the HuggingFace library (the full list of checkpoints can be found in the Supplementary Material).

4.3. Baselines

We compare the performance of **CBVLM** against several baselines, including:

(i) **standard CBM** [14] trained on each of the four datasets considered using the official implementation. It is worth noting that, for training the CBM on the Derm7pt dataset, we considered 11 clinical concepts. Specifically, 3 out of the 5 annotated concepts have three options each (e.g., Pigment Network can be Absent, Regular, or Irregular), as described in 4.1. To address this, and given that the concept layer predicts the presence or absence of each concept, we treated each option as an independent concept. For instance, for Pigment Network, we defined the concepts: Absent Pigment Network, Regular Pigment Network, and Irregular Pigment Network;

(ii) **CLAT** [41] - a concept-based framework specifically designed for retinal disease diagnosis. For DDR, the pre-computed textual embeddings of the clinical concepts provided in the official CLAT repository were used. However, since the original embeddings were obtained from a retinal foundation model [42], for the other datasets, we replaced it

² Mistral-7B-Instruct-v0.3




Fig. 3. Concept detection performance of LVLMs across different n -shot settings. Each bar corresponds to an n -shot scenario ($n = \{0, 1, 2, 4\}$). Filled colored bars denote BACC, whereas the hatched bars indicate F1-scores. Red crosses indicate the percentage of unknowns, i.e. the proportion of samples whose LVLMs' responses do not contain sufficient information to answer the posed question..

with MedImageInsight [43], an embedding model for medical imaging trained on a mixture of X-ray, computed tomography, dermatology, and pathology datasets. For Derm7pt, we follow the same strategy described above, i.e. we transformed the 5 concepts into 11 independent concepts;

(iii) **supervised black-box models**, namely ImageNet-pretrained ResNet50 [44] and ViT Base [45] fine-tuned on the disease diagnosis task;

(iv) **task-specific black-box models**, namely [17] for the Derm7pt dataset, [46] for the SkinCon dataset, [47] for the CORDA dataset, and [48] for the DDR dataset, reporting classification performance on the respective dataset.

4.4. Retrieval module

For the retrieval strategy used to perform ICL, image similarity is determined using cosine similarity between features, which are extracted from two distinct encoders, namely MedImageInsight [43] and the (vision) encoder of the corresponding LVLm.

4.5. Evaluation metrics

For both concept prediction and disease diagnosis, we report Balanced Accuracy (BACC) and F1-score, as these metrics are tailored to deal with imbalanced datasets, which is the case in medical image classification. For concept prediction, we measure the BACC and F1-score averaged over each concept.

5. Results and analysis

This section presents the key results of *CBVLM*, comparing its performance against CBMs and task-specific supervised methods in terms of BACC and F1-score across two tasks: (i) Concept Detection (Figs. 3, 4, 5, and 6); and (ii) Disease Diagnosis (Figs. 7 and 8).

5.1. Concept detection performance

We start by assessing the predictive performance of *CBVLM* in the task of concept detection. Fig. 3 presents the results across different datasets under various n -shot settings, $n \in \{0, 1, 2, 4\}$. For comparative analysis, we also report the performance of CBM and CLAT. The results for $n > 0$ were obtained considering the best vision encoder, in this case, MedImageInsight, selected after an ablation study with different vision encoders (see Fig. 9).

Few-shot prompting boosts performance. The results presented in Fig. 3 and the analysis depicted in Fig. 4 (left) reveals that few-shot prompting significantly boosts performance by incorporating contextually relevant information within the prompts [1]. Notably, there is a consistent trend across all datasets, demonstrating improved performance as the number of demonstrations increases. For example, the performance of *CBVLM* when adopting CheXagent, LLaVA-OV, Qwen2-VL, and SkinGPT-4 improves by over 15% when incorporating 4 demonstration examples into the prompt compared to the zero-shot scenario ($n = 0$). Remarkably, *CBVLM* outperforms both CBM and CLAT across all datasets except Derm7pt, where its performance is on par with CBM. Impressively, *CBVLM* achieves these results without any training, relying on only four demonstration examples or less.




Fig. 4. Concept detection results per dataset averaged over all models (left) and over generic and medical LVLMs (right). Each bar corresponds to the number of shots ($n = \{0, 1, 2, 4\}$). Filled colored bars denote BACC, whereas the hatched bars indicate F1-scores.




Fig. 5. Per-concept F1-scores across datasets averaged over all (blue), generic (green), and medical (orange) LVLMs, in the 4-shot scenario. Values in parentheses indicate the number of samples in the dataset for the corresponding concept.

Medical LVLMs outperform generic LVLMs in few-shot settings. From the results depicted in Fig. 4 (right), in the 0-shot scenario, generic LVLMs consistently outperform medical LVLMs across all datasets. However, in the few-shot setting, medical LVLMs outperform their generic counterparts in all datasets using only one or two demonstration examples ($n \in \{1, 2\}$). The stronger performance of generic LVLMs in the zero-shot setting can be attributed to their training on large-scale generic image datasets, allowing for better generalization in this setting. As demonstration examples are introduced into the prompt, medical LVLMs improve, benefiting from their pre-training on specialized medical data, which enhances their ability to generalize in these scenarios. However, in the 4-shot setting, generic LVLMs tend to achieve higher concept detection performance than their medical counterparts. This might stem from the fact that generic LVLMs tend to allow for longer prompts and are trained with interleaved image-text data, while medical LVLMs are not.

Underrepresented concepts are harder to predict. Fig. 5 presents a detailed comparison of the per-concept F1-score for each dataset. The

reported results, averaged across all, generic, and medical LVLMs in the 4-shot setting, reveal that underrepresented concepts are the most challenging to detect, as expected. For instance, in the CORDA dataset, none of the models correctly predicts the presence of the “pleural other” concept (which only appears in 4 images of the entire dataset). The concepts “atelectasis” and “lung lesion” are also consistently difficult to detect. The full results per-LVLM can be found in the Supplementary Material.

CBVLM outperforms CBMs. CBVLM outperforms CBMs and CLAT in concept prediction across all datasets, except for Derm7pt, where CBM achieves the best result with a margin of $\approx 5\%$. This may be due to the fact that in Derm7pt, some concepts have three available options, whereas in the other datasets, each concept is binary, which may make it easier for the LVLM to predict the presence or absence of a concept more accurately.

CBVLM outperforms CBM when using only 10% annotated data. Although CBVLM is able to outperform CBMs with only 4-shots, it should be noted that, in total, more than 4 images are being used, as the same




Fig. 6. Influence of the example set size on the concept prediction performance of *CBVLM*. Even when only 10% of the example images per class are available, *CBVLM* outperforms CBM, except for Derm7pt. Nevertheless, with only 10% of the examples, the decrease in performance is negligible for Derm7pt. Thus, *CBVLM* indeed requires few annotated examples.




Fig. 7. Disease diagnosis performance of LVLMs across different n -shot settings. Each bar corresponds to an n -shot scenario ($n = \{0, 1, 2, 4\}$). "0 w/o" corresponds to the 0-shot experiment, where no concepts are used for the disease diagnosis. Filled colored bars denote BACC whereas the hatched bars indicate F1-scores. Red crosses indicate the percentage of unknowns, i.e. the proportion of samples whose LVLMs responses do not contain sufficient information to answer the posed question.

4 example images will not be the most similar for all query images in the test set. To assess the impact of the number of images per class in the examples set, we progressively reduce the percentage of images per class, limiting the number of available examples for selection for few-shot prompting. This experiment is conducted using the best-performing model for each dataset: CheXagent 4-shots for Derm7pt, and Open-Flamingo 2-shots for SkinCon, CORDA and DDR. The results, shown in Fig. 6 and averaged across 3 runs, demonstrate that even with just 10% of the images per class in the examples set, *CBVLM* consistently outperforms CBM across all datasets, except for Derm7pt where even with 100% of the images per class, the performance is lower than CBM, as discussed earlier. Nevertheless, the results shown in Fig. 6 underscore a key observation: few-shot prompting with a good example selection strategy can boost performance by up to 20%

(cf. BACC at 0% and 10%), even in situations where annotated data is severely limited, which is the case of medical data. Thus, LVLMs can indeed be successfully used for concept detection with very limited amounts of annotated data, thus reducing the annotation cost inherent to CBMs.

5.2. Disease diagnosis performance

Fig. 7 shows the performance of *CBVLM* in terms of BACC and F1-score for the task of disease classification across different datasets under various n -shots ($n = \{0, 1, 2, 4\}$). Similarly to Section 5.1, we also report the performance of CBM, CLAT and a Black-box (Bbox) model for comparison. Additionally, we include results for the setting where




Fig. 8. Disease diagnosis results per dataset averaged over all models (left) and over generic and medical LVLMs (right). Each bar corresponds to the number of shots ($n = \{0, 1, 2, 4\}$). “0 w/o” corresponds to the 0-shot experiment where no concepts are used for the disease diagnosis. Filled colored bars denote BACC, whereas the hatched bars indicate F1-scores.




Fig. 9. Concept detection results with 1-shot and several example selection strategies. “random” corresponds to random selection of 1 example, while “random_per_class” corresponds to selecting 1 example per class randomly. “rices” is related to the RICES methodology [24], where the demonstration example is chosen based on the cosine similarity to the query image’s features. “rices_medii” corresponds to using RICES when MedImageInsight [43] is the vision encoder used to extract the features from the query and demonstration images, while in “rices_model” the encoder is the vision model included in the corresponding LVLm.

the CBVLM predicts the disease label directly without prior concept information (“0 w/o” bar in Fig. 7).

Few-shot prompting boosts performance. Similar to the trend observed in the concept detection task, the results shown in Fig. 8 (left) confirm our initial assumption and align with findings reported in the literature: few-shot prompting significantly enhances the performance of LVLms in disease classification as the number of demonstrations increases, across all datasets.

Using concepts improves performance and promotes explainability. The results show that incorporating clinical concepts into the prompt improves both BACC and F1-score compared to the scenario without concept guidance (i.e., “0” vs “0 w/o” bars in Fig. 7). Furthermore,

grounding the final diagnosis on a set of clinical concepts promotes transparency and explainability, aligning more closely with human reasoning, unlike the concept-free setting (“0 w/o”).

Medical LVLms outperform generic LVLms in few-shot settings. The results in Fig. 8 (right) confirm that generic LVLms consistently outperform their medical counterparts across all datasets in the scenario without concepts (“0 w/o”). This trend is also observed in the 0-shot scenario with concepts, in general. In the few-shot setting, generic LVLms exhibit superior performance only in SkinCon in terms of BACC, with a marginal average improvement of 1% across scenarios with $n > 0$. Conversely, when averaging across all few-shot scenarios ($n > 0$), medical LVLms consistently outperform generic models across all datasets, achieving F1-score improvements of 25.74%, 14.48%, 4.62%,

and 19.81% for Derm7pt, SkinCon, CORDA, and DDR, respectively. This phenomenon can be attributed to two factors: i) the stronger domain knowledge of medical LVLMs, and (ii) the initialization of their vision encoders from medical-specific CLIP versions, such as Biomed-CLIP [49] for LLaVA-Med, or larger improved CLIP architectures like EVA-CLIP-g [50], used in CheXagent and SkinGPT.

CBVLM outperforms CBMs. *CBVLM* surpasses CBMs across all datasets. With the top-performing LVLM, we observe F1-score improvements of 7.42%, 39.70%, 7.25%, and 63.13% for Derm7pt, SkinCon, CORDA, and DDR, respectively, compared to the standard CBM. It is worth recalling that *CBVLM* achieves these results without any kind of training, whereas CBM needs to be trained for each dataset. When comparing *CBVLM* with CLAT [41], a method specifically optimized for retinal disease diagnosis, our experiments show that *CBVLM* outperforms it on the diabetic retinopathy dataset, DDR, in terms of F1-score. Furthermore, *CBVLM* also outperforms CLAT on the other datasets, achieving an average relative improvement of $\approx 6\%$.

CBVLM outperforms supervised black-box models. Finally, when compared to black-box approaches, a similar trend is observed. *CBVLM* outperforms task-specific black-box models by 2.15% in BACC for Derm7pt and achieves comparable performance on CORDA, with a slight decrease of -2.32% . In terms of F1 score, the best performing *CBVLM* exceeds black-box task-specific models by 0.36% on SkinCon and by 1% on DDR. Thus, *CBVLM* achieves state-of-the-art results on Derm7pt, SkinCon and DDR.

5.3. Ablation study: Strategies for ICL example selection

In order to validate the importance of the *Retrieval Module* (cf. Fig. 2), which employs RICES [24] (i.e. ranks examples based on the cosine similarity between image features extracted from an arbitrary vision encoder), we conduct an ablation study on the task of concept detection for the 1-shot scenario.

In this study, we first compare random ICL example selection with RICES, and then assess different vision encoders to perform RICES. For random selection, we compare choosing only 1 random example against choosing 1 random example per class. The results of Fig. 9 show that, except for Open-Flamingo and MiniCPM-V, using 1 example per class does not improve performance. In fact, in the majority of cases it slightly decreases it, when comparing to simply choose 1 random example.

We then explore two different encoders for RICES: MedImageInsight [43] and the encoder integrated within the respective LVLM. Although in some models (e.g., Open-Flamingo and MiniCPM-V), choosing the examples randomly brings better performance, in the majority of cases it improves both BACC and F1-score (especially for Medical LVLMs). More importantly, we are not concerned with the average results, but rather on maximizing them, even if just for one LVLM, to be able to find the best configuration of *CBVLM* to compete against other approaches like CBM and supervised black-box models. Thus, using RICES beats random selection and using MedImageInsight instead of the LVLM encoder yields the best results on all datasets. As such, all experiments are conducted using RICES (i.e. the *Retrieval Module*) with MedImageInsight as the feature extractor.

5.4. General discussion

Table 2 presents a detailed comparison of our proposed *CBVLM* with concept-based models, supervised black-box models, and task-specific models. As highlighted in the results, *CBVLM* consistently achieves either the best or second-best performance across all evaluations. Notably, this is accomplished without any additional training and by leveraging just four examples, enabling concept-based explanations and enhanced interpretability and transparency in the model responses. In summary, *CBVLM* using CheXagent with 4 demonstration examples

emerges as the top-performing combination for Derm7pt. For SkinCon, *CBVLM* with the generic Open-Flamingo model and 4 demonstration examples achieves the best results in both disease diagnosis and concept detection tasks. For CORDA, the best performance is achieved by *CBVLM* with CheXagent with 2 demonstration examples. For DDR, *CBVLM* with Idefics3 and 4 demonstrations proves to be the best-performing combination.

Overall, the results underscore the effectiveness of *CBVLM* in enabling arbitrary LVLMs to predict the presence of clinical concepts from an image and use those concepts to ground the final diagnosis, enhancing interpretability and transparency in the model's responses, all without requiring any training and with just a few annotated examples. Notably, *CBVLM* outperforms CBM-related approaches across all datasets and even surpasses black-box models when employing ICL.

5.5. Limitations

The main drawback of the proposed methodology for generating concept-based explanations with LVLMs lies in the independent prediction of each concept, which is a concern when applied to datasets with a large number of concepts. A similar problem arises with datasets with a large number of classes, as each class is provided as a separate option in the prompt. This can result in excessively long prompts that would require alternative ways of querying the model. Although in medical applications the number of concepts is typically limited, a potential future direction to address concerns regarding the scalability for datasets with numerous concepts could be the adoption of tailored selection policies that prioritize only a small subset of the most informative concepts, as demonstrated in recent works [51].

Additionally, our approach is limited to discrete-valued concepts (e.g., absent/present or, for some concepts of Derm7pt, absent/regular/irregular), contrasting with traditional CBMs, where each concept is associated with a probability. Nevertheless, this limitation could be addressed by prompting the LVLM to not only answer if a given concept is present in the input image but also to provide a confidence score, as is done in other works [7]. We leave this exploration to future work.

Finally, *CBVLM* is only as good as the LVLM it uses. In a few cases, some of the LVLMs we tested struggled to provide answers to the questions (e.g., instead of answering the question, the LVLM simply described the input image), leading to an increase in unknown responses (as shown in Figs. 3 and 7). This phenomenon is particularly common with SkinGPT, which may be due to its pretraining on very specific prompts designed to elicit desired responses [4]. LLaVA-Med also faced challenges in providing valid responses, especially in few-shot scenarios where $n > 1$. This difficulty may stem from the fact that it is not trained with interleaved image-text data; in such models, image data is given at the input in the form of tokens, so as the number of shots increases, so does the number of tokens, making it more difficult for the model to pay attention to the full input sequence [52].

Another relevant concern involves the robustness, security, and inherent biases of LVLMs [53]. While a comprehensive analysis of these aspects is beyond the scope of this work, it remains a critical consideration for future studies aimed at safe and reliable integration of LVLMs into clinical workflows. Collaboration between clinicians and AI researchers is crucial to ensure transparency and explainability of AI models deployed into clinical practice [54].

6. Conclusions and future work

We propose a simple yet effective methodology, *CBVLM*, that leverages off-the-shelf LVLMs to produce concept-based explanations and predict disease diagnoses grounded on those explanations, thus offering a more transparent and explainable decision-making process, something that is critical in high-stakes scenarios such as medical use-cases. Through extensive experiments across four medical datasets, and open-source LVLMs, we show that *CBVLM* outperforms the traditional CBM

Table 2
Comparison of the best performing CBVLM with state-of-the-art supervised models and CBMs for disease diagnosis and concept detection. BACC refers to Balanced Accuracy, and F1 refers to the F1-score. The best performance is highlighted in **bold**, while the second-best performance is indicated with underlining.

Dataset	Method	Disease Diagnosis		Concept Detection	
		BACC (%)	F1 (%)	BACC (%)	F1 (%)
Derm7pt	CBM [14]	74.01	75.66	75.68	67.11
	CLAT [41]	79.67	<u>82.98</u>	66.69	54.76
	Black-box (ResNet50)	75.33	80.76	-	-
	Black-box (ViT Base)	82.26	82.04	-	-
	Black-box (Task-Specific) [17]	83.20	-	-	-
	CBVLM (CheXagent & 4-shot)	85.35	83.08	<u>70.21</u>	<u>59.55</u>
SkinCon	CBM [14]	49.97	46.51	<u>72.81</u>	<u>56.33</u>
	CLAT [41]	68.21	86.76	62.64	45.52
	Black-box (ResNet50)	73.98	81.49	-	-
	Black-box (ViT Base)	<u>80.25</u>	85.85	-	-
	Black-box (Task-Specific) [46]	-	77.80	-	-
	CBVLM (Open-Flamingo & 4-shot)	85.61	<u>86.21</u>	79.78	57.91
CORDA	CBM [14]	70.75	70.98	<u>76.20</u>	<u>61.92</u>
	CLAT [41]	75.64	<u>77.96</u>	<u>70.76</u>	<u>53.06</u>
	Black-box (ResNet50)	73.17	73.28	-	-
	Black-box (ViT Base)	<u>78.03</u>	77.63	-	-
	Black-box (Task-Specific) [47]	80.00	-	-	-
	CBVLM (CheXagent & 2-shot)	77.68	78.23	81.38	64.33
DDR	CBM [14]	23.62	21.25	59.05	39.06
	CLAT [41]	78.87	72.81	<u>76.64</u>	<u>64.53</u>
	Black-box (ResNet50)	33.70	57.00	-	-
	Black-box (ViT Base)	39.09	59.37	-	-
	Black-box (Task-Specific) [48]	-	<u>83.38</u>	-	-
	CBVLM (Idefics3 & 4-shot)	66.65	84.38	82.69	75.73

and even task-specific black-box models across all benchmarks. This is achieved without requiring any training and using just a few annotated examples, thus tackling a very common problem in medical applications, where the annotation process requires a clinician’s expertise. Overall, *CBVLM* is able to leverage the best of both worlds: (i) its LVLm-based training-free nature presents a major advantage over CBMs, as these models require additional training and are restricted by predefined concepts, and (ii) its two-stage process keeps the interpretability inherent to CBMs but that LVLms lack.

Although *CBVLM* was tested on medical datasets, it is a general methodology that can be applied in other domains. In the future, we would like to do so, e.g., on the CUB dataset [55]. Given the promising results obtained in this work with open-source LVLms of around 8 billion parameters, we would also like to experiment with bigger and/or proprietary LVLms.

CRedit authorship contribution statement

Cristiano Patrício: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Conceptualization. **Isabel Rio-Torto:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Conceptualization. **Jaime S. Cardoso:** Writing – review & editing, Validation, Supervision, Methodology, Conceptualization. **Luís F. Teixeira:** Writing – review & editing, Validation, Supervision, Methodology, Conceptualization. **João C. Neves:** Writing – review & editing, Validation, Supervision, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was funded by the Portuguese Foundation for Science and Technology (FCT) under the PhD grants “2020.07034.BD” (<https://doi.org/10.54499/2020.07034.BD>) and “2022.11566.BD” (<https://doi.org/10.54499/2022.11566.BD>), and supported by UID/04516/NOVA Laboratory for Computer Science and Informatics (NOVA LINC) with the financial support of FCT/IP.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.compbiomed.2025.111145>.

References

[1] M. Moor, Q. Huang, S. Wu, M. Yasunaga, Y. Dalmia, J. Leskovec, C. Zakka, E.P. Reis, P. Rajpurkar, Med-flamingo: a multimodal medical few-shot learner, in: Proceedings of the 3rd Machine Learning for Health Symposium, in: Proceedings of Machine Learning Research, vol. 225, PMLR, 2023, pp. 353–367.

[2] C. Li, C. Wong, S. Zhang, N. Usuyama, H. Liu, J. Yang, T. Naumann, H. Poon, J. Gao, LLaVA-Med: Training a large language-and-vision assistant for biomedicine in one day, Adv. Neural Inf. Process. Syst. (NeurIPS) 36 (2024).

[3] Z. Chen, M. Varma, J.-B. Delbrouck, M. Paschali, L. Blankemeier, D. Van Veen, J.M.J. Valanarasu, A. Youssef, J.P. Cohen, E.P. Reis, et al., CheXagent: Towards a foundation model for chest X-Ray interpretation, 2024, arXiv:2401.12208.

[4] J. Zhou, X. He, L. Sun, J. Xu, X. Chen, Y. Chu, L. Zhou, X. Liao, B. Zhang, S. Afvari, et al., Pre-trained multimodal large language model enhances dermatological diagnosis using SkinGPT-4, Nat. Commun. 15 (1) (2024) 5649.

[5] T. Han, L.C. Adams, S. Nebelung, J.N. Kather, K.K. Bressen, D. Truhn, Multimodal large language models are generalist medical image interpreters, MedRxiv (2023) 2012–2023.

[6] M.-H. Van, P. Verma, X. Wu, On large visual language models for medical imaging analysis: An empirical study, in: 2024 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies, CHASE, 2024, pp. 172–176.

[7] D. Ferber, G. Wölflein, I.C. Wiest, M. Liger, S. Sainath, N. Ghaffari Laleh, O.S. El Nahhas, G. Müller-Franzes, D. Jäger, D. Truhn, et al., In-context learning enables multimodal large language models to classify cancer pathology images, Nat. Commun. 15 (1) (2024) 10104.

- [8] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nat. Mach. Intell.* 1 (5) (2019) 206–215.
- [9] H. Liu, C. Li, Q. Wu, Y.J. Lee, Visual instruction tuning, in: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 36, 2023, pp. 34892–34916.
- [10] OpenAI, GPT-4 technical report, 2023, arXiv:2303.08774.
- [11] C. Royer, B. Menze, A. Sekuboyina, MultiMedEval: A benchmark and a toolkit for evaluating medical vision-language models, 2024.
- [12] J. Liu, W. Wang, S. Yihang, J. Huang, Y. Zhang, C.-Y. Li, W. Chen, X. Xing, K.-J. Chang, L. Shen, M.R. Lyu, Asclepius: A spectrum evaluation benchmark for medical multi-modal large language models, in: *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2025, pp. 24181–24201.
- [13] P. Xia, Z. Chen, J. Tian, Y. Gong, R. Hou, Y. Xu, Z. Wu, Z. Fan, Y. Zhou, K. Zhu, W. Zheng, Z. Wang, X. Wang, X. Zhang, C. Bansal, M. Niethammer, J. Huang, H. Zhu, Y. Li, J. Sun, Z. Ge, G. Li, J. Zou, H. Yao, CARES: A comprehensive benchmark of trustworthiness in medical vision language models, in: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 37, 2024, pp. 140334–140365.
- [14] P.W. Koh, T. Nguyen, Y.S. Tang, S. Mussmann, E. Pierson, B. Kim, P. Liang, Concept bottleneck models, in: *Proceedings of the International Conference on Machine Learning, ICML, PMLR*, 2020, pp. 5338–5348.
- [15] C. Patrício, J.C. Neves, L.F. Teixeira, 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 Workshops, CVPRW*, 2023, pp. 3798–3807.
- [16] Y. Bie, L. Luo, H. Chen, MICA: Towards explainable skin lesion diagnosis via multi-level image-concept alignment, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, (2) 2024, pp. 837–845.
- [17] C. Patrício, L.F. Teixeira, J.C. Neves, Towards concept-based interpretability of skin lesion diagnosis using vision-language models, in: *Proceedings of the IEEE International Symposium on Biomedical Imaging, ISBI, IEEE*, 2024, pp. 1–5.
- [18] S. Ghosh, K. Yu, F. Arabshahi, K. Batmanghelich, Bridging the gap: From post hoc explanations to inherently interpretable models for medical imaging, in: *Proceedings of the ICML Workshop on Interpretable Machine Learning in Healthcare*, 2023.
- [19] OpenAI, ChatGPT: Optimizing language models for dialogue, 2024, <https://chat.openai.com/>.
- [20] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, in: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 1877–1901.
- [21] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds, et al., Flamingo: a visual language model for few-shot learning, in: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, 2022, pp. 23716–23736.
- [22] G. Zhou, Z. Han, S. Chen, B. Huang, L. Zhu, S. Khan, X. Gao, L. Yao, Adapting large multimodal models to distribution shifts: The role of in-context learning, 2024, arXiv:2405.12217.
- [23] A. Awadalla, I. Gao, J. Gardner, J. Hessel, Y. Hanafy, W. Zhu, K. Marathe, Y. Bitton, S. Gadre, S. Sagawa, et al., OpenFlamingo: An open-source framework for training large autoregressive vision-language models, 2023, arXiv:2308.01390.
- [24] Z. Yang, Z. Gan, J. Wang, X. Hu, Y. Lu, Z. Liu, L. Wang, An empirical study of GPT-3 for few-shot knowledge-based VQA, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, (3) 2022, pp. 3081–3089.
- [25] S. Chen, Z. Han, B. He, M. Buckley, P. Torr, V. Tresp, J. Gu, Understanding and improving in-context learning on vision-language models, in: *ICLR 2024 Workshop on Mathematical and Empirical Understanding of Foundation Models*, 2024.
- [26] X. Wang, B. Ma, C. Hu, L. Weber-Genzel, P. Röttger, F. Kreuter, D. Hovy, B. Plank, “My answer is C”: First-token probabilities do not match text answers in instruction-tuned language models, in: *Findings of the Association for Computational Linguistics: ACL 2024*, 2024, pp. 7407–7416.
- [27] J. Kawahara, S. Daneshvar, G. Argenziano, G. Hamarneh, Seven-point checklist and skin lesion classification using multitask multimodal neural nets, *IEEE J. Biomed. Heal. Informatics (JBHI)* 23 (2) (2019) 538–546.
- [28] R. Daneshjoui, M. Yuksekgonul, Z.R. Cai, R. Novoa, J.Y. Zou, SkinCon: A skin disease dataset densely annotated by domain experts for fine-grained model debugging and analysis, in: *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, 2022, pp. 18157–18167.
- [29] M. Alesina, C.A. Barbano, C. Berzovini, M. Busso, M. Calandri, A.D. Pascale, A. Fiandrotti, P. Fonio, M. Grangetto, et al., CORDA dataset, 2023, <http://dx.doi.org/10.5281/zenodo.7821611>.
- [30] T. Li, Y. Gao, K. Wang, S. Guo, H. Liu, H. Kang, Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening, *Inform. Sci.* 501 (2019) 511–522.
- [31] M. Groh, C. Harris, L. Soenksen, F. Lau, R. Han, A. Kim, A. Koochek, O. Badri, Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset, in: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPRW, 2021, pp. 1820–1828.
- [32] C.A. Barbano, E. Tartaglione, C. Berzovini, M. Calandri, M. Grangetto, A two-step radiologist-like approach for Covid-19 computer-aided diagnosis from chest X-Ray images, in: *Image Analysis and Processing – ICIAP 2022*, 2022, pp. 173–184.
- [33] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghighi, R. Ball, K. Shpanskaya, et al., CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, (01) 2019, pp. 590–597.
- [34] H. Laurençon, A. Marafioti, V. Sanh, L. Tronchon, Building and better understanding vision-language models: insights and future directions, in: *Workshop on Responsibly Building the Next Generation of Multimodal Foundational Models*, 2024.
- [35] J. Lin, H. Yin, W. Ping, P. Molchanov, M. Shoenybi, S. Han, VILA: On pre-training for visual language models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, 2024, pp. 26689–26699.
- [36] B. Li, Y. Zhang, D. Guo, R. Zhang, F. Li, H. Zhang, K. Zhang, P. Zhang, Y. Li, Z. Liu, C. Li, LLaVA-OneVision: Easy visual task transfer, *Trans. Mach. Learn. Res.* (2025).
- [37] J. Bai, S. Bai, S. Yang, S. Wang, S. Tan, P. Wang, J. Lin, C. Zhou, J. Zhou, Qwen-VL: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023, arxiv:2308.12966.
- [38] Y. Yao, T. Yu, A. Zhang, C. Wang, J. Cui, H. Zhu, T. Cai, H. Li, W. Zhao, Z. He, et al., Minicpm-v: A GPT-4V level MLLM on your phone, 2024, arxiv:2408.01800.
- [39] Z. Chen, W. Wang, Y. Cao, Y. Liu, Z. Gao, E. Cui, J. Zhu, S. Ye, H. Tian, Z. Liu, L. Gu, X. Wang, Q. Li, Y. Ren, Z. Chen, J. Luo, J. Wang, T. Jiang, B. Wang, C. He, B. Shi, X. Zhang, H. Lv, Y. Wang, W. Shao, P. Chu, Z. Tu, T. He, Z. Wu, H. Deng, J. Ge, K. Chen, K. Zhang, L. Wang, M. Dou, L. Lu, X. Zhu, T. Lu, D. Lin, Y. Qiao, J. Dai, W. Wang, Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling, 2025, arXiv:2412.05271.
- [40] J. Ye, H. Xu, H. Liu, A. Hu, M. Yan, Q. Qian, J. Zhang, F. Huang, J. Zhou, mPLUG-Owl3: Towards long image-sequence understanding in multi-modal large language models, 2024, arXiv:2408.04840.
- [41] C. Wen, M. Ye, H. Li, T. Chen, X. Xiao, Concept-based lesion aware transformer for interpretable retinal disease diagnosis, *IEEE Trans. Med. Imaging* 44 (1) (2025) 57–68.
- [42] J. Silva-Rodríguez, H. Chakor, R. Kobbi, J. Dolz, I. Ben Ayed, A foundation language-image model of the retina (FLAIR): encoding expert knowledge in text supervision, *Med. Image Anal.* 99 (2025) 103357.
- [43] N.C. Codella, Y. Jin, S. Jain, Y. Gu, H.H. Lee, A.B. Abacha, A. Santamaria-Pang, W. Guyman, N. Sangani, S. Zhang, et al., MedImageInsight: An open-source embedding model for general domain medical imaging, 2024, arXiv:2410.06542.
- [44] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2016, pp. 770–778.
- [45] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An image is worth 16x16 words: Transformers for image recognition at scale, in: *International Conference on Learning Representations, ICLR*, 2021.
- [46] J. Hou, J. Xu, H. Chen, Concept-attention whitening for interpretable skin lesion diagnosis, in: *Medical Image Computing and Computer Assisted Intervention, MICCAI*, 2024, pp. 113–123.
- [47] C.A. Barbano, R. Renzulli, M. Grosso, D. Basile, M. Busso, M. Grangetto, AI-assisted diagnosis for Covid-19 CXR screening: from data collection to clinical validation, in: *Proceedings of the IEEE International Symposium on Biomedical Imaging, ISBI*, 2024, pp. 1–4.
- [48] S. Madarapu, S. Ari, K. Mahapatra, A multi-resolution convolutional attention network for efficient diabetic retinopathy classification, *Comput. Electr. Eng.* 117 (2024) 109243.
- [49] S. Zhang, Y. Xu, N. Usuyama, H. Xu, J. Bagga, R. Tinn, S. Preston, R. Rao, M. Wei, N. Valluri, C. Wong, A. Tupini, Y. Wang, M. Mazzola, S. Shukla, L. Liden, J. Gao, A. Crabtree, B. Piening, C. Bifulco, M.P. Lungren, T. Naumann, S. Wang, H. Poon, A multimodal biomedical foundation model trained from fifteen million image-text pairs, *NEJM AI* 2 (1) (2025) A1042400640.
- [50] Q. Sun, Y. Fang, L. Wu, X. Wang, Y. Cao, EVA-CLIP: Improved training techniques for CLIP at scale, 2023, arxiv:2303.15389.
- [51] K. Chauhan, R. Tiwari, J. Freyberg, P. Shenoy, K. Dvijotham, Interactive concept bottleneck models, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, (5) 2023, pp. 5948–5955.

- [52] T. Li, G. Zhang, Q.D. Do, X. Yue, W. Chen, Long-context LLMs struggle with long in-context learning, *Trans. Mach. Learn. Res.* (2025).
- [53] H. Zhang, W. Shao, H. Liu, Y. Ma, P. Luo, Y. Qiao, N. Zheng, K. Zhang, B-AVIBench: Toward evaluating the robustness of large vision-language model on black-box adversarial visual-instructions, *IEEE Trans. Inf. Forensics Secur.* 20 (2025) 1434–1446.
- [54] M. Phung, V. Muralidharan, V. Rotemberg, R.A. Novoa, A.S. Chiou, C.Y. Sadée, B. Rapaport, K. Yekrang, J. Bitz, O. Gevaert, et al., Best practices for clinical skin image acquisition in translational artificial intelligence research, *J. Invest. Dermatol.* 143 (7) (2023) 1127–1132.
- [55] C. Wah, S. Branson, P. Welinder, P. Perona, S. Belongie, The caltech-UCSD birds-200–2011 dataset, 2011.