



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Mammo-CLIP: A Vision Language Foundation Model to Enhance Data Efficiency and Robustness in Mammography

Shantanu Ghosh¹[0000-0003-4085-541X](✉), Clare B. Poynton², Shyam Visweswaran³, and Kayhan Batmanghelich¹

- ¹ Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA
shawn24@bu.edu
- ² Boston University Chobanian & Avedisian School of Medicine, Boston, MA, USA
- ³ Department of Biomedical Informatics, University of Pittsburgh, Pittsburgh, PA, USA

Abstract. The lack of large and diverse training data on Computer-Aided Diagnosis (CAD) in breast cancer detection has been one of the concerns that impedes the adoption of the system. Recently, pre-training with large-scale image text datasets via Vision-Language models (VLM) (*e.g.*, CLIP) partially addresses the issue of robustness and data efficiency in computer vision (CV). This paper proposes Mammo-CLIP, the first VLM pre-trained on a substantial amount of screening mammogram-report pairs, addressing the challenges of dataset diversity and size. Our experiments on two public datasets demonstrate strong performance in classifying and localizing various mammographic attributes crucial for breast cancer detection, showcasing data efficiency and robustness similar to CLIP in CV. We also propose Mammo-FActOR, a novel feature attribution method, to provide spatial interpretation of representation with sentence-level granularity within mammography reports. Code is available publicly: <https://github.com/batmanlab/Mammo-CLIP>.

1 Introduction

Breast cancer remains a leading cause of death among women, with a global death toll of 670K in 2022. Creating AI tools for cancer in mammograms has been a central focus for CAD [1, 18, 19]. Yet, training robust CAD models requires a large and diverse dataset, which is expensive due to the high cost of collecting large-scale annotations. VLMs (*e.g.*, CLIP [17]) pre-trained on large image-text pairs partially improve robustness in CV. This paper introduces Mammo-CLIP, the first VLM trained with mammogram-report pairs. Also, we propose Mammo-FActOR: Mammograms for Feature Atribution Connecting Observations and Reports, novel feature attribution method, enhancing interpretability by finding textual attribute-aligned visual representations.

VLMs align image-text representations in joint embedding space, providing significant benefits. Pretraining on large datasets enhances generalizability [3, 11]

and data-efficiency [8,17] for downstream tasks. The language encoder reduces reliance on costly attribute annotations for developing interpretable models [16,24]. In medical imaging, radiology reports are a source of weak labels, aiding disease localization [14,26]. Generalizability, interpretability, data efficiency, and weak supervision are crucial to improving diagnostic precision and insightful analysis in medical imaging.

Application of VLM in medical imaging has been mostly limited in chest X-ray (CXR) domain [4,22,23,25,27] due to availability of image-report dataset *e.g.*, MIMIC-CXR [6]. VLMs in other domains are mostly created with images and captions from PubMed or public forums [5,7,10]. Although fine-tuning the CLIP model on mammograms remains an option, fine-tuning reduces the image resolution to be consistent with the input dimension of CLIP. Reducing the resolution results in the loss of critical semantic visual cues essential for accurate interpretation and diagnosis. Also, identifying mammographic attributes requires the interpretation of high-resolution images, which CLIP does not offer.

Contribution. Our contributions are two fold: (1) We present Mammo-CLIP, the first VLM in mammography; (2) We introduce Mammo-FActOR, a novel feature attribution method, to align visual representations with attributes in mammography reports. Mammo-CLIP leverages a dataset of screening mammogram-report pairs. This dataset is collected by the University of Pittsburgh Medical Center (UPMC). Mammo-CLIP employs two key approaches for enhancing its learning capability: 1) a data augmentation strategy to expand the training data by incorporating external datasets containing images and mammographic attributes (*e.g.*, mass, calcification etc.) but lacking reports. We synthesize reports based on these attributes. 2) multi-view supervision (MVS) [8], to learn rich representations from limited data. Our experiments on various breast cancer detection and localization tasks, conducted across two publicly available datasets, exhibit data efficiency and robustness similar to those observed in the CV community. Although Mammo-CLIP is trained on the screening dataset, where the cancer is not explicitly mentioned, fine-tuning its representation can bootstrap data-scare classification and abnormality localization experiments in completely out-of-training distributions. The final contribution, Mammo-FActOR, successfully localizes the attributes without relying on the ground truth bounding boxes.

2 Method

The goal is to align image (\mathbf{x}^I) and text (\mathbf{x}^T) representations via image ($f^I(\cdot)$) and text encoders ($f^T(\cdot)$). Our method consists of 3 components: 1) To enhance generalizability from limited data, Mammo-CLIP is based on Multi-View Supervision (MVS) [8]. MVS leverages cross-modal self-supervision across original and augmented image-text pairs. 2) We use both instance and dataset augmentation strategies to bootstrap the generalizability further. 3) Finally, we propose Mammo-FActOR, a novel interpretable module to map the visual representation to the textual attributes. We discuss the architectures of each component in detail in the experiment section. Fig 1 illustrates our approach.

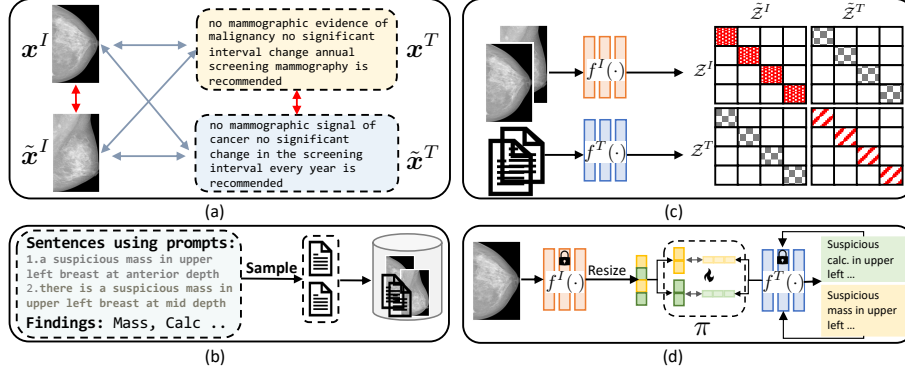


Fig. 1. Schematic view of our method. (a) Image-text augmentation for MVS. (b) Dataset augmentation by synthesizing reports using image-attribute datasets. (c) Mammo-CLIP pretraining strategy. (d) Feature attribution using Mammo-FACtoR.

Notation. Let (x_i^I, x_i^T) and $(\tilde{x}_i^I, \tilde{x}_i^T)$ be the original and augmented image-text pair of patient i respectively. Let $\mathcal{Z}^I = \{z_i^I\}_{i=1}^B$ and $\mathcal{Z}^T = \{z_i^T\}_{i=1}^B$ be the normalized representation of the image and text respectively of a batch of size B . Similarly, $\tilde{\mathcal{Z}}^I$ and $\tilde{\mathcal{Z}}^T$ denote the representation of the augmented image-text pairs. We obtain the normalized representations by projecting the outputs from Mammo-CLIP’s image and text encoders, $f^I(\cdot)$ and $f^T(\cdot)$, to the same dimensions, followed by ℓ_2 normalization.

2.1 Mammo-CLIP

Given a batch of representation pairs $\mathcal{Z}, \tilde{\mathcal{Z}}$, the following contrastive loss aligns the two representations by pulling paired representations closer while pushing unpaired ones,

$$\mathcal{L}(\mathcal{Z}, \tilde{\mathcal{Z}}) = - \sum_{z \in \mathcal{Z}} \log \frac{\exp(\langle z, \tilde{z} \rangle / \tau)}{\sum_{\tilde{z} \in \tilde{\mathcal{Z}}} \exp(\langle z, \tilde{z} \rangle / \tau)}, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ and τ are the similarity function and a learnable temperature to scale the logits, respectively. Mammo-CLIP utilizes self-supervision across every pair of representations using the following MVS loss,

$$\mathcal{L} = \sum_{\mathcal{Z}, \tilde{\mathcal{Z}} \in \{\mathcal{Z}^I, \tilde{\mathcal{Z}}^I, \mathcal{Z}^T, \tilde{\mathcal{Z}}^T\}, \mathcal{Z} \neq \tilde{\mathcal{Z}}} \mathcal{L}(\mathcal{Z}, \tilde{\mathcal{Z}}). \quad (2)$$

In our experiments, we down-weight $\mathcal{L}(\mathcal{Z}^T, \tilde{\mathcal{Z}}^T)$, the loss for the original and augmented text representation pair by half.

Table 1. Classification performance on VinDr dataset. We report Area Under the Curve (AUC) AUC for binary classification tasks to classify calcification and mass. We report accuracy for multi-class classification tasks to predict density.

Label	Model	Pre-training Data	Zero-shot (ZS)	Linear Probe (LP)		Finetune (FT)	
			100%	10%	50%	100%	100%
Calcification	CLIP w/ RN-50	UPMC	0.37	0.42	0.51	0.55	0.68
	CLIP w/ EN-B5	UPMC	0.53	0.53	0.60	0.65	0.90
	Mammo-CLIP w/ EN-B2	UPMC	0.68	0.90	0.92	0.92	0.98
	Mammo-CLIP w/ EN-B5	UPMC	0.61	<u>0.91</u>	<u>0.93</u>	<u>0.94</u>	0.98
	Mammo-CLIP w/ EN-B5	UPMC, VinDr	0.62	0.92	0.94	0.96	0.98
Mass	CLIP w/ RN-50	UPMC	0.20	0.34	0.42	0.49	0.57
	CLIP w/ EN-B5	UPMC	0.32	0.52	0.57	0.59	0.78
	Mammo-CLIP w/ EN-B2	UPMC	<u>0.58</u>	0.69	0.72	0.73	<u>0.85</u>
	Mammo-CLIP w/ EN-B5	UPMC	0.48	<u>0.73</u>	<u>0.78</u>	<u>0.79</u>	<u>0.85</u>
	Mammo-CLIP w/ EN-B5	UPMC, VinDr	0.76	0.80	0.84	0.86	0.88
Density	CLIP w/ RN-50	UPMC	0.04	0.55	0.61	0.69	0.72
	CLIP w/ EN-B5	UPMC	0.10	0.76	0.76	0.78	0.83
	Mammo-CLIP w/ EN-B2	UPMC	0.13	<u>0.80</u>	0.82	0.84	<u>0.85</u>
	Mammo-CLIP w/ EN-B5	UPMC	0.15	0.83	<u>0.84</u>	<u>0.85</u>	0.88
	Mammo-CLIP w/ EN-B5	UPMC, VinDr	0.15	0.83	0.86	0.86	0.88

2.2 Instance and Dataset Augmentation

Instance augmentation. For i^{th} patient, \mathbf{x}_i^I and $\tilde{\mathbf{x}}_i^I$ denote the CC and MLO views respectively. We also use affine and elastic transformations to extend the augmented sets. Likewise, \mathbf{x}_i^T and $\tilde{\mathbf{x}}_i^T$ denote the impression and findings sections of the mammography reports. We followed the technique proposed in [25] that uses translation between two languages to generate augmented text.

Dataset augmentation. In addition to the dataset containing paired mammograms and reports, we leverage a separate dataset containing mammograms and attributes (*e.g.*, *mass*, *calcification*) to mitigate the challenge of the limited dataset further. To synthesize reports from the attribute, we ask a board-certified radiologist to construct a set of prompts using the values of the attributes (*positive*, *negative* etc.), subtypes of the attribute (*e.g.*, *suspicious*, *obscured* etc.), laterality of the breast (*i.e.*, *right* and *left*), depth of the breast (*i.e.*, *anterior*, *mid*, *posterior*) and position of the breast (*e.g.*, *upper*, *lower* etc.). We synthesize the report by substituting the attribute values for each patient to a randomly selected prompt (see Fig. 1 in the Appendix).

2.3 Mammo-FACTOR

Attributes mentioned in the radiology report can be viewed as weak labels. The Mammon-CLIP learns to align image and text representation at the global level. To obtain a spatial interpretation of the learned representation at the granularity of the attribute, we develop Mammo-FACTOR. The Mammo-FACTOR learns to align sentences in the radiology reports containing an attribute to the channels of the frozen image encoder, $f^I(\cdot)$. Let's assume we have K different attributes and $f^I(\mathbf{x}_i^I) \in \mathbb{R}^{C \times H \times W}$ is the output of the frozen image encoder and $\mathbf{t}_i^k \in \mathbb{R}^d$ is the representation of the sentence containing the k 'th attribute

Table 2. Classification performance on RSNA dataset to classify malignancy.

Model	Pre-training Data	Zero-shot (ZS)	Finetune (FT)			Linear Probe (LP)
		100%	10%	50%	100%	100%
CLIP w/ RN-50	UPMC	0.28	0.62	0.67	0.72	0.58
CLIP w/ EN-B5	UPMC	0.42	0.80	0.84	0.86	0.70
Mammo-CLIP w/ EN-B2	UPMC	<u>0.60</u>	<u>0.82</u>	0.86	0.89	0.72
Mammo-CLIP w/ EN-B5	UPMC	0.62	0.85	<u>0.89</u>	<u>0.90</u>	<u>0.75</u>
Mammo-CLIP w/ EN-B5	UPMC, VinDr	<u>0.60</u>	0.85	0.90	0.91	0.79

in the report of the patient i . We use \mathcal{X}_-^k to denote all mammography images not containing the k 's attribute. For attribute k , the Mammo-FACTOR learns a mapping $\pi^k(\mathbf{t}_i^k, \mathbf{x}_i^I) = (\text{MLP}(f^I(\mathbf{x}_i^I); \theta_k) \mathbf{t}_i^k) \in \mathbb{R}^C$ representing the similarity of the channels with \mathbf{t}_i^k . The $\text{MLP}(\cdot; \theta_k)$ is a Multi-layer Perception parametrized by θ_k . To learn the parameters of the MLP, we minimize the contrastive loss between images with and without attributes:

$$\mathcal{L}^{\text{fac}} = - \sum_{i,c,k} \log \frac{\exp([\pi^k(\mathbf{t}_i^k, \mathbf{x}_i^I)]_c / \tau)}{\exp([\pi^k(\mathbf{t}_i^k, \mathbf{x}_i^I)]_c / \tau) + \sum_{x \in \mathcal{X}_-^k} \exp([\pi^k(\mathbf{t}_i^k, \mathbf{x})]_c / \tau)}, \quad (3)$$

where $[\cdot]_c$ denotes c 'th element of the input vector argument and τ is the temperature. After training, we weigh the representations as, $\pi^k(\cdot) f^I(\cdot) \in \mathbb{R}^{H \times W}$ to construct the textual aligned heatmap for attribute k .

3 Experiments

We aim to answer the following research questions: **RQ1.** Does Mammo-CLIP enhance zero-shot (ZS) capabilities and labeling efficiency in classification tasks? **RQ2.** Does Mammo-CLIP learn robust representations? **RQ3.** Does Mammo-CLIP boost the performance of label-efficient localization? **RQ4.** Does Mammo-FACTOR enhance interpretability through attribute-aligned representation?

3.1 Datasets

UPMC. The UPMC dataset includes 13,829 patient-report pairs, resulting in 25,355 screening mammograms (BI-RADS 0-2) with patients having at least one CC or MLO view. We use 0.8/0.2 for the train/val split. **VinDr.** The VinDr-Mammo [15] is a publicly available dataset comprising 5,000 exams (20,000 images) from Vietnam, each with four views, breast level BI-RADS assessment category (1-5), breast density category (A-D), annotating and location of mammographic attributes (*i.e.*, mass, calcifications, etc.). We use the train-test split in [15]. **RSNA.** The RSNA-Mammo⁴ is a publicly available dataset having 11913 patients with 486 cancer cases. We use train/val/test split as 0.7/0.2/0.1.

⁴ <https://www.kaggle.com/competitions/rsna-breast-cancer-detection>

Table 3. Localization performance (mAP) on VinDr dataset.

Label	Model	Pre-training Data	Finetune			Freeze Encoder
			10%	50%	100%	100%
Calcification	CLIP w/ RN-50	UPMC	0.03	0.11	0.17	0.08
	CLIP w/ EN-B5	UPMC	0.09	0.23	0.29	0.10
	Mammo-CLIP w/ EN-B2	UPMC	0.08	0.20	<u>0.32</u>	0.17
	Mammo-CLIP w/ EN-B5	UPMC	0.12	0.28	0.35	0.17
	Mammo-CLIP w/ EN-B5	UPMC, VinDr	<u>0.10</u>	<u>0.25</u>	0.35	0.17
Mass	CLIP w/ RN-50	UPMC	0.23	0.29	0.36	0.15
	CLIP w/ EN-B5	UPMC	0.34	0.45	0.49	0.28
	Mammo-CLIP w/ EN-B2	UPMC	0.38	0.50	<u>0.55</u>	<u>0.37</u>
	Mammo-CLIP w/ EN-B5	UPMC	<u>0.41</u>	<u>0.52</u>	0.58	0.35
	Mammo-CLIP w/ EN-B5	UPMC, VinDr	0.43	0.54	0.58	0.39

3.2 Experimental details

Pre-processing the images. We use a rule-based approach to find the breast’s ROI for all the datasets. We set values less than 40 to 0 and eliminate the consistently identical rows and columns, as they denote the background. It results in images with an average aspect ratio of 1:1.6 to 2, finally resized to 1520×912 .

Pre-training. We develop 2 variant of Mammo-CLIP models: one pre-trained with the UPMC dataset only and the other using both the UPMC and VinDr datasets. Mammo-CLIP utilizes BioClinicalBERT [2] as the text encoder and EfficientNet(EN)-B2 and B5 [20] with ImageNet pre-trained weights as the image encoder. Image augmentations consist of 1) affine transformations with rotations up to 20° , a minimum translation of 0.1%, scaling factors [0.8, 1.2], and shearing by 20° ; and 2) elastic transformations with ($\alpha = 10$, $\sigma = 5$). Text augmentations include sentence swaps and back-translation from Italian to English⁵. We train Mammo-CLIP for 10 epochs in a mixed-precision manner and optimize by AdamW [13] with an initial learning rate $5e-5$ and a weight decay $1e-4$. We use a cosine-annealing learning-rate scheduler [12] with a warm-up for 1 epoch. **Baseline.** Using UPMC dataset and CLIP objective [17], we construct two baselines: 1) an image encoder w/ ResNet (RN)-50 initialized with CLIP weights and fine-tuned with 224×224 images, 2) EfficientNet (EN)-B5 fine-tuned using the same pre-processed images as Mammo-CLIP. Both the baselines are pre-trained using the UPMC dataset, as CLIP only uses an image-text dataset, not an image-attribute dataset. **Mammo-FACTOR.** Our Mammo-Factor is the same light-weight neural network [21] with linear layer \rightarrow ReLU \rightarrow linear layer. We train it with a learning rate 0.0001 and temperature (τ) of 0.007 for 20 epochs.

Downstream evaluation. Evaluation of the downstream tasks utilizes both the Mammo-CLIP’s image representations to 1) classify mass, calcification, and density in the VinDr dataset and cancer in the RSNA dataset; 2) localize mass and calcification in the VinDr dataset on a held-out test set. We evaluate the classification tasks in 3 settings: zero-shot (ZS), linear probe (LP), and fine-tuning

⁵ <https://huggingface.co/Helsinki-NLP>

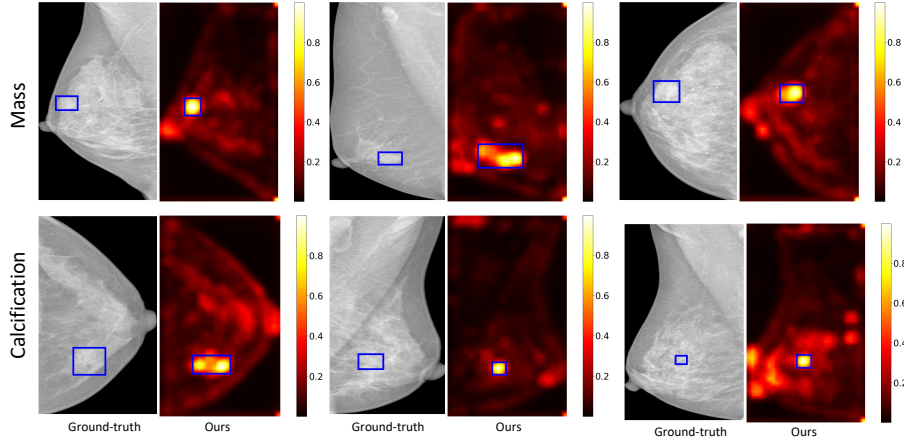


Fig. 2. Mammo-FActoR localizes mass and calcification w/o the ground-truth bboxes from the VinDr dataset. For each pair, the left image denotes the original image with the ground-truth bbox, while the right one is the bbox predicted by Mammo-Factor.

(FT). In ZS and LP, we freeze both the encoders of Mammo-CLIP. For FT, we jointly fine-tune the image encoder of Mammo-CLIP and the linear classifier. For zero-shot (ZS) evaluation, we use prompts as $\{\text{No } \langle E \rangle, \langle E \rangle\}$, where E denotes mass, calcification, and cancer. Refer to Tab. 1 in the Appendix for prompts for density. LP on VinDr utilizes varying amounts of the training data (10%, 50%, or 100%). For FT on VinDr, we use 100% data. As the screening mammogram reports do not mention malignancy explicitly, data-efficient cancer classification for the RSNA dataset involves fine-tuning with varying proportions of the training data (10%, 50%, or 100%). We also report LP for the RSNA dataset using the entire training set. The localization task includes fine-tuning the image encoder of Mammo-CLIP using RetinaNet [9] ($\alpha = 0.6$, $\gamma = 2.0$) model and 10%, 50%, and 100% of training data from VinDR. We also report the localization performance of RetinaNet by freezing Mammo-CLIP’s image encoder with 100% of training data. Both setups in localization use only VinDr’s suspicious findings. **Metric.** We report AUC scores for binary classification tasks to predict mass, calcification, and cancer and accuracy for multi-class classification tasks to predict density. For localization, we report mAP with IoU=0.5 and detector confidence thresholded to 0.05 [9] unless specified.

4 Results

Classification (RQ1 & RQ2). Tab. 1 and 2 illustrate the classification results for VinDr and RSNA, respectively. The Mammo-CLIP w/ EN-B5, pre-trained on the UPMC and VinDr datasets, outperforms all the other models in most of the evaluation settings. For VinDr, Mammo-CLIP’s LP efficacy with 10%

Table 4. Weakly supervised localization on VinDr using Mammo-FActOR.

Model	Mass		Calcification	
	IoU@0.25	IoU@0.50	IoU@0.25	IoU@0.50
Mammo-CLIP w/ EN-B2	0.38	0.22	0.2	0.13
Mammo-CLIP w/ EN-B5	0.45	0.37	0.25	0.17

of data surpasses the FT performance of the baselines with 100% data. Also, Mammo-CLIP excels in cancer classification on the RSNA dataset, showing its robustness, even when cancer is not explicitly mentioned in the reports.

CLIP w/ RN-50 backbone fails to capture the visual cues from low-resolution images. CLIP w/ EN-B5 addresses this by employing high-resolution images but struggles to learn rich representation primarily due to its optimization strategy focusing solely on the contrastive loss of the original image-text pairs. High-resolution images and MVS optimization strategy allow Mammo-CLIP w/ EN backbones to extract richer representations from limited data than baselines, improving the classification performance. Recall that the UPMC dataset consists of mammograms from a different population than the VinDr dataset. The image encoder from Mammo-CLIP pre-trained with the UPMC dataset shows excellent performance on both LP and FT settings on the VinDR dataset, showcasing the robustness of the learned representations. ZS results on some tasks (especially density) exhibit the need for more data.

Localization (RQ2 & RQ3). Tab. 3 shows localization results in a supervised setting, utilizing the bounding boxes of VinDr. Mammo-CLIP w/EN-B5 pre-trained with the UPMC and VinDr dataset outperforms other models in most settings. Refer to the Classification subsection (described above) for insights into such improvement. Also, the utilization of VinDr during the pre-training improves the performance of both downstream tasks.

Evaluating Mammo-FActOR (RQ4). We evaluate the Mammo-FActOR using a weakly supervised localization task to localize the mass and calcification of the VinDr dataset without using the ground truth bounding boxes (bboxes). We construct heatmaps using the Mammo-FActOR for both attributes. Following [26], we extract isolated regions from the heatmap where pixel values are greater than the 95% quantile of the heatmap’s pixel value distribution and generate bboxes. We estimate mAP by comparing the IoU between generated and ground truth bounding boxes. A generated box is true positive when $\text{IoU} > T(\text{IoU})$, where $T(\cdot)$ is a threshold. Tab. 4 demonstrates the localization results. Mammo-CLIP w/ EN-B5 for $T(\text{IoU}) = 0.25$ outperforms the other variants for both the attributes. Fig. 2 shows the qualitative results. Mammo-FActOR aligns image representations with attributes-specific textual descriptions from reports. Thus, it successfully localizes the attributes without relying on ground-truth bounding boxes during finetuning. Also, Fig. 2 in the Appendix shows the top 3 feature maps by Mammo-FActOR for detecting these attributes.

5 Conclusion

This paper proposes a Mammo-CLIP and Mammo-FActOR to learn a cross-modal embedding, improving generalizability, data efficiency, and interpretability. Future directions include further leveraging vision transformers and cross-attention mechanisms to enhance Mammo-CLIP and Mammo-FActOR.

Acknowledgments. This work was partially supported by the Pennsylvania Department of Health, NIH Award Number 1R01HL141813-01, and funding from the Hariri Institute for Computing, Boston University. We are grateful for the computational resources from Pittsburgh Super Computing grant number TG-ASC170024.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alberdi, E., Povyakalo, A., Strigini, L., Ayton, P., Hartswood, M., Procter, R., Slack, R.: Use of computer-aided detection (cad) tools in screening mammography: a multidisciplinary investigation. *The British journal of radiology* **78**(suppl_1), S31–S40 (2005)
2. Alsentzer, E., Murphy, J.R., Boag, W., Weng, W.H., Jin, D., Naumann, T., McDermott, M.: Publicly available clinical bert embeddings. *arXiv preprint arXiv:1904.03323* (2019)
3. Eslami, S., de Melo, G., Meinel, C.: Does clip benefit visual question answering in the medical domain as much as it does in the general domain? *arXiv preprint arXiv:2112.13906* (2021)
4. Huang, S.C., Shen, L., Lungren, M.P., Yeung, S.: Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3942–3951 (2021)
5. Huang, Z., Bianchi, F., Yuksekgonul, M., Montine, T.J., Zou, J.: A visual–language foundation model for pathology image analysis using medical twitter. *Nature medicine* **29**(9), 2307–2316 (2023)
6. Johnson, A., Lungren, M., Peng, Y., Lu, Z., Mark, R., Berkowitz, S., Horng, S.: Mimic-cxr-jpg-chest radiographs with structured labels. *PhysioNet* (2019)
7. Kim, C., Gadgil, S.U., DeGrave, A.J., Cai, Z.R., Daneshjou, R., Lee, S.I.: Fostering transparent medical image ai via an image-text foundation model grounded in medical literature. *medRxiv* (2023)
8. Li, Y., Liang, F., Zhao, L., Cui, Y., Ouyang, W., Shao, J., Yu, F., Yan, J.: Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In: *International Conference on Learning Representations* (2022), <https://openreview.net/forum?id=zq1iJkNk3uN>
9. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2980–2988 (2017)
10. Lin, W., Zhao, Z., Zhang, X., Wu, C., Zhang, Y., Wang, Y., Xie, W.: Pmc-clip: Contrastive language-image pre-training using biomedical documents. *arXiv preprint arXiv:2303.07240* (2023)

11. Liu, J., Zhang, Y., Chen, J.N., Xiao, J., Lu, Y., A Landman, B., Yuan, Y., Yuille, A., Tang, Y., Zhou, Z.: Clip-driven universal model for organ segmentation and tumor detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21152–21164 (2023)
12. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
13. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
14. Müller, P., Meissen, F., Brandt, J., Kaissis, G., Rueckert, D.: Anatomy-driven pathology detection on chest x-rays. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 57–66. Springer (2023)
15. Nguyen, H.T., Nguyen, H.Q., Pham, H.H., Lam, K., Le, L.T., Dao, M., Vu, V.: Vindr-mammo: A large-scale benchmark dataset for computer-aided diagnosis in full-field digital mammography. *Scientific Data* **10**(1), 277 (2023)
16. Oikarinen, T., Das, S., Nguyen, L.M., Weng, T.W.: Label-free concept bottleneck models. arXiv preprint arXiv:2304.06129 (2023)
17. 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)
18. Srivastava, S., Sharma, N., Singh, S.K., Srivastava, R.: Design, analysis and classifier evaluation for a cad tool for breast cancer detection from digital mammograms. *International journal of Biomedical Engineering and technology* **13**(3), 270–300 (2013)
19. Srivastava, S., Sharma, N., Singh, S., Srivastava, R.: Quantitative analysis of a general framework of a cad tool for breast cancer detection from mammograms. *Journal of Medical Imaging and Health Informatics* **4**(5), 654–674 (2014)
20. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
21. Varma, M., Delbrouck, J.B., Hooper, S., Chaudhari, A., Langlotz, C.: Villa: Fine-grained vision-language representation learning from real-world data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22225–22235 (2023)
22. Wang, Z., Wu, Z., Agarwal, D., Sun, J.: Medclip: Contrastive learning from unpaired medical images and text. arXiv preprint arXiv:2210.10163 (2022)
23. Wu, C., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: Medklip: Medical knowledge enhanced language-image pre-training. Proceedings of the IEEE/CVF International Conference on Computer Vision (2023)
24. 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. pp. 19187–19197 (2023)
25. You, K., Gu, J., Ham, J., Park, B., Kim, J., Hong, E.K., Baek, W., Roh, B.: Cxr-clip: Toward large scale chest x-ray language-image pre-training. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 101–111. Springer (2023)
26. Yu, K., Ghosh, S., Liu, Z., Deible, C., Batmanghelich, K.: Anatomy-guided weakly-supervised abnormality localization in chest x-rays. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 658–668. Springer (2022)

27. Zhang, Y., Jiang, H., Miura, Y., Manning, C.D., Langlotz, C.P.: Contrastive learning of medical visual representations from paired images and text. In: Machine Learning for Healthcare Conference. pp. 2–25. PMLR (2022)