

CLEFT: Language-Image Contrastive Learning with Efficient Large Language Model and Prompt Fine-Tuning

Yuexi Du¹, Brian Chang¹, Nicha C. Dvornek^{1,2}

¹ Department of Biomedical Engineering,

² Department of Radiology & Biomedical Imaging,
Yale University, New Haven, CT, USA

{yuexi.du, brian.chang, nicha.dvornek}@yale.edu

Abstract. Recent advancements in Contrastive Language-Image Pre-training (CLIP) [21] have demonstrated notable success in self-supervised representation learning across various tasks. However, the existing CLIP-like approaches often demand extensive GPU resources and prolonged training times due to the considerable size of the model and dataset, making them poor for medical applications, in which large datasets are not always common. Meanwhile, the language model prompts are mainly manually derived from labels tied to images, potentially overlooking the richness of information within training samples. We introduce a novel language-image Contrastive Learning method with an Efficient large language model and prompt Fine-Tuning (CLEFT) that harnesses the strengths of the extensive pre-trained language and visual models. Furthermore, we present an efficient strategy for learning context-based prompts that mitigates the gap between informative clinical diagnostic data and simple class labels. Our method demonstrates state-of-the-art performance on multiple chest X-ray and mammography datasets compared with various baselines. The proposed parameter efficient framework can reduce the total trainable model size by 39% and reduce the trainable language model to only 4% compared with the current BERT encoder.³

Keywords: Deep Learning · Multi-Modal · Contrastive Learning · Chest X-ray · Mammography

1 Introduction

Contrastive learning [9] has emerged as a pivotal paradigm in deep learning due to its ability to construct a robust feature space and generalize well to downstream tasks. Traditional contrastive learning methods [9,6,5,20] focus on building positive pairs from different views of the same input while distinguishing it from other data in the feature space. Such a contrastive paradigm allows the

³ The official implementation is available at <https://github.com/XYPB/CLEFT>.




Fig. 1. Zero-shot Performance on CheXpert-5x200. We compare the performance of our method with multiple baselines on zero-shot CheXpert-5x200 classification. For each model, the diameter denotes the total number of trainable parameters, the color shows the number of trainable text encoder parameters, and the number reports accuracy. Our method outperforms all baselines with better parameter efficiency.

model to learn a robust representation when an exhaustive amount of training data is provided. Recently, Contrastive Language-Image Pre-training (CLIP) [21], which utilizes both visual and textual data insights, has extended contrastive learning to multi-modality data and reaped immense benefits from advancements in language models [22,3,25]. CLIP supervises the image encoder with a language encoder trained simultaneously. By contrasting the features learned from both the image and language encoder, the CLIP model learns to bridge the textual and visual data in the high-dimensional embedding space, allowing it to use knowledge learned from the language model to guide visual encoder training. A properly pre-trained contrastive visual encoder can be adapted to multiple downstream tasks with a minimum amount of labeled data required.

However, difficulties arise when adapting CLIP from the natural image-text pair to the medical domain, where access to data can be severely restricted by various factors, including security and privacy concerns, difficulty in obtaining expert annotations, and expensive imaging. The limited image-text pairs in the medical domain constrain the potential of CLIP models trained from scratch. While existing medical CLIP methods [12,28,26,27] use a pre-trained BERT language model [1], its limited model size constrained the expression ability in the embedding space and further constrained the pre-training capability. Also, the common approach of handcrafting textual prompts for CLIP [21,12,27] leads to a lack of diversity in text training prompts, which can result in the catastrophic forgetting phenomenon [19] in the text encoder and limit the model’s performance.

In light of these challenges, we introduce a novel language-image Contrastive Learning method with Efficient LLM and prompt context Fine-Tuning (CLEFT) to boost overall performance. We leverage the strengths of vast pre-trained large language models (LLMs) and visual models, adapting them to the medical domain to counterbalance the scarcity of medical data and address the constraints due to language model size. Our work is the first in the realm of medical imaging to scale up the language model in the language-image pre-training to the billion parameter




Fig. 2. Proposed Method Framework. (a) Language-image contrastive learning with an LLM by utilizing PEFT. Fixed handcrafted prompts are used in this stage. (b) Prompt context learning with the pre-trained image and text encoder via classification.

level, leading to greatly improved model generalization ability. Further, our approach is parameter efficient, solely focusing on optimizing smaller adaptation layers [11,15] in the LLM, thus conserving GPU resources without compromising knowledge acquired from natural language. Our model reduces the total number of trainable parameters by 39% with only 4% of the trainable language model parameters compared to the vanilla CLIP [22] (Fig. 1). In addition, our method learns a context-based prompt with prompt fine-tuning [30] to mitigate the bias introduced by the undiversified hand-crafted prompt and generalize to different unseen data. The proposed method shows state-of-the-art performance on two public chest X-ray datasets [13,24] and a public mammography dataset [14].

2 Methods

The proposed CLEFT framework is in Fig. 2. We first efficiently incorporate an LLM into the CLIP [21] framework. We then train the learnable prompt context with frozen pre-trained text and visual encoders to further improve generalization.

2.1 Boosting CLIP with an LLM

Contrastive Language-Image Pre-training. The conventional CLIP [21] framework includes a vision encoder and a text encoder with the corresponding projection head that encodes the image-text pair (x_{I_i}, x_{T_i}) sampled from the training data to feature (I_i, T_i) (Fig. 2(a)). The projection head maps the embedding from two different modalities into the same feature space and therefore allows the model to bridge two modalities. The multi-modal contrastive learning

optimizes both image-to-text and text-to-image InfoNCE [9] loss symmetrically:

$$\mathcal{L} = -\frac{1}{2N} \left[\underbrace{\sum_{i=0}^N \log \frac{\exp(I_i \cdot T_i / \tau)}{\sum_{j \neq i}^N \exp(I_i \cdot T_j / \tau)}}_{\mathcal{L}_{I2T}} + \underbrace{\sum_{i=0}^N \log \frac{\exp(T_i \cdot I_i / \tau)}{\sum_{j \neq i}^N \exp(T_i \cdot I_j / \tau)}}_{\mathcal{L}_{T2I}} \right], \quad (1)$$

where N is the number of samples within a batch, and τ is the learnable softmax temperature. Optimizing the loss will reduce the cosine distance between paired image and text features while repelling unpaired samples. However, proper contrastive learning requires a large amount of training data as well as a large negative sample size (*i.e.*, batch size), and it is hard to gather large datasets in the medical domain and GPU resources are limited. Thus, we propose to replace the vanilla encoders trained from scratch with fully pre-trained models to alleviate this issue. Initializing the text encoder with a pre-trained model provides higher-quality supervision since the pre-trained model already provides a robust feature space [4], which simplifies the contrastive learning procedure.

PEFT LLM as Text Encoder. To further explore the potential of the CLIP framework, we take advantage of a medical LLM that was pre-trained with a large amount of text data. We use a GPT-2-based causal language model as our text encoder rather than a BERT-based language model, which was widely used in previous medical CLIP models [28,12,26,27], since the causal LLM has shown a better capability as it scales up to over a billion parameters [3]. A stronger text encoder allows the model to embed the input into a more robust feature space with less training. However, LLMs are more likely to overfit on undiversified text data given their strong expression ability. To avoid this issue and maintain the robust pre-trained knowledge within the LLM efficiently, we introduce the parameter-efficient fine-tuning (PEFT) module to the frozen LLM (Fig. 2(a)), where a small set of trainable parameters are injected into each transformer block, adjusting the original output of the attention layers slightly. This reduces the number of trainable parameters during training to no more than 1% of the full large language model size. Common PEFT methods like LoRA [11] and IA3 [16] either adjust the attention output with low-rank bottleneck matrices or scale the key, query, and value outputs. Also, prefix fine-tuning [15] introduces extra trainable prefix tokens to the language model to influence the behavior of the attention layers. The nature of these PEFT methods ensures the fine-tuned output will not deviate too much from the original model and helps avoid the catastrophic forgetting phenomenon [19]. To further merge the domain gap between pre-training text data and CLIP prompt, we unlock the LLM’s embedding layer and update the corresponding token embedding during pre-training.

CLIP as Knowledge Distillation. We further argue that contrastive language-image learning with a pre-trained text encoder can be viewed as a knowledge distillation process, where the numerator in Eq. (1) minimizes the distance

between I_i and T_i . Given that the language encoder is largely frozen, the language encoder serves more as a teacher model since it already has a well-established embedding space that can distinguish between different samples, and its output only changes slightly during training. The fully optimized visual encoder then focuses on aligning its output to the language model’s embedding space, which serves as the student model. This allows us to distill the knowledge within the well-pre-trained language model into a generally much smaller vision model and, therefore, further improve its performance. Additionally, since the “teacher model” is still optimized during pre-training, the negative pairs are still necessary.

Model Architecture. We choose GPT-2 [22] with 32 causal transformer blocks as the text encoder and the ViT-Base [8] with a patch size of 14 as the visual encoder (Fig. 2(a)). Similar to the original CLIP [21], we use a randomly initialized linear projection layer to map the embeddings from each encoder to a unified embedding space with the same size. For the text encoder, we use the output embedding of the first [EOS] token since the model is causal and this token encodes all the information from the input. For the PEFT module, we experiment with LoRA [11], IA3 [16], and prefix fine-tuning [15]. For the visual encoder, we use the averaged embedding of all visual tokens from ViT [8] as the visual embedding. We remove the last layer norm for better training stability.

2.2 Learning the Context-Based Prompt

To further address the issue of the lack of diversity in the hand-crafted prompts, we introduce a second stage of training that only optimizes a learnable context-based token of length L (Fig. 2(b)). After pre-training, we freeze both encoders and replace the original hand-crafted prompt with a series of trainable tokens that then feed into the language model. The same context-based prompt tokens are used from all classes, which ensures the generalization ability of these tokens. Different from the pre-training stage, we optimize the trainable context tokens with a zero-shot classification cross-entropy loss. This allows the prompt tokens to adapt to different classes evenly and avoid the potential shortcut issue. We further initialize these tokens with the embedding of the original hand-crafted caption. If L is longer than the original caption, we instead initialize the first few tokens according to the random uniform distribution.

3 Experiments

3.1 Datasets

We evaluate our CLEFT model on two major applications in medical imaging, chest X-ray and mammography. We use the **CheXpert-1.0** [13] for pre-training following GLORIA [12]. The dataset has 223,415 images from 65,240 patients with corresponding class labels for 14 different classes. We only use frontal chest radiographs for consistency. We leave out 5% of data for validation. For

evaluation, we use the in-domain **CheXpert-5x200** [12] and out-of-domain **RSNA** [24] datasets. The CheXpert-5x200 [12] is a multi-class classification subset of CheXpert-1.0 with 5 different classes and 200 images in each class. These images are removed from the training set. The RSNA [24] provides a collection of pneumonia and non-pneumonia images for binary classification. Following MedCLIP [27], we sample a 1:1 subset with 8,486 training and 3,538 testing images for this dataset. For the mammography data, we use the **EMBED** [14] dataset which contains 364,515 2D mammograms with both BI-RADS [23] and breast density labels. We split the data into 70%/10% for training/validation. We evaluate our model on both BI-RADS and density prediction tasks with two balanced subsets sampled from the 20% remaining data. The BI-RADS classification test set contains 7 classes with 200 samples per class. The density classification test set contains 4 classes each with 500 samples.

3.2 Baselines and Evaluation Metrics

We compare multiple state-of-the-art baselines. To demonstrate the effectiveness of the CLIP pre-training, we compare with the same **ViT** [20] model with random initialization and Image-Net [7] pre-training. We further compare our model with conventional **CLIP** [21], ResNet50 [10] based medical CLIP method **ConVIRT** [28] and **GLoRIA** [12]. We also compare with recent medical CLIP baselines including **MGCA** [26], **MRM** [29], and **MedCLIP** [27] with Swin-Transformer [17]. We choose these baselines as they provide either their pre-trained model or full training code. However, the ConVIRT [28] model does not provide a pre-trained model, so we report results directly drawn from Wang *et al.* [27].

We evaluate all models under zero-shot, linear-probing, and full fine-tuning settings. We report accuracy for zero-shot classification and both accuracy and area under the receiver operating characteristic curve (AUC) for the two fine-tuning settings. We further evaluate the data efficiency during full fine-tuning of our model and compare the model size.

3.3 Implementation Details

We choose BioMedLM-3B [2] and DiNOv2 [20] to initialize our encoders. During the pre-training, we use a batch size of 72 and learning rate of 4×10^{-5} for 40,000 steps. We use the cosine annealing scheduler with a 4,000-step linear warm-up and AdamW [18] optimizer with weight decay of 0.2. We select the model with the smallest validation loss as the final model. During prompt context learning, we use a batch size of 36 and a learning rate of 1×10^{-3} for 4,000 steps. Prompt context length is set to $L = 30$. We use the same scheduler with a 1,000-step warm-up and SGD optimizer. For the PEFT strategy, we experiment with LoRA [11], IA3 [16] and Prefix-tuning [15]. For linear probing and full fine-tuning, we optimize cross-entropy loss using a batch size of 36 and learning rate of 5×10^{-4} for 8,000 steps with weight decay of 1×10^{-3} . All models are trained with BFloat-16-mix precision with 2 NVIDIA A5000 GPUs using PyTorch.

Table 1. Main Evaluation Results. We evaluate the proposed method with multiple baselines on the CheXpert-5×200 [13] and RSNA [24] datasets. We evaluate zero-shot (ZS) classification, linear-probing (LP), and full-finetuning (FT) tasks. All values are percentages. We highlight the top result in bold and the second-best with an underline.

Method	CheXpert 5 × 200 [13]					RSNA [24]				
	ZS-Acc	LP-Acc	LP-AUC	FT-Acc	FT-AUC	ZS-Acc	LP-Acc	LP-AUC	FT-Acc	FT-AUC
Random-ViT [20]	-	20.62	57.85	20.32	63.68	-	65.15	72.22	72.70	79.89
ImageNet-ViT [20]	-	38.54	75.44	56.46	85.71	-	74.96	83.72	77.44	85.24
CLIP-ViT [21]	20.16	35.34	65.14	44.84	77.59	49.89	69.47	75.90	77.08	83.53
ConVIRT-R50 [28]	41.88*	47.70*	-	-	-	47.31*	78.46*	-	-	-
GLoRIA-R50 [12]	43.28*	51.65	83.15	57.56	87.11	33.06*	76.77	86.53	78.55	87.15
MGCA-ViT [†] [26]	51.35	26.63	63.94	56.96	86.33	<u>69.90</u>	75.35	85.63	79.79	88.11
MRM-ViT [29]	-	58.26	84.48	56.56	87.41	-	76.43	85.48	78.77	86.63
MedCLIP-Swin [†] [27]	59.42*	54.55	85.75	57.46	<u>87.85</u>	74.47*	<u>79.26</u>	88.26	78.80	87.36
Ours-Prefix	60.55	<u>61.56</u>	<u>87.20</u>	<u>61.16</u>	87.73	64.07	78.75	<u>88.30</u>	79.34	88.52
Ours-IA3	<u>65.61</u>	<u>60.16</u>	<u>86.48</u>	<u>61.06</u>	86.81	64.08	78.80	<u>87.74</u>	<u>79.99</u>	<u>88.59</u>
Ours-LoRA	66.74	63.46	87.76	63.96	88.22	64.93	79.40	88.34	80.36	88.72

* Result directly drawn from Wang *et al.* [27]

[†] Method pre-trained with a different dataset with 2× greater size

3.4 Main Results

Zero-shot Classification. As shown in Tab. 1, our model with LoRA [11] outperforms other baselines with 7% improvement on the CheXpert-5x200 [13] dataset and with the smallest number of trainable language parameters and overall small trainable model size (Fig. 1, Tab. S3). Note that while our model falls behind MGCA [26] and MedCLIP [27] on the RSNA [24] evaluation, these two baselines were pre-trained with a different, larger dataset with twice the size of CheXpert-1.0 [13]. Compared with the other baselines that were pre-trained with the same data, our method performs best on the out-of-domain RSNA data.

Linear probing. Under the linear probing condition, our model with LoRA [11] achieves the best performance on both datasets (Tab. 1). We highlight the 5% gap in the CheXpert-5x200 [13] experiment. This indicates our model has a more robust embedding space that can distinguish input data even without task-specific fine-tuning. We further note that with linear probing, our method now surpasses MGCA [26] and MedCLIP [27] on RSNA, even with less pre-training data.


Full Fine-tuning. Our model outperforms all other baselines when the model is fully fine-tuned (Tab. 1). Our model shows impressive improvement on the CheXpert [13] dataset and also beats other baselines on the RSNA [24] dataset. We suggest it is the proposed PEFT language model that provides a better quality of supervision. Meanwhile, the vastly pre-trained encoders allow the model to properly adapt to out-of-domain tasks. We further evaluate each model’s data efficiency with different ratios of training data in Tab. S1 and Tab. S2. A more robust pre-trained model should be able to generalize easily to the target task even with a small amount of training data. Notably, our model also outperforms

Table 2. EMBED Evaluation. All values are percentages.

Method	BI-RADS			Density		
	ZS-Acc	LP-Acc	FT-Acc	ZS-Acc	LP-Acc	FT-Acc
Random [20]	-	17.86	17.43	-	34.65	54.45
ImageNet [20]	-	18.29	20.21	-	56.95	67.45
CLIP [21]	20.93	16.57	21.21	70.45	36.70	50.50
Ours-LoRA	33.29	30.86	23.79	74.40	74.80	74.95

Table 3. Ablation Evaluation. All values are percentages.

Method	LM #Param.	Zero-shot Acc.	
		CheXpert 5 × 200 [13]	RSNA [24]
Ours-LoRA	2.6M	66.74 (± 0.00)	64.63 (± 0.00)
- Prompt FT.	2.6M	63.46 (-3.28)	57.15 (-7.48)
- Freeze LM	0.0M	64.44 (-2.23)	65.51 ($+0.88$)
+ Full FT. LM	2.7B	67.03 ($+0.29$)	67.58 ($+2.95$)

**Fig. 3. Zero-shot Accuracy vs. Number of Trainable Prompt Tokens.**

the baselines even with much less training data. We highlight that for RSNA, our accuracy drops by $< 3\%$ when using 1% compared to 100% of the training data.

Mammography Evaluation. We report the performance of our model with LoRA [11] and baselines on the EMBED mammography data in Tab. 2. Our model clearly surpasses the compared baselines with a considerable gap. These initial results on the mammography dataset suggest the proposed model has the potential to be applied to other medical domains.

3.5 Ablation Experiments

Results of ablation experiments with our method are presented in Tab. 3. Without the second stage prompt fine-tuning, the accuracy drops by $\sim 3\%$ for both datasets. Using a fully frozen language model harms performance on in-domain data but improves performance with out-of-domain data. The fully fine-tuned model improves the performance even with a much smaller batch size; however, this improvement comes with the cost of ~ 4 times more GPU memory cost and only 1/20 batch size with 2 times longer training. We also evaluate the influence of using different lengths for the prompt context (Fig. 3), and we choose $L = 30$ as our best model in the other evaluations. Increasing the prompt length does not always improve the performance according to this experiment.

4 Discussion and Conclusion

We propose a novel contrastive language-image pre-training framework for medical images, called CLEFT, which incorporates a pre-trained medical LLM, parameter-efficient fine-tuning for the LLM, and prompt context learning. The proposed

method successfully outperforms multiple baselines in the experiments conducted on three public datasets. Furthermore, the proposed method greatly reduces the training cost of incorporating an LLM while maintaining robust performance. We provide a new view that treats the existing CLIP task as knowledge distillation from the LLM. The proposed framework can be easily merged with other related CLIP methods like MedCLIP [27] to further improve the performance.

While the proposed method demonstrated enhanced performance, additional methodological improvements should be explored, *e.g.*, different initialization methods, more robust prompt finetuning methods that can adapt to the medical image domain, and generating more diverse and detailed prompting during training. Finally, more experiments including text-image retrieval and benchmarking of different visual encoders and LLM models should be conducted to further evaluate the methods.

Acknowledgments This work was supported by NIH grant R21EB032950.

Disclosure of Interests The authors have no competing interests in this work and other related research.

References

1. Alsentzer, E., Murphy, J., Boag, W., Weng, W.H., Jin, D., Naumann, T., McDermott, M.: Publicly available clinical BERT embeddings. In: Proceedings of the 2nd Clinical Natural Language Processing Workshop. pp. 72–78. Association for Computational Linguistics, Minneapolis, Minnesota, USA (Jun 2019). <https://doi.org/10.18653/v1/W19-1909>, <https://www.aclweb.org/anthology/W19-1909>
2. Bolton, E., Hall, D., Yasunaga, M., Lee, T., Manning, C., Liang, P.: Biomedlm, <https://huggingface.co/stanford-crfm/BioMedLM>, accessed: 2023-03-02
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Nee-lakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
4. Chen, C., Zhong, A., Wu, D., Luo, J., Li, Q.: Contrastive masked image-text modeling for medical visual representation learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 493–503. Springer (2023)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
6. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297* (2020)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
8. 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)

9. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. arXiv preprint arXiv:1911.05722 (2019)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
11. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021)
12. 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)
13. 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. 33-01, pp. 590–597 (2019)
14. Jeong, J.J., Vey, B.L., Bhimireddy, A., Kim, T., Santos, T., Correa, R., Dutt, R., Mosunjac, M., Oprea-Ilie, G., Smith, G., et al.: The emory breast imaging dataset (embed): A racially diverse, granular dataset of 3.4 million screening and diagnostic mammographic images. *Radiology: Artificial Intelligence* **5**(1), e220047 (2023)
15. Li, X.L., Liang, P.: Prefix-tuning: Optimizing continuous prompts for generation. arXiv preprint arXiv:2101.00190 (2021)
16. Liu, H., Tam, D., Muqeeth, M., Mohta, J., Huang, T., Bansal, M., Raffel, C.: Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. arXiv preprint arXiv:2205.05638 (2022)
17. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
18. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
19. Luo, Y., Yang, Z., Meng, F., Li, Y., Zhou, J., Zhang, Y.: An empirical study of catastrophic forgetting in large language models during continual fine-tuning. arXiv preprint arXiv:2308.08747 (2023)
20. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Howes, R., Huang, P.Y., Xu, H., Sharma, V., Li, S.W., Galuba, W., Rabbat, M., Assran, M., Ballas, N., Synnaeve, G., Misra, I., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision (2023)
21. 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)
22. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners. OpenAI blog (2019)
23. Sickles, E.A., D’Orsi, C.J., Bassett, L.W., et al.: ACR BI-RADS mammography. In: ACR BI-RADS Atlas, Breast Imaging Reporting and Data System. American College of Radiology, Reston, VA, 5th edn. (2013)
24. Stein, A., et.al., T.X.: Rsna pneumonia detection challenge (2018), <https://kaggle.com/competitions/rsna-pneumonia-detection-challenge>

25. Sun, Q., Fang, Y., Wu, L., Wang, X., Cao, Y.: Eva-clip: Improved training techniques for clip at scale. arXiv preprint arXiv:2303.15389 (2023)
26. Wang, F., Zhou, Y., Wang, S., Vardhanabhuti, V., Yu, L.: Multi-granularity cross-modal alignment for generalized medical visual representation learning. *Advances in Neural Information Processing Systems* **35**, 33536–33549 (2022)
27. Wang, Z., Wu, Z., Agarwal, D., Sun, J.: Medclip: Contrastive learning from unpaired medical images and text. arXiv preprint arXiv:2210.10163 (2022)
28. 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)
29. Zhou, H.Y., Lian, C., Wang, L., Yu, Y.: Advancing radiograph representation learning with masked record modeling. arXiv preprint arXiv:2301.13155 (2023)
30. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *International Journal of Computer Vision* **130**(9), 2337–2348 (2022)