

Enhancing Biomedical Multi-modal Representation Learning with Multi-scale Pre-training and Perturbed Report Discrimination

Xinliu Zhong

*Department of Biomedical Engineering
National University of Singapore
Singapore
xinliuzhong@u.nus.edu*

Kayhan Batmanghelich

*Department of Electrical and Computer Engineering
Boston University
Boston, USA
batman@bu.edu*

Li Sun

*Department of Electrical and Computer Engineering
Boston University
Boston, USA
lisun@bu.edu*

Abstract—Vision-language models pre-trained on large scale of unlabeled biomedical images and associated reports learn generalizable semantic representations. These multi-modal representations can benefit various downstream tasks in the biomedical domain. Contrastive learning is widely used to pre-train vision-language models for general natural images and associated captions. Despite its popularity, we found biomedical texts have complex and domain-specific semantics that are often neglected by common contrastive methods. To address this issue, we propose a novel method, perturbed report discrimination, for pre-train biomedical vision-language models. First, we curate a set of text perturbation methods that keep the same words, but disrupt the semantic structure of the sentence. Next, we apply different types of perturbation to reports, and use the model to distinguish the original report from the perturbed ones given the associated image. Parallel to this, we enhance the sensitivity of our method to higher level of granularity for both modalities by contrasting attention-weighted image sub-regions and sub-words in the image-text pairs. We conduct extensive experiments on multiple downstream tasks, and our method outperforms strong baseline methods. The results demonstrate that our approach learns more semantic meaningful and robust multi-modal representations.

Index Terms—AI for healthcare, multi-modal model, self-supervised learning

I. INTRODUCTION

In the ever-evolving landscape of the biomedical field, the advent of Artificial Intelligence (AI) has unfurled a new era, particularly with its deep learning (DL) tendrils extending into diverse healthcare applications. DL's transformative role in medical imaging, embracing a spectrum of tasks from image classification to segmentation, has been metaphorically hailed as 'A third eye for doctors' [1]. Natural language processing (NLP) techniques have also been applied in the medical domain and achieved great success in many tasks, including radiology natural language inference, medical question answering,

and others. Beyond imagery and linguistics, AI's prowess also manifests in the nuanced handling of varied data streams, including audio, video, and complex high-dimensional data, underscoring its multifaceted role in healthcare.

While AI has revolutionized healthcare, its advances have predominantly focused on single modalities. However, the intrinsic nature of healthcare information necessitates exploring multi-modal data. Multi-modal learning (MML), employing tokenization and embedding approaches, treats inputs from any modality as a token sequence to learn a joint embedding space. This approach not only mitigates data scarcity but also enables the development of diverse tasks and facilitates zero-shot learning [2] without specific fine-tuning.

Vision-language models (VLM) are one typical type of multi-modal models, with popular methods such as CLIP [3] and ALIGN [4] being successful. Although showing good performance, they grapple with challenges in relational understanding, fusion, and alignment between text and image modalities. VLMs also struggle with limited transferability, efficiency, universality, and interpretability. Within healthcare, these challenges are compounded by growing reporting backlogs, pressure on clinicians, such as strict radiology workflows slowing down case throughput especially as patient volumes rise, the absence of pretrained foundational models impeding rapid integration and deployment, and limited size of clinical datasets restricting the development of robust and versatile AI solutions. A particular hurdle is the unique healthcare discourse style and textual semantics, like the use of negation to rule out conditions (e.g., 'no evidence of pneumonia'), which could be ignored by AI models without sufficient data fed in to interpret complex medical data such as paired chest X-rays (CXR) and corresponding radiological reports written by experienced physicians.

Aware of the deficiency of clinical semantics understanding of current language models as well as vision-language models, [5] developed CXR-BERT tailored specifically in outstanding CXR reports by additional pre-training steps focused on radiology reports to capture their dense and complex semantics. While this specialization significantly enhanced vocabulary-specific understanding, it did not fully address the intricacies of clinical sentence structure. Together with the specialized text model, they presented a self-supervised vision-language approach BioViL using contrastive learning, demonstrating state-of-the-art performance in various biomedical benchmarks.

In this paper, we aim to propose a novel self-supervised VLM based on contrastive learning of jointly pretrained language and image models tailored for fine-grained understanding of clinical semantics and sentence structure, resulting in stronger vision-language alignment. Regarding high level of granularity of medical modalities, for example, the model should learn localized representations, where scene-level global information does not suffice. This is solved by introducing a local attentive contrastive loss which contrasts attention-weighted image sub-regions and words in the Image-Text pairs. Besides, the model should avoid the ‘object bias’ [6] similar to ‘bag of objects’ caused by CLIP-loss computation by our data-driven approach of loss design. We designed a set of diverse perturbations for every text in the original dataset, and the loss is computed reflecting the model’s ability to discriminate these generated texts together with the original text, which share almost or totally same retrieval results. To examine whether our proposed model has better compositional representations of objects, attributes, and relations for medical text encoders, we tailored an evaluation protocol using the aforementioned text perturbations. The main contributions of this paper can be summarized as follows:

- 1) We develop a novel self-supervised vision-language model that manipulates clinical texture well using diverse text perturbations derived from the standard image-text paired dataset in contrastive learning paradigm.
- 2) We introduce a hierarchical contrastive learning strategy, including a image-level contrastive loss to capture global structure, and a local contrastive loss which focus on fine details.
- 3) We further design a set of evaluation tasks to evaluate the model’s ability to understand clinical language representation other than just ‘bag of words’ based on the former perturbations.
- 4) Empirically, our proposed model outperforms the baseline vision-language model on our proposed evaluation protocol of sentence structure, medical natural language inference benchmarks including MedNLI [7] and RadNLI [8]. Especially, our model succeeds in multi-task image classification on a chest radiograph benchmark CheXpert [9], and outperforms baseline methods, demonstrating its strong ability in cross-modal correlation.

II. RELATED WORK

Our work lies in the evolving field of biomedical VLMs, which centers on extracting information from combination of medical images and reports. This area has recently gained traction, as evidenced by works like [5]. The majority of current methods rely on contrastive learning to bridge visual and linguistic elements. ConVIRT [10] firstly used bidirectional contrastive losses to project the two modalities into a shared space, and performed image classification task well while requiring much fewer annotated training data as an ImageNet-initialized counterpart. A shift in focus towards local rather than global information was introduced by LoVT [11], which targets localized medical imaging tasks like semantic segmentation or object detection by introducing a local contrastive loss to align image regions or report sentences while encouraging spatial smoothness and sensitivity. Wang et al. [12] extracted both image labels and entities extracted from reports to deepen understanding of medical semantics, eliminating the false negative noises and making use of vast image-only and text-only unpaired datasets. With the help of strong backbones, models such as BioViL [5] and ChexZero [13], which leverage a pretrained radiology-specific CXR-BERT text encoder and a pretrained CLIP [3] model respectively, also achieve comparable results.

Inspired by these prior studies, our work aims to comprehensively investigate the distinct clinical semantics as well as the high granularity of both medical vision and language that are essential to understanding biomedical multi-modal representation.

III. METHOD

Our proposed self-supervised vision-language learning architecture integrates unimodal feature extractors for both images and texts pretrained separately in advance and a contrastive projection approach to fuse the cross-modal embeddings into joint space. Our local-attentive contrastive loss is designed to boost precision in medical image-text matching by aligning specific image sub-regions with relevant text fragments. We also introduce report perturbation sensitivity loss to understand clinical semantics, focusing on sentence structure and part of speech relevant to clinical contexts, paired with the standard image report matching contrastive loss which primarily focuses on lexical identification. A series of text perturbations, integral to this process, are also employed in our subsequent evaluation protocol. The overall architecture is shown in Fig. 1.

A. Pretrained Unimodal Feature Extractor

With pre-trained text feature extractors carrying radiological knowledge by pretraining at hand, we choose to transfer powerful domain-specific language information from them instead of training the models from scratch which is time and energy-consuming. For each modality, we obtain the global features for subsequent calculation of image report matching contrastive loss and report perturbation sensitivity loss, as well as the local features for local attentive contrastive loss. For

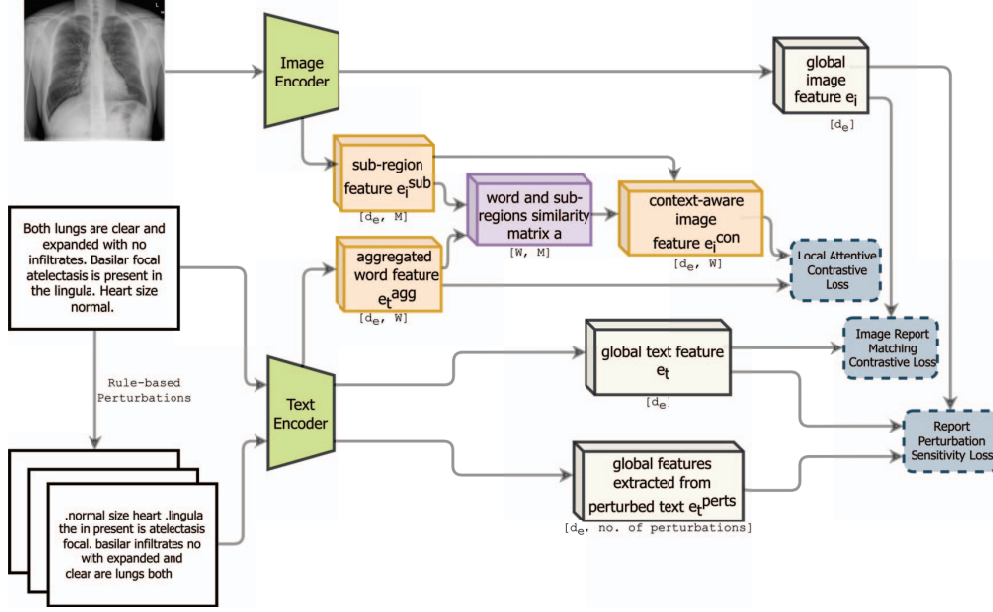


Fig. 1. Overall architecture of our proposed self-supervised vision-language learning framework.

text data, we used the CXR-BERT pretrained via Masked Language Modelling (MLM) on the PubMed [14], MIMIC-III [15] and MIMIC-CXR corpora as well as radiology section matching (RSM) which sequence prediction tasks were performed on MIMIC-CXR to match ‘Impression’ to ‘Findings’ sections of the same study. The output of the pretrained model become the global features of the text $e_t \in \mathbb{R}^{d_e}$. Considering the common abbreviations and typographical errors in medical reports, we employ word-piece level tokenization, resulting a sub-word feature $e_i^{sub} \in [K, N]$ for every report with W words, K dimension for every word-piece feature, and N as the sum of the number of the sub-words in each word, output from the encoder. The ultimate word-level local text embedding $e_t^{local} \in \mathbb{R}^{d_e}$ is aggregated by the average of the sub-word embeddings for each word, projected to a vector sized of d_e .

B. Joint Contrastive Projection

Based on the gained unimodal information in radiology domain from aforementioned feature extractors, we build our own contrastive projection model to fuse the vision and language embeddings into the joint space. Target at tackling the issue of high granularity of medical region of interests and unique clinical texture and semantics style which differs from the general language, we incorporated the local attentive contrastive loss to learn the salient image sub-region to attend to, report perturbation sensitive loss to learn the linguistic characteristics as well as the image report matching contrastive loss to perform the simple lexical retrieval task.

1) *Image Report Matching Contrastive Loss*: The basic image report matching contrastive loss [3] can be defined

as a symmetric cross-entropy loss based on cosine similarity between image and text features. Denoting the normalized image and text global embeddings as $e_i \in \mathbb{R}^{d_e}$ and $e_t \in \mathbb{R}^{d_e}$ correspondingly, cross-entropy loss for Image-Text and Text-Image of the scaled pairwise similarity for every batch B can be defined as below, where τ is the scaling parameter:

$$\mathcal{L}_{global} = -\log \frac{\exp(e_{i_j}^T \cdot e_{t_j} / \tau)}{\exp(e_{i_j}^T \cdot e_{t_j} / \tau) + \sum_{k \neq j} \exp(e_{i_k}^T \cdot e_{t_j} / \tau)}$$

By employing this image report matching loss, the matched Image-Text pair were attracted together, while unmatched pairs were pushed away.

2) *Local Attentive Contrastive Loss*: Medical images’ Region of Interests (ROI) are identified by more subtle cues than the natural images. This means only a small portion in a single graph aligns with the key word in the given report, requiring learning of local granules in addition to global vectors for Image-Text retrieval. In our local attentive contrastive loss [16], We adopted a context-aware image embedding $e_i^{con} \in [d_e, W]$ for every image, which is the attention-weighted sum of all M sub-region’s significance $e_i^{sig} \in [M, W]$ to every given word. The attention weight a_{jk} is the normalized similarity between a given word j and a given sub-region k across all image sub-regions:

$$a_{jk} = \frac{\exp(e_{i_k}^{localT} \cdot e_{t_j}^{local})}{\sum_l \exp(e_{i_l}^{localT} \cdot e_{t_j}^{local})}$$

The aggregated similarities S between all word embeddings and their corresponding context-aware image embeddings de-

finds the local attentive matching score for a Image-Text pair $[I, T]$:

$$S_{I,T} = \log\left(\sum_{k=1}^W \left(\exp\left(\sum_{j=1}^M a_{jk} \cdot v_k\right)\right)^T \cdot e_{t_k}^{local}\right)$$

Addition of the symmetrical cross-entropy loss for Image-Text and Text-Image matching becomes the final contrastive loss:

$$\mathcal{L}_{local} = -\log \frac{\exp(S_{I_i, T_i}/\tau)}{\exp(S_{I_i, T_i}/\tau) + \sum_{j \neq i} \exp(S_{I_i, T_j}/\tau)}$$

The model learns the fine-grained local semantic alignments between the image sub-regions and word pieces by employing this attentive contrastive loss.

3) *Report Perturbation Sensitivity Loss*: We present a report perturbation sensitivity loss to teach our model to understand the rich semantics in the radiological setting. Without expensive manual annotations or specific encoding of order and composition cues in advance, our method can enhance the text encoder’s preference for reports with correct word ordering and semantics. This is done by performing perturbations of every caption systematically based on our customized rules using spaCy [17].

a) *Generating Text Perturbations*: In [18], the author conducted a series of ablation experiments on language models, demonstrating several manipulations of the contexts are crucial in degrading the models’ performance. Inspired by that, we define a set of rules and generate 9 different perturbations as negatives by shuffling the words of specific entity types or according to specific patterns while keeping words unchanged. These perturbations, shown in Table I, are used both in the calculation of report perturbation sensitivity loss and subsequent evaluation.

b) *Contrastive Loss*: Deriving from the original captions, we use the generated perturbations as negative text samples in contrast with the original reports. The global image embeddings are denoted $e_{i_{ori}} \in \mathbb{R}^{d_e}$, and the global text embeddings from the original texts and the 9 generated perturbations are denoted as $e_{t_{ori}} \in \mathbb{R}^{d_e}$ and $e_{t_{Perts}} \in [d_e, 9]$. For a single image, the report perturbation contrastive loss is calculated as the cross-entropy loss of the similarity with the original aligned text against similarity with all other negatives:

$$\mathcal{L}_{pert} = -\log \frac{\exp(e_{i_{ori}}^T \cdot e_{t_{ori}}/\tau)}{\exp(e_{i_{ori}}^T \cdot e_{t_{ori}}/\tau) + \sum_k^{Perts} \exp(e_{i_{ori}}^T \cdot e_{t_k}/\tau)}$$

This way, the model aims to pick the original reports for an associated image, thus understanding the clinical contexts better.

Therefore, our final loss is:

$$\mathcal{L} = \mathcal{L}_{global} + \alpha \mathcal{L}_{local} + \beta \mathcal{L}_{pert}$$

IV. EXPERIMENT

To validate our proposed model, we conducted various downstream evaluation experiments including fine-tuning clinical text classification tasks on RadNLI and MedNLI, fine-tuning multi-task image classification task on CheXpert, as

well as our own clinical semantic structure evaluation task on Open-I dataset after pretraining.

A. Datasets

1) *Open-I [19]*: We used the Open-I dataset to train our vision-language model. The chest X-ray dataset contains 3,996 radiology reports associated with 8,121 images. The *findings* part of the reports are used, and 6469 Image-Text pairs remains after removing empty values.

2) *RadNLI [8] and MedNLI [7]*: RadNLI and MedNLI benchmarks contain 960 and 11k labeled hypothesis and premise pairs, respectively. MedNLI is collected in a broad clinical domain, while RadNLI is more radiology-specific. The task, also named Natural Language Inference (NLI), is to predict the label from the three label categories: entailment, contradiction, and neutral given the sentence pair.

3) *CheXpert [9]*: The CheXpert dataset contains a total of 224,316 chest radiographs from 65,240 patients, where each radiograph is paired with the corresponding radiology reports. Each radiograph is labeled for the presence of 14 total medical observations.

B. Implementation details

In all scenarios, we pretrained our multi-modal model for 150 epochs, with a batch size of 64. We use SGD [20] optimizer, and set the learning rate to 0.0015, momentum factor to 0.9, and weight decay to $5e-4$. For image preprocessing, we follow [5], using the same augmentations as used in the aforementioned pretrained image encoder. The dimensions of our input images are 224×224 pixels. To fine-tune the impact of each loss term and the margin in our model, we set the scaling hyper-parameters α and β at 0.1, and τ at 0.07. These values were selected based on outcomes from a targeted grid search.

C. Evaluation of extracted representation

We compare our model to its version that removes the local attentive contrastive loss, as well as two strong baselines under several configurations. ConVIRT [10] (Contrastive Visual Representation Learning from Text) is an efficient method of image representation learning from natural language supervision by contrasting only global representations of image and report pairs. In addition, within the same radiological domain, we also compared our architecture to GLoRIA [16] (Global-Local Representations for Images using Attention). This method also demonstrated high-performance and label-efficiency for various downstream medical image recognition tasks with limited labels. For a fair comparison, we keep the pretraining procedure unchanged, using the same 6000 Image-Text pairs of the Open-I dataset as input. The pretrained unimodal encoders are kept the same, using the same published checkpoints through the comparison. We conducted the same protocols on our downstream task as well.

TABLE I
LIST OF TEXT PERTURBATIONS FOR AN EXAMPLE RADIOLOGICAL REPORT

Perturbation Type	Example
Original Text	the lungs are clear there is no pleural effusion or pneumothorax
Shuffle All Words	pneumothorax are lungs the there is pleural or no effusion clear
Swap Adjacent Words in the Sentence	lungs the clear are is there pleural no or effusion pneumothorax
Reverse Sentence	pneumothorax or effusion pleural no is there clear are lungs the
Shuffle Within Trigrams	lungs the are there is clear pleural effusion no or pneumothorax
Shuffle Trigrams	or pneumothorax no pleural effusion the lungs are clear there is
Shuffle Nouns and Adjectives	the pneumothorax are clear there is no pleural lungs or effusion
Shuffle All but Nouns and Adjectives	there lungs is clear are the or pleural effusion no pneumothorax
Shuffle Nouns, Verbs, and Adjectives	the is pneumothorax lungs there pleural no are effusion or clear
Replace Adjectives with Antonyms	the lungs are unclear there is no pleural effusion or pneumothorax

TABLE II
RESULTS OF FINE-TUNED MULTI-TASK IMAGE CLASSIFICATION ON THE CHEXPert BENCHMARK

Accuracy	Consolidation	Pleural Effusion	Average
ConVIRT [10]	28.80%	43.60%	36.20%
GLoRIA [16]	71.11%	28.89%	50.00%
Ours (w/o local loss)	33.80%	77.80%	55.80%
Ours	93.40%	51.80%	72.60%

1) *Text classification*: For text classification, we fed the MedNLI and RadNLI datasets into the text encoder of the multi-modal model after pretraining and fine-tuning separately. Table III displays the accuracy of the different models on MedNLI and RadNLI benchmarks.

TABLE III
RESULTS OF FINE-TUNED TEXT CLASSIFICATION ON THE RADNLI AND MEDNLI BENCHMARKS

Method	MedNLI	RadNLI
ConVIRT [10]	86.80%	68.50%
GLoRIA [16]	86.64%	68.33%
Ours (w/o local loss)	87.62%	66.67%
Ours	85.79%	68.96%

2) *Multi-task image classification*: For fine-tuned multi-task image classification, we only used frontal images, and we use 20,596 images for training and 500 images for testing. We benchmark the performance on classification of *Consolidation* and *Pleural Effusion*. Similarly, we fed the dataset into the image encoder of the multi-modal model after pretraining and fine-tuning. Table II displays the accuracy of the different models on both separate tasks and averages.

3) *Clinical Semantic Structure Evaluation*: In our clinical semantic structure evaluation task, we used the aforementioned text perturbations as negatives. The cosine similarity is calculated for every caption inside a candidate caption list which includes an original text and its perturbations with the corresponding image using the global embeddings:

$$\mathcal{S}_{I,T} = \frac{e_I \cdot e_T}{\|e_I\|_2 \|e_T\|_2}$$

If the similarity between the original aligned Image-Text pair is the highest, then the decision is considered as cor-

rect. This means the model is able to retrieve the correctly structured report given an image. This evaluation task incorporates the representations of negative samples, outperforming the classical Image-Text retrieval task by introducing more difficult candidates.

We use 469 image-text pairs from the Open-I dataset as validation set, then feeding them into the uni-modal encoders of the multi-modal model, deriving the accuracy illustrated before. The results are shown in Table IV. The results show that our proposed method outperforms baseline methods.

TABLE IV
RESULTS OF ZERO-SHOT CLINICAL SEMANTIC STRUCTURE EVALUATION ON THE OPEN-I BENCHMARK

Method	Open-I
ConVIRT [10]	43.10%
GLoRIA [16]	44.30%
Ours (w/o local loss)	46.30%
Ours	49.00%

V. CONCLUSION

In this paper, we propose a novel pre-trained vision-Language model that enhances fine-grained clinical semantics understanding by increasing sensitivity to caption perturbations and focusing on local attention. This approach allows the model to capture not just broad image-level information, but also intricate clinical sub-region and word-piece details within both images and texts. By contrasting generated negative examples with original medical reports, our model gains a deeper understanding of medical report semantics and structure. This method has demonstrated significant improvements across various clinical datasets in multiple downstream benchmarks and our custom-designed composition sensitivity evaluation

task. In the future, we will explore pre-training our proposed model with images of higher resolution, and extending our model to other modalities, including electronic health records (EHR).

REFERENCES

- [1] A. Fourcade and R. H. Khonsari, "Deep learning in medical image analysis: A third eye for doctors," *Journal of stomatology, oral and maxillofacial surgery*, vol. 120, no. 4, pp. 279–288, 2019.
- [2] Y. Xian, C. H. Lampert, B. Schiele, and Z. Akata, "Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 9, pp. 2251–2265, 2018.
- [3] A. Radford, J. W. Kim, C. Hallacy, *et al.*, "Learning transferable visual models from natural language supervision," *International conference on machine learning*, pp. 8748–8763, 2021.
- [4] C. Jia, Y. Yang, Y. Xia, *et al.*, "Scaling up visual and vision-language representation learning with noisy text supervision," *International conference on machine learning*, pp. 4904–4916, 2021.
- [5] B. Boecking, N. Usuyama, S. Bannur, *et al.*, "Making the Most of Text Semantics to Improve Biomedical Vision–Language Processing," *en, arXiv.org*, Apr. 2022. DOI: 10.1007/978-3-031-20059-5_1.
- [6] M. Yuksekogonul, F. Bianchi, P. Kalluri, D. Jurafsky, and J. Zou, "When and why vision-language models behave like bags-of-words, and what to do about it?" *The Eleventh International Conference on Learning Representations*, 2022.
- [7] A. Romanov and C. Shivade, "Lessons from natural language inference in the clinical domain," *arXiv preprint arXiv:1808.06752*, 2018.
- [8] Y. Miura, Y. Zhang, E. B. Tsai, C. P. Langlotz, and D. Jurafsky, "Improving factual completeness and consistency of image-to-text radiology report generation," 2021. arXiv: 2010.10042 [cs.CL].
- [9] J. Irvin, P. Rajpurkar, M. Ko, *et al.*, "Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison," *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, pp. 590–597, 2019.
- [10] Y. Zhang, H. Jiang, Y. Miura, C. D. Manning, and C. P. Langlotz, "Contrastive learning of medical visual representations from paired images and text," *Machine Learning for Healthcare Conference*, pp. 2–25, 2022.
- [11] P. Müller, G. Kaissis, C. Zou, and D. Rueckert, "Joint learning of localized representations from medical images and reports," in 2022, pp. 685–701.
- [12] Z. Wang, Z. Wu, D. Agarwal, and J. Sun, "Medclip: Contrastive learning from unpaired medical images and text," *arXiv preprint arXiv:2210.10163*, 2022.
- [13] E. Tiu, E. Talius, P. Patel, C. P. Langlotz, A. Y. Ng, and P. Rajpurkar, "Expert-level detection of pathologies from unannotated chest x-ray images via self-supervised learning," *Nature Biomedical Engineering*, vol. 6, no. 12, pp. 1399–1406, 2022.
- [14] "Pubmed data download," *en, PubMed*, [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/download/>.
- [15] A. E. W. Johnson, T. J. Pollard, L. Shen, *et al.*, "MIMIC-III, a freely accessible critical care database," *en, Scientific Data*, vol. 3, no. 1, p. 160 035, May 2016, Number: 1 Publisher: Nature Publishing Group, ISSN: 2052-4463. DOI: 10.1038/sdata.2016.35.
- [16] S.-C. Huang, L. Shen, M. P. Lungren, and S. Yeung, "Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition," 2021, pp. 3942–3951.
- [17] M. Honnibal and I. Montani, "Spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. neural machine translation," 2017, pp. 688–697.
- [18] J. O'Connor and J. Andreas, "What context features can transformer language models use?" *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 851–864, 2021.
- [19] D. Demner-Fushman, S. Antani, M. Simpson, and G. Thoma, "Design and development of a multimodal biomedical information retrieval system," *Journal of Computing Science and Engineering*, vol. 6, Jun. 2012. DOI: 10.5626/JCSE.2012.6.2.168.
- [20] S. Ruder, "An overview of gradient descent optimization algorithms," 2017. arXiv: 1609.04747 [cs.LG].