

MAKE: Multi-Aspect Knowledge-Enhanced Vision-Language Pretraining for Zero-shot Dermatological Assessment

Siyuan Yan^{1,2*}, Xieji Li^{2*}, Ming Hu^{1,2}, Yiwen Jiang^{1,2}, Zhen Yu²,
Zongyuan Ge²(✉)

¹ Faculty of Engineering, Monash University, Melbourne, Australia

² AIM for Health Lab, Monash University, Victoria, Australia

{siyuan.yan,zongyuan.ge}@monash.edu

Abstract. Dermatological diagnosis represents a complex multimodal challenge that requires integrating visual features with specialized clinical knowledge. While vision-language pretraining (VLP) has advanced medical AI, its effectiveness in dermatology is limited by text length constraints and the lack of structured texts. In this paper, we introduce MAKE, a **M**ulti-**A**spect **K**nowledge-**E**nhanced vision-language pretraining framework for zero-shot dermatological tasks. Recognizing that comprehensive dermatological descriptions require multiple knowledge aspects that exceed standard text constraints, our framework introduces: (1) a multi-aspect contrastive learning strategy that decomposes clinical narratives into knowledge-enhanced subtexts through large language models, (2) a fine-grained alignment mechanism that connects subtexts with diagnostically relevant image features, and (3) a diagnosis-guided weighting scheme that adaptively prioritizes different subtexts based on clinical significance prior. Through pretraining on 403,563 dermatological image-text pairs collected from the internet, MAKE significantly outperforms state-of-the-art VLP models on seven datasets across zero-shot skin disease classification, concept annotation, and cross-modal retrieval tasks. Our code is available at <https://github.com/SiyuanYan1/MAKE>.

Keywords: Dermatology · Vision-language · Knowledge augmentation.

1 Introduction

Dermatological diagnosis represents a complex multimodal challenge [14], requiring clinicians to simultaneously interpret visual features of skin lesions alongside patient history and various clinical concept descriptions [1,12,23] for accurate assessment. While deep learning has advanced automated diagnosis systems [3], standard approaches face significant limitations when applied to dermatology. Specifically, supervised [3,14] and self-supervised techniques [22,18] require extensive labeled data for different tasks and struggle to capture the rich, mul-

* Equal contribution

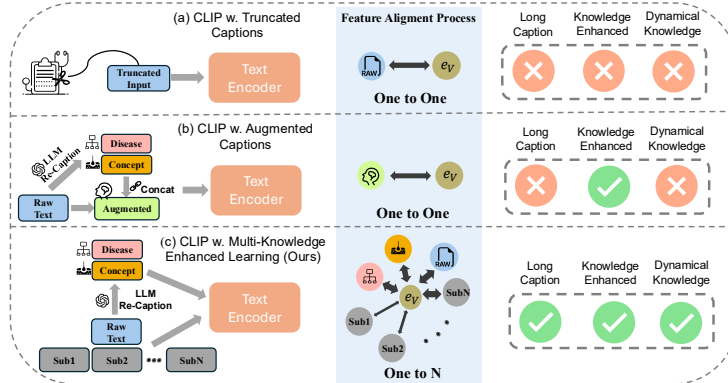


Fig. 1. Comparison between training strategies. Our framework utilizes CLIP with multi-knowledge enhanced learning, addressing long caption modeling limitations. It enables dynamic knowledge modeling between each image and its multiple corresponding captions, each capturing diverse aspects of crucial dermatological knowledge.

timodal nature of dermatological knowledge. This multimodal complexity necessitates more sophisticated approaches that can effectively bridge visual and linguistic understanding for dermatological diagnosis.

In parallel, vision-language pretraining (VLP) [16] has emerged as a powerful paradigm for multimodal tasks, demonstrating remarkable generalization capabilities without extensive fine-tuning, a capability formally known as zero-shot learning. By learning from rich textual information rather than single-label annotations, these models offer a way to alleviate the heavy dependence on labeled data. However, applying VLP models in dermatology faces two significant challenges. First, conventional VLP frameworks like CLIP [16] typically limit text input to a fixed token length (e.g., 77 tokens), truncating any longer description (Fig 1a). This truncation oversimplifies the rich clinical narrative and discards vital diagnostic details, ultimately constraining the model’s ability to capture the nuanced clinical concepts of skin lesions. Second, dermatology lacks standardized image-text pairs that are crucial for effective VLP training. Unlike other medical specialties such as radiology [10,9] where structured reports provide well-organized image-text pairs, dermatology often relies on unstructured clinical narratives. Recent works [8,12,7] attempt to crawl image-text pairs from web sources, which frequently yield noisy text. Some approaches [4,20] have explored knowledge augmentation via large language models (LLMs) to generate more comprehensive knowledge-enhanced captions, as illustrated in Fig. 1b. Yet, VLP models trained using these methods still struggle to model the complex inter-relationships among multiple aspects of clinical knowledge—such as lesion morphology, standardized disease descriptions, and associated symptoms—within a single short-length description. Additionally, existing methods [27] treat all in-

formation equally, neglecting the varying contributions of different aspects of knowledge to the final diagnosis.

To overcome these dermatology-specific challenges, we propose MAKE, a **M**ulti-**A**spect **K**nowledge-**E**nhanced vision-language pretraining framework specifically designed for zero-shot dermatological tasks (Fig. 1c). Our framework introduces three key innovations: First, a multi-aspect knowledge-image contrastive learning strategy that decomposes complex dermatological descriptions into multiple sub-captions, each capturing distinct aspects of clinical knowledge such as morphology, distribution patterns, and associated symptoms. This approach not only mitigates the text length constraint but also enables precise alignment between visual features of skin lesions and various aspects of clinical knowledge, critical for differential diagnosis. Second, a fine-grained alignment mechanism that associates multiple sub-captions with diagnostically relevant image patches of skin lesions, enabling different aspects of dermatological knowledge to jointly characterize the salient visual features crucial for accurate diagnosis. Third, a diagnosis-guided weighting scheme that adaptively prioritizes different aspects of knowledge based on their diagnostic relevance in dermatology practice, better reflecting how dermatologists assign varying importance to different clinical attributes during the diagnostic reasoning process.

In summary, our contributions include: (1) introducing MAKE, the first vision-language pretraining framework for dermatology; (2) proposing three complementary technical innovations described above that enable fine-grained knowledge-enhanced visual-textual learning for dermatological applications; and (3) verifying our method through pretraining on Derm1M dataset [21], which comprises 403,563 dermatological image-text pairs. Extensive experiments demonstrate that MAKE significantly outperforms state-of-the-art VLP models on zero-shot skin disease classification, concept annotation, and cross-modal retrieval tasks across eight datasets..

2 Methodology

As illustrated in Fig. 2, our MAKE framework comprises three core components: multi-aspect knowledge-image contrastive learning, fine-grained alignment, and diagnosis knowledge-guided weighting. We detail each of them below.

2.1 Encoding Stage

The original dataset consists of image-text pairs $\mathcal{D} = \{(I_i, T_i^r)\}_{i=1}^M$, where I_i denotes the i -th skin image and T_i^r represents its associated raw text description. Traditional VLMs like CLIP [16] are constrained by text length limitations and cannot effectively leverage the rich clinical knowledge in these descriptions.

To address this limitation, we expand each image-text pair (I_i, T_i^r) into a richer multi-aspect representation through two complementary augmentation methods: 1) **Knowledge extraction**: We use LLMs to extract and generate two specialized knowledge aspects from the original text: - *Disease aspect text*

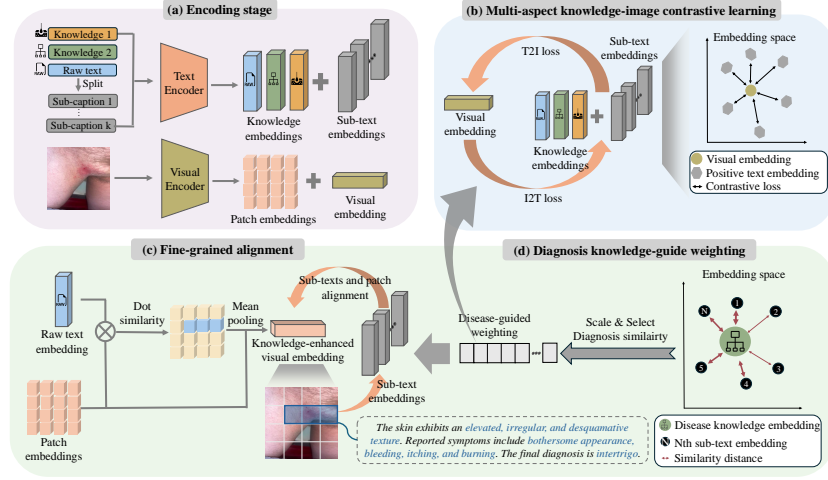


Fig. 2. Overview of our MAKE framework. (a) Encoding multi-aspect clinical knowledge. (b) The process of aligning visual embeddings with multiple positive text embeddings. **Knowledge 1:** *concept-aspect text*; **Knowledge 2:** *disease-aspect text*. (c) A fine-grained alignment process, which matches each subtext embedding with knowledge-enhanced visual embeddings. (d) Diagnosis similarity-based weights that modulate alignment between subtexts and visual embeddings.

(t_i^d): Contains standardized disease terminology, synonyms, and hierarchical relationships (e.g., melanoma’s superclass is “malignant”). - *Concept aspect text* (t_i^c): Captures interpretable clinical descriptors from dermatological lexicons (e.g., “plaque”, “scale”, “erosion”) crucial for diagnosis. Detailed prompts are available at the link ³. 2) **Sentence decomposition**: We preserve the raw text’s detailed content by splitting T_i^r into multiple sentences, creating a subtext set $S_i = \{S_i^j\}_{j=1}^K$ of size K focusing on different aspects of the raw description. To this end, we obtain two complementary sets: a *knowledge set* $\{T_i^r, t_i^d, t_i^c\}$ containing the raw text and derived knowledge aspects, and a *subtext set* S_i from sentence decomposition.

Given a batch $\mathcal{B} = \{I_i, T_i^r\}_{i=1}^N$ of N samples, our framework transforms each image-text pair into an enhanced representation:

$$\mathcal{B}_{\text{enhanced}} = \{I_i, (T_i^r, t_i^d, t_i^c, \{S_i^j\}_{j=1}^K)\}_{i=1}^N \quad (1)$$

This approach overcomes text length constraints while capturing multi-aspect clinical knowledge.

Then, we process inputs through dedicated encoders. Using the vision encoder E_V , we obtain the normalized visual embedding $e_i^v = E_V(I_i)$ and patch embeddings $e_i^p = [v_i^1, \dots, v_i^{HW}]$. For textual part, we employ the text encoder E_T

³ <https://github.com/SiyuanYan1/MAKE/blob/main/assets/prompt.md>

to project each subtext from both the *knowledge set* and *subtext set* individually, resulting in $K + 3$ text embeddings:

$$e_i^t = (e_i^r, e_i^d, e_i^c, \{e_i^{s_j}\}_{j=1}^K) = (E_T(T_i^r), E_T(t_i^d), E_T(t_i^c), \{E_T(S_i^j)\}_{j=1}^K) \quad (2)$$

where K is the number of subtexts from *subtext set*, and 3 represents the embeddings of raw, *disease aspect*, and *concept aspect text* from *knowledge set*.

2.2 Multi-aspect Knowledge-Image Contrastive Learning

To optimize multiple aspect texts with their corresponding single image, we apply multi-positive contrastive learning [4], as shown in Fig. 2(b). We align the visual embedding with all associated $K + 3$ text embeddings from both the *knowledge set* and *subtext set* in the shared embedding space using an image-to-text multi-positive contrastive learning loss:

$$\mathcal{L}_{i2t}^{mkcl} = - \sum_{i=1}^N \sum_{j=1}^{K+3} \log \frac{\exp(\text{sim}\langle e_i^v, e_{ij}^t \rangle / \tau)}{\sum_{n=1}^N \exp(\text{sim}\langle e_n^v, e_{ij}^t \rangle / \tau)} \quad (3)$$

where e_{ij}^t represents the j -th text embedding of the i -th sample with $j \in \{r, d, c, s_1, \dots, s_K\}$ as defined in Eq. 2, $\text{sim}\langle \cdot, \cdot \rangle$ denotes cosine similarity, and τ is a learnable temperature parameter. Similarly, the text-to-image loss is:

$$\mathcal{L}_{t2i}^{mkcl} = - \sum_{i=1}^N \sum_{j=1}^{K+3} \log \frac{\exp(\text{sim}\langle e_{ij}^t, e_i^v \rangle / \tau)}{\sum_{n=1}^N \exp(\text{sim}\langle e_{ij}^t, e_n^v \rangle / \tau)} \quad (4)$$

The final multi-aspect knowledge-image contrastive learning loss is defined as $\mathcal{L}^{mkcl} = (\mathcal{L}_{t2i}^{mkcl} + \mathcal{L}_{i2t}^{mkcl})/2$.

2.3 Fine-grained Alignment

To enhance the fine-grained alignment capability of VLMs, we draw inspiration from dermatologists who leverage multiple knowledge aspects to characterize a skin lesion. As shown in Fig. 2(c), we align all subtexts from *subtext set* with specific knowledge-enhanced patches to improve fine-grained alignment.

Specifically, we first calculate dot product similarity between patch embeddings $e_i^p = [v_i^1, \dots, v_i^{HW}]$ and raw text embedding e_i^r to generate a normalized similarity map $z_i = e_i^r \cdot (e_i^p)^T$. Next, we compute the dot product between this similarity map and patch embeddings to highlight patches with strong knowledge-semantic relevance. Finally, we apply mean pooling to align the dimensionality of knowledge-enhanced visual embeddings with subtext embeddings. The knowledge-enhanced visual embedding is formulated as:

$$e_i^k = \sum_{n=1}^{HW} v_i^n \cdot \frac{z_i^n}{\sum_{j=1}^{HW} z_i^j} \quad (5)$$

where z_i^n is the n -th element of the similarity map z_i , v_i^n represents the n -th patch embedding, and HW is the total number of image patches.

To align each subtext embedding with the knowledge-enhanced visual embedding, we define the fine-grained alignment loss as:

$$\mathcal{L}_{slra} = - \sum_{i=1}^N \sum_{j=1}^K \log \frac{\exp(\text{sim}\langle e_i^{s_j}, e_i^k \rangle / \tau)}{\sum_{n=1}^N \exp(\text{sim}\langle e_i^{s_j}, e_n^k \rangle / \tau)} \quad (6)$$

where $e_i^{s_j}$ denotes the embedding of the j -th subtext for the i -th sample and e_i^k is the knowledge-enhanced visual embedding from Eq. 5.

2.4 Diagnosis Knowledge-guided Weighting

Mimicking how dermatologists prioritize clinical information, our approach adaptively weights text elements by diagnostic relevance. As shown in Fig. 2(d), we introduce a weighting mechanism reflecting clinical decision-making. For each sample, we compute subtext weights for *subtext set* by measuring semantic similarity between subtext embeddings $\{e_i^{s_j}\}_{j=1}^K$ and *disease aspect* embedding e_i^d :

$$w_i = \frac{\{e_i^d \cdot (e_i^{s_j})^T\}_{j=1}^K}{\max(\{e_i^d \cdot (e_i^{s_j})^T\}_{j=1}^K)} \quad (7)$$

This yields weight vector $w_i = [w_i^1, \dots, w_i^K]$ for each sample's K subtexts, normalized by maximum similarity. Batch-wide weights are denoted as $\hat{w} = \{w_i\}_{i=1}^N$. *Knowledge set* embeddings (raw text, disease, and concept aspects) receive default weights of 1 as they already contain rich diagnostic information. Our final loss integrates contrastive and fine-grained alignment losses, modulated by these diagnosis-guided weights.

$$\mathcal{L}_{total} = \hat{w}_{mkcl} \mathcal{L}_{mkcl} + \lambda \hat{w}_{slra} \mathcal{L}_{slra} \quad (8)$$

where λ balances the two loss terms, while \hat{w}_{mkcl} and \hat{w}_{slra} represent the weights applied to each respective loss component derived from Eq. 7.

3 Experimental details

Experiment Setup: We conduct pretraining on Derm1M⁴ [21], a dataset of 403,563 skin image-text pairs, including diverse data from PubMed, YouTube, Twitter, and clinical notes. We denote Derm1M[†] as our knowledge-augmented version where *disease aspect* and *concept aspect texts* are pre-pended to raw text. For evaluation, we use seven downstream datasets (PAD [15], Fitzpatrick17K [5], SD-128 [19], SNU-134 [6], SkinCon [1], Derm7pt [11], and SkinCAP [28]) across three categories: (1) Zero-shot disease classification for skin cancer and general

⁴ <https://github.com/SiyuanYan1/Derm1M>

Table 1. Zero-shot performance comparison for disease and concept classification.

Method	Pretrain Data	Disease Classification (ACC)					Concept Annotation (AUROC)		
		PAD	F17K	SD-128	SNU-134	Average	SkinCon	Derm7pt	Average
Class number		6	114	128	134		32	7	
Test size		2,298	16,577	5,619	2,101		3,855	1,011	
CLIP [16]	CLIP400M	0.4330	0.0628	0.0728	0.0733	0.1604	0.6642	0.5594	0.6118
BiomedCLIP [26]	PMC-15M	0.4295	0.0890	0.1321	0.0971	0.1869	0.6817	0.6092	0.6455
PMC-CLIP [13]	PMC-OA	0.4312	0.0268	0.0443	0.0390	0.1353	0.6251	0.5820	0.6036
MONET [12]	PubMed+TextBook	0.4308	0.1409	0.2072	0.1333	0.2280	0.7502	0.6889	0.7196
CLIP [16]	Derm1M	0.5957	0.2595	0.3349	0.2689	0.3648	0.7233	0.6707	0.6970
SigLIP [25]	Derm1M	0.5113	0.2718	0.3574	0.2875	0.3570	0.7649	0.6867	0.7258
CoCa/ViT-B-32 [24]	Derm1M	0.5635	0.2019	0.2862	0.1942	0.3115	0.6431	0.6289	0.6360
CLIP [16]	Derm1M [†]	0.5631	0.2559	0.3189	0.2485	0.3466	0.7348	0.6900	0.7124
SigLIP [25]	Derm1M [†]	0.5892	0.2369	0.3045	0.2385	0.3423	0.6860	0.5701	0.6281
CoCa/ViT-B-32 [24]	Derm1M [†]	0.5688	0.2040	0.2557	0.2289	0.3144	0.6715	0.6408	0.6562
MAKE (Ours)	Derm1M [†]	0.5953	0.3242	0.3914	0.3270	0.4095	0.7873	0.6864	0.7369

Table 2. Cross-modal retrieval performance comparison on the SkinCAP dataset.

Method	Pretrain Data	Image-to-Text			Text-to-Image			Average
		R@10	R@50	R@100	R@10	R@50	R@100	
CLIP [16]	CLIP400M	0.0913	0.2354	0.3407	0.0592	0.1860	0.2760	0.1981
BiomedCLIP [26]	PMC-15M	0.1359	0.3429	0.4698	0.1238	0.3304	0.4578	0.3101
PMC-CLIP [13]	PMC-OA	0.0672	0.1908	0.2783	0.0649	0.1855	0.2652	0.1753
MONET [12]	PubMed+TextBook	0.1421	0.3384	0.4568	0.1492	0.3490	0.4756	0.3185
CLIP [16]	Derm1M	0.1532	0.3590	0.4851	0.1552	0.3715	0.4728	0.3328
SigLIP [25]	Derm1M	0.1757	0.3783	0.4871	0.1843	0.3896	0.4974	0.3521
CoCa/ViT-B-32 [24]	Derm1M	0.1193	0.2865	0.3833	0.1291	0.3078	0.4089	0.2723
CLIP [16]	Derm1M [†]	0.1587	0.3700	0.4788	0.1592	0.3537	0.4675	0.3313
SigLIP [25]	Derm1M [†]	0.1835	0.3803	0.5006	0.1750	0.3793	0.4951	0.3523
CoCa/ViT-B-32 [24]	Derm1M [†]	0.1406	0.3241	0.4214	0.1429	0.3269	0.4202	0.2960
MAKE (Ours)	Derm1M [†]	0.2096	0.4440	0.5613	0.1995	0.4420	0.5628	0.4032

skin condition diagnosis; (2) Zero-shot concept annotation for identifying clinically relevant concepts that aid diagnosis and interpretability [12,1]; and (3) Cross-modal retrieval for both image-to-text and text-to-image retrieval. Following CLIP [16], we employ zero-shot evaluation without fine-tuning.

Implementation Details: Following CLIP [16], we use ViT-B/16 [2] as the image encoder and GPT2 [17] with a context length of 77 as the text encoder. Our proposed MAKE leverages all three text types (raw, *disease aspect*, and *concept aspect text*) of Derm1M[†]. To ensure a fair comparison with state-of-the-art VLM methods, we trained each baseline model in two configurations: one using only raw text (denoted as training on the Derm1M dataset) and another using knowledge-augmented text (denoted as training on the Derm1M[†] dataset). Models are trained for 15 epochs with a batch size of 2048, a learning rate of $1e-4$, and a 1500-step warm-up with a weight decay of 0.1 on a single NVIDIA H200 GPU (141GB). The loss weighting factor λ is 0.7, and images are processed at 224×224 resolution. For all models, we use the final checkpoint and conduct extensive hyperparameter tuning to find the optimal model.

4 Results

We compare our MAKE method with three groups of approaches across seven datasets on disease classification, concept annotation, and cross-modal retrieval.

General VLMs includes foundation models like BiomedCLIP [26], PMC-CLIP [13], and MONET [12] trained on natural, biomedical, or dermatological image-text pairs. *Standard VL methods* comprises SOTA vision-language approaches like CLIP [16], SigLIP [25], and CoCa [24] pretrained on Derm1M dataset. *Knowledge-enhanced VL methods* includes the same approaches as Group 2 but pretrained on the knowledge-augmented version (Derm1M[†]).

Zero-shot Skin Disease Classification and Concept Annotation: Table 1 presents results across six datasets for zero-shot disease classification and concept annotation. MAKE outperforms all methods on most datasets, with 3.95% accuracy improvement on SNU-134 and 5.24% on F17K over the best baseline. Overall, MAKE delivers 4.47% higher average accuracy on disease classification and 1.11% better average AUROC on concept annotation than the best baseline. Three key findings emerge: (1) *Standard VL Methods* trained on Derm1M substantially outperform *General VLMs*; (2) VLMs trained on Derm1M[†] often perform worse than on Derm1M, showing conventional VLMs cannot effectively utilize knowledge-augmented data; (3) In contrast, our MAKE framework effectively leverages the knowledge-augmented data and achieves superior performance through multi-aspect knowledge contrastive learning.

Cross-modal Retrieval: Table 2 evaluates VLMs’ zero-shot image-text and text-image retrieval capabilities on the SkinCAP [28] dataset. MAKE outperforms all baselines, achieving 44.4% and 44.2% R@50 for image-to-text and text-to-image retrieval, respectively. This represents a 6.37% improvement over the best baseline (SigLIP on Derm1M[†]) for image-to-text and a 5.24% gain for text-to-image retrieval, demonstrating MAKE’s superior capability in aligning visual and textual representations in the dermatology domain.

Ablation Study: We perform ablation studies to analyze each component of our model, as shown in Table 3. The first row represents our baseline, which uses the CLIP architecture. The second row shows CLIP with our multi-aspect knowledge-image contrastive learning loss (mkcl) using only *knowledge set*, improving the baseline by 2.03% in average accuracy. The third row incorporates mkcl with both *knowledge set* and *subtext set*, further improving performance by 0.96% compared to mkcl with only *knowledge set*, and by 2.99% compared to the baseline. When adding local alignment loss (slra) to mkcl, the average accuracy further increases by 0.76%. Finally, incorporating diagnosis knowledge-guided weighting (dkw) achieves the best average performance of 40.95%, which is 0.72% higher than using only mkcl and slra, and 4.47% higher than the baseline. These results demonstrate the effectiveness of all proposed components.

5 Conclusion

In this paper, we introduce MAKE, a Multi-Aspect Knowledge-Enhanced vision-language pretraining framework for dermatology that addresses limitations of

Table 3. Ablation study on different components of the MAKE framework. We report classification accuracy per dataset. # denotes mkcl using only raw text, disease aspect, and concept aspect texts without spitted text for training.

Module			Disease Classification (ACC)				
mkcl	slra	dkw	PAD	F17K	SD-128	SNU-134	Average
			0.5957	0.2595	0.3349	0.2689	0.3648
✓#			0.6023	0.3142	0.3465	0.2775	0.3851
✓			0.6062	0.3210	0.3533	0.2984	0.3947
✓	✓		0.5653	0.3191	0.4019	0.3227	0.4023
✓	✓	✓	0.5953	0.3242	0.3914	0.3270	0.4095

conventional medical VLP models through three innovations: multi-aspect knowledge-image contrastive learning, fine-grained alignment, and diagnosis-guided weighting. Experiments on diverse benchmarks demonstrate our framework’s superior performance over SOTA VLP models. We hope our work inspires further research on knowledge-enhanced vision-language pretraining for medical domains.

Disclosure of Interests. The authors declare that they have no competing interests.

References

1. Daneshjou, R., Yuksekgonul, M., Cai, Z.R., Novoa, R., Zou, J.Y.: Skincon: A skin disease dataset densely annotated by domain experts for fine-grained debugging and analysis. *Advances in Neural Information Processing Systems* **35**, 18157–18167 (2022)
2. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
3. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. *nature* **542**(7639), 115–118 (2017)
4. Fan, L., Krishnan, D., Isola, P., Katabi, D., Tian, Y.: Improving clip training with language rewrites. *Advances in Neural Information Processing Systems* **36**, 35544–35575 (2023)
5. Groh, M., Harris, C., Soenksen, L., Lau, F., Han, R., Kim, A., Koochek, A., Badri, O.: Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1820–1828 (2021)
6. Han, S.S.: SNU dataset + Quiz. Preprint (3 2019). <https://doi.org/10.6084/m9.figshare.6454973.v12>
7. Hu, M., Yuan, K., Shen, Y., Tang, F., Xu, X., Zhou, L., Li, W., Chen, Y., Xu, Z., Peng, Z., et al.: Ophclip: Hierarchical retrieval-augmented learning for ophthalmic surgical video-language pretraining. *arXiv preprint arXiv:2411.15421* (2024)
8. Ikezogwo, W., Seyfioglu, S., Ghezloo, F., Geva, D., Sheikh Mohammed, F., Anand, P.K., Krishna, R., Shapiro, L.: Quilt-1m: One million image-text pairs for histopathology. *Advances in neural information processing systems* **36**, 37995–38017 (2023)

9. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghighi, B., Ball, R., Shpanskaya, K., 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, pp. 590–597 (2019)
10. Johnson, A.E., Pollard, T.J., Berkowitz, S.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Mark, R.G., Horng, S.: Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data* **6**(1), 317 (2019)
11. Kawahara, J., Daneshvar, S., Argenziano, G., Hamarneh, G.: Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *IEEE Journal of Biomedical and Health Informatics* **23**(2), 538–546 (mar 2019). <https://doi.org/10.1109/JBHI.2018.2824327>
12. Kim, C., Gadgil, S.U., DeGrave, A.J., Omiye, J.A., Cai, Z.R., Daneshjou, R., Lee, S.I.: Transparent medical image ai via an image–text foundation model grounded in medical literature. *Nature Medicine* **30**(4), 1154–1165 (2024). <https://doi.org/10.1038/s41591-024-02887-x>, <https://doi.org/10.1038/s41591-024-02887-x>
13. Lin, W., Zhao, Z., Zhang, X., Wu, C., Zhang, Y., Wang, Y., Xie, W.: Pmc-clip: Contrastive language-image pre-training using biomedical documents. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 525–536. Springer (2023)
14. Liu, Y., Jain, A., Eng, C., Way, D.H., Lee, K., Bui, P., Kanada, K., de Oliveira Marinho, G., Gallegos, J., Gabriele, S., et al.: A deep learning system for differential diagnosis of skin diseases. *Nature medicine* **26**(6), 900–908 (2020)
15. Pacheco, A.G., et al.: Pad-ufes-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones. *Data in Brief* **32**, 106221 (2020). <https://doi.org/https://doi.org/10.1016/j.dib.2020.106221>
16. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PmLR (2021)
17. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019)
18. Shen, Y., Li, H., Sun, C., Ji, H., Zhang, D., Hu, K., Tang, Y., Chen, Y., Wei, Z., Lv, J.: Optimizing skin disease diagnosis: harnessing online community data with contrastive learning and clustering techniques. *NPJ Digital Medicine* **7**(1), 28 (2024)
19. Sun, X., Yang, J., Sun, M., Wang, K.: A benchmark for automatic visual classification of clinical skin disease images. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *Computer Vision – ECCV 2016*. pp. 206–222. Springer International Publishing, Cham (2016)
20. Xie, Y., Chen, Q., Wang, S., To, M.S., Lee, I., Khoo, E.W., Hendy, K., Koh, D., Xia, Y., Wu, Q.: Pairaug: What can augmented image-text pairs do for radiology? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11652–11661 (2024)
21. Yan, S., Hu, M., Jiang, Y., Li, X., Fei, H., Tschandl, P., Kittler, H., Ge, Z.: Derm1m: A million-scale vision-language dataset aligned with clinical ontology knowledge for dermatology. *arXiv preprint arXiv:2503.14911* (2025)
22. Yan, S., Yu, Z., Primiero, C., Vico-Alonso, C., Wang, Z., Yang, L., Tschandl, P., Hu, M., Ju, L., Tan, G., et al.: A multimodal vision foundation model for clinical dermatology. *Nature Medicine* pp. 1–12 (2025)

23. Yan, S., Yu, Z., Zhang, X., Mahapatra, D., Chandra, S.S., Janda, M., Soyer, P., Ge, Z.: Towards trustable skin cancer diagnosis via rewriting model's decision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11568–11577 (2023)
24. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917 (2022)
25. Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L.: Sigmoid loss for language image pre-training. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 11975–11986 (2023)
26. Zhang, S., Xu, Y., Usuyama, N., Xu, H., Bagga, J., Tinn, R., Preston, S., Rao, R., Wei, M., Valluri, N., et al.: Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. arXiv preprint arXiv:2303.00915 (2023)
27. Zheng, K., Zhang, Y., Wu, W., Lu, F., Ma, S., Jin, X., Chen, W., Shen, Y.: Dreamlip: Language-image pre-training with long captions. In: European Conference on Computer Vision. pp. 73–90. Springer (2024)
28. Zhou, J., Sun, L., Xu, Y., Liu, W., Afvari, S., Han, Z., Song, J., Ji, Y., He, X., Gao, X.: Skincap: A multi-modal dermatology dataset annotated with rich medical captions. arXiv preprint arXiv:2405.18004 (2024)