Multimodal Image-Text Matching Improves Retrieval-based Chest X-Ray Report Generation

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2023

Abstract

Automated generation of clinically accurate radiology reports can improve patient care. Previous report generation methods that rely on image captioning models often generate incoherent and incorrect text due to their lack of relevant domain knowledge. Retrieval-based attempts frequently retrieve reports that are irrelevant to the input image. In this work, we propose Contrastive X-Ray REport Match (X-REM), a novel retrieval-based radiology report generation module that uses an image-text matching score to measure the similarity of a chest X-ray image and radiology report at the retrieval step. We observe that computing the image-text matching score with a language-image model can effectively capture the fine-grained interaction between image and text that is often lost when using cosine similarity. X-REM outperforms multiple prior radiology report generation modules in terms of both natural language and clinical metrics. Human evaluation of the generated reports suggests that X-REM increased the number of zero-error reports and decreased the average error severity compared to the baseline retrieval approach.

Keywords: Medical Report Generation, Vision-Language Modeling, Clinical Expert Evaluation

1. Introduction

An AI-based model that automatically generates radiology reports can dramatically improve both clinical workflows and patient care (Boag et al., 2020; Hartung et al., 2020). Such models can triage cases, lessen radiologist workloads, and reduce delays in diagnosis(Nsengiyumva et al., 2021; Dyer et al., 2021; Annarumma et al., 2019). Additionally, AI tools, whether acting autonomously or in collaboration with radiologists, can improve human accuracy (Seah et al., 2021; Sim et al., 2020).

However, current report generation models are yet to match expert-level performance and are not ready for translation into clinical practice (Jing et al., 2018; Chen et al., 2020). While many prior methods adopt image-captioning models to directly generate reports from an image input, the generated texts often contain hallucinated information and self-contradictory claims (Yan et al., 2021; Miura et al., 2021). This issue rises from the fact that the natural language generative models lack the medical background knowledge necessary to self-evaluate the sanity of the generated reports by themselves, therefore being prone to hallucination. Another stream of work approaches medical report generation as an image-text retrieval task that retrieves, from a medical corpus of radiology reports, a report that best describes the input image as a human-written report would be clinically coherent (Endo et al., 2021). Yu et al. (2022) has shown that while the theoretical upper bound for the performance of a retrieval-based generation method is high, prior approaches are nowhere

© 2023 Anonymous.

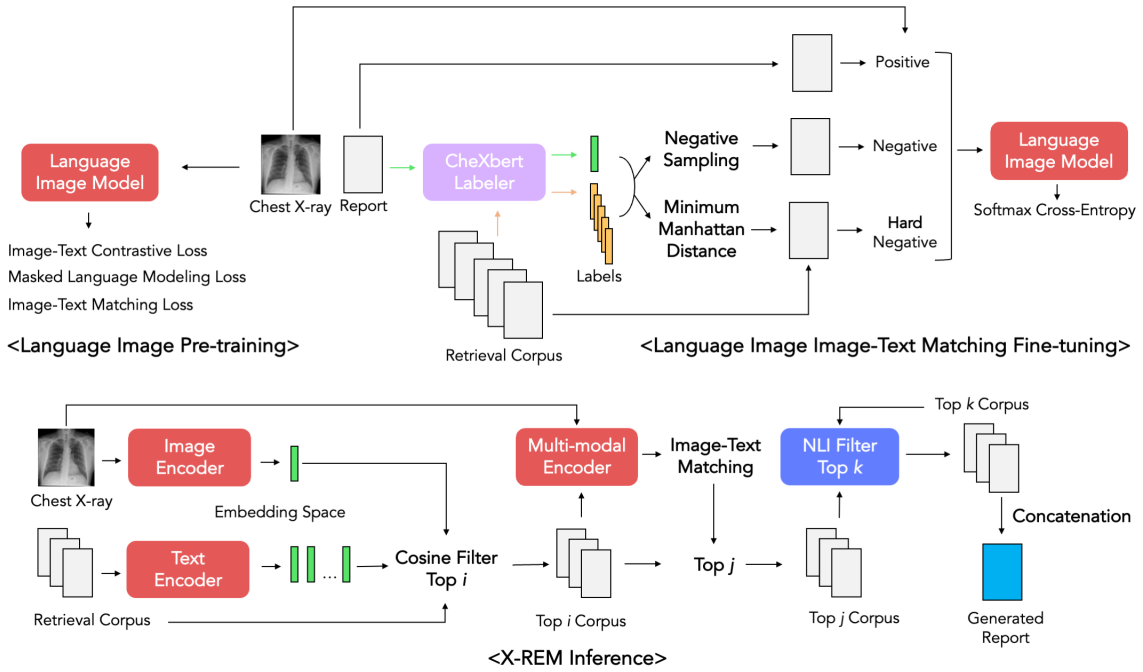

Figure 1: An overview of the pre-training, fine-tuning, and inference steps of X-REM

near the upper bound as they often fail to select the ideal report from a pool of retrieval corpus.

To this end, we propose Contrastive X-Ray REport Match (X-REM), an innovative, retrieval-based radiology report generation method that differs from the current paradigm in two key ways. First, our approach emphasizes multimodality, leveraging a language-image model to acquire a joint representation of the image and text, as opposed to separately representing them using two unimodal encoders. Secondly, at the retrieval step, we use an image-text matching score as the main similarity metric. In addition to aligning the image and text embeddings in the unsupervised pre-training phase, we also perform supervised contrastive learning that fine-tunes the model to match image and text with the same clinical label. Image-text matching score that is further tuned on the domain-specific features better captures the complicated interaction between a medical image and a report than the conventional cosine similarity, which primarily relies on the alignment of the latent representations.

When evaluated on the pre-processed test split of MIMIC-CXR (Johnson et al., 2019), X-REM shows a noticeable improvement over previous radiology report generation modules on both clinical and natural language metrics. Additionally, we asked radiologists to provide annotations of error scores on reports generated by X-REM and the baseline retrieval-based approach. The human evaluation study demonstrates that our model makes significant improvements upon the baseline model.

2. Related Works

Earlier works in radiology report generation have focused heavily on performing image-captioning with a transformer architecture. R2Gen (Chen et al., 2020) and $\mathcal{M}^2$Trans (Miura et al., 2021) use a memory-driven Transformer and a Meshed-Memory Transformer (Cornia et al., 2019), respectively, to directly generate a radiology report. WCL (Yan et al., 2021) shows that incorporating a weakly supervised contrastive loss while training the encoder-decoder models improves the model performance. CvT2DistilGPT2 (Nicolson et al., 2022) demonstrates that pre-trained checkpoints for traditional computer vision and natural language tasks can be helpful for radiology report generation as well. However, one main shortcoming of the image-captioning models is the presence of medically inconsistent information frequently found in the generated reports. To address this issue, CXR-RePaiR (Endo et al., 2021) adopts an image-text retrieval method that retrieves a report whose CLIP (Radford et al., 2021) text embedding scores the highest cosine similarity with the chest X-ray's CLIP image embedding. X-REM improves upon CXR-RePaiR by utilizing the image-text matching score computed with a multimodal encoder to acquire a finer representation of the similarity of the inputs.

3. Method

As shown in figure 1, X-REM first acquires the input image's embedding using the language-image model's image encoder. Next, it narrows down the retrieval corpus to the top i reports with the highest cosine similarity with the input image. X-REM then selects the top j reports with the highest image-text matching score with the input image. X-REM traverses the j reports in the order of decreasing image-text matching score and selects at most k reports, using a natural language inference (NLI) score to filter out reports with repetitive information. Finally, the model concatenates the retrieved reports into a single document.

3.1. Image-Text Matching

Metric Image-text matching is a binary classification task that predicts whether the image and the text describe the same event. As image-text matching score has been fine-tuned to learn the multimodal interaction, it can inherently make a more sophisticated judgment than cosine similarity which lacks such additional tuning. In this study, we propose using the logit value for the task's `positive` class as the similarity score of the pair. Formally, given a multimodal encoder fine-tuned on image-text matching task M_{ITM}, input chest X-ray image x, and a radiology report r, we define the image-text matching score function as

$$f_{ITM}(x, r) = M_{ITM}(x, r)[\texttt{positive}]. \tag{1}$$

Negative Samples As conventional medical datasets consist solely of valid image-text pairs, synthetic generation of negative samples is necessary to construct a fine-tuning dataset for image-text matching. After generating a 14-dimensional clinical label of the reports in MIMIC-CXR's train split using CheXbert labeler (Smit et al., 2020), we generate two types of negative samples: a sample that randomly matches images and reports with different CheXbert labels and a sample that matches images and reports with the smallest nonzero

Manhattan CheXbert label distance. The intention of introducing hard negatives is to boost the model's precision on predicting `positive`.

3.2. Language Image Model

In this study, we choose ALBEF (Li et al., 2021) as the backbone language-image model for X-REM. ALBEF directly aligns the embeddings of its image and text encoders, allowing us to its encoders both jointly to compute the image-text matching scores and separately to compute the cosine similarity. Moreover, ALBEF does not require an object detection module in its image encoder, which is desirable given the unavailability of bounding-box annotated chest X-ray images needed to train the object detection module.

3.3. Cosine Similarity Filter

As the computation for image-text matching score may be a resource-heavy process, we take advantage of ALBEF's already-aligned unimodal embeddings and first narrow down the pool of candidate reports to top i number of reports that score high cosine similarity with the input image. Then we compute the image-text matching scores for the i reports to select top j reports with the highest image-text matching scores.

3.4. Natural Language Inference Filter

We prevent the model from retrieving multiple reports with either overlapping or conflicting content by appending a filter that rejects a candidate report if it is either entailed or contradicted by the generated report. Formally, let k represent the desired number of reports retrieved and $B = \{r_1, \ldots, r_m\}$ the batch of m reports we have selected so far. For a given premise p and hypothesis h, let $g(p, h)$ be a language model that classifies the relationship between p and h into one of {`contradiction, entailment, neutral`}. Then, we add $r \in B_{ITM}$ to B if

$$m < k \wedge g(B, r) = \texttt{neutral}. \tag{2}$$

3.5. Learning Objectives

X-REM is a three-layered module that uses unimodal image and text encoders for measuring cosine similarity, a multimodal encoder for measuring image-text matching score, and a language model for classifying NLI relationship. For training the two unimodal encoders, we follow the three pre-training objectives proposed by Li et al. (2021): $\mathcal{L}_{\text{itc}}$, an image-text contrastive loss that directly aligns the unimodal representations of the image and text, $\mathcal{L}_{\text{mlm}}$, a masked language modeling loss that predicts the randomly-masked tokens, and $\mathcal{L}_{\text{itm}}$, an image-text matching loss that predicts whether the given image and text are positively matched. The main difference between the image-text matching loss in the pre-training and the fine-tuning step is the usage of supervised clinical labels for generating the negative samples. We fine-tune both the multi-modal encoder and the language model on image-text matching and natural language inference, respectively, to perform a n-way classification with a cross-entropy loss.

3.6. Data

MIMIC-CXR (Johnson et al., 2019) is a de-identified, publicly available dataset of chest radiographs and semi-structured radiology reports that contains over 225,000 studies from over 65,000 patients. We follow the pre-processing steps suggested by Endo et al. (2021) to combine the train and validation splits into a bigger train split and then extract the impression sections of the radiology reports, removing reports that do not contain the impressions section. Additionally, we reserve 60 studies from the training set for human evaluation. Our final training data has 185,538 studies, consisting of 313,777 pairs of chest X-ray images and radiology reports. The same training set serves as the retrieval corpus in the inference step. For our test set, we refer to Yu et al. (2022) and remove multiple samples that are from the same study. While filtering out samples from a duplicate study, we prioritize samples with images in either PA or AP view, randomly choosing between them if both views are present. Our final test set consists of 2,192 studies and 2,192 pairs of chest X-ray images and reports.

4. Experiments

Implementation For the language-image model, we use the default ALBEF architecture that consists of a $\text{BERT}_{\text{base}}$ (Devlin et al., 2019) with 123.7M parameters and a ViT-B/16 (Dosovitskiy et al., 2020) with 85.8M parameters. After loading the published pre-trained weights (ALBEF 4M), we further pre-train ALBEF on MIMIC-CXR for 60 epochs with a learning rate of 1×10^{-5}. Then, we fine-tune the language-image model on the image-text matching task for 8 epochs with the default learning rate of 2×10^{-5}. We resize the input chest X-ray images into the shape of $(256, 256, 3)$ for pre-training and $(384, 384, 3)$ for fine-tuning, adding black padding to the boundary if necessary. We also remove the original language-image model's on-the-fly data augmentation steps such as random cropping to ensure that the model always has access to the entire, uncorrupted chest X-ray. It took us 2 days to pre-train the language-image model and 4 days to fine-tune it on 4 NVIDIA RTX 8000 GPUs. For the language model, we use $\text{BERT}_{\text{base}}$ with the published checkpoint fine-tuned on MedNLI (Romanov and Shivade, 2018) and RadNLI(Miura et al., 2021).

Evaluation Metrics We refer to Yu et al. (2022) and use RadCliQ (Yu et al., 2022), RadGraph F_1 (Jain et al., 2021), CheXbert vector similarity (Smit et al., 2020), BERTScore (Zhang et al., 2019), and BLEU (Papineni et al., 2002) as the main metrics for evaluation. The two natural language metrics we use are BLEU as BLEU-2 bigrams and BERTScore with distillroberta-base as the base BERT model. RadGraph F_1 is a clinical metric that converts each of the two input radiology reports into a knowledge graph and then measures the degree of overlap between the two graphs. CheXbert vector similarity first generates a 14-dimensional vector that indicates the presence of 13 common symptoms and the no-finding observation for each of the reports and then computes the cosine similarity between the two vectors. RadCliQ predicts the radiologist-determined error by combining BLEU-2 and RadGraph F_1. As the RadCliQ score is an estimation of the error, smaller RadCliQ values are preferable. Yu et al. (2022) has recently shown that RadCliQ is the most aligned with the radiologists' evaluation of the reports, and we therefore optimize the model on RadCliQ.

		RadCliQ	RadGraph F_1	CheXbert	BERTScore	BLEU2
1	$\mathcal{M}^2$ Trans*	3.087	0.111	0.268	0.227	**0.087**
2	R2Gen*	3.232	0.057	0.203	0.186	0.059
3	WCL*	3.205	0.068	0.218	0.188	0.064
4	CvT2DistilGPT2*	3.191	0.073	0.264	0.193	0.066
5	BLIP	3.304	0.046	0.309	0.190	0.030
6	CXR-RePaiR*	3.181	0.091	0.379	0.191	0.055
7	**X-REM**	**3.054**	**0.133**	**0.384**	**0.287**	0.084

Table 1: X-REM outperforms previous report generation models on multiple metrics. Results with * are from a previous work by Yu et al. (2022) who evaluated the models on the identically-preprocessed MIMIC-CXR test set.

4.1. Comparison

In Table 1, we compare the performance of X-REM with multiple prior radiology report generation modules $\mathcal{M}^2$ Trans, R2GEN, CXR-RePaiR, WCL, and CvT2DistilGPT2 on the preprocessed test split of MIMIC-CXR. Moreover, in order to directly compare the retrieval approach with the image-captioning method, we generate baseline image-captioning results with BLIP (Li et al., 2022), an image-captioning model that was pre-trained and fine-tuned on MIMIC-CXR for 50 epochs and 30 epochs respectively. The results indicate that X-REM shows a notable improvement from the previous models measured on BERTScore, CheXbert vector similarity, RadGraph F_1, and RadCliQ. The wide gap between the scores for CXR-RePaiR and X-REM suggests that image-text matching score serves as a better similarity metric than cosine similarity. The poor performance of BLIP on the metrics demonstrates the limitations of an image-captioning model without the additional learning objectives introduced in the prior radiology report generation modules.

4.2. Expert Evaluation

We conduct a human evaluation on a subset of 60 studies from MIMIC-CXR and recruit four radiologists to provide evaluations. We note that our study has been approved by an institutional review board. For each study image, we compile three reports: one generated from our method (X-REM), one generated from a machine learning baseline (CXR-RePaiR) trained on the same MIMIC-CXR training set, and one taken from a human benchmark (MIMIC-CXR). To each radiologist, we present one image and one report for each of the 60 studies. Each report is randomly and independently chosen from one of the three sources. The radiologist is blinded to the source. We ask each radiologist to assess the error severity of their assigned reports.

Error Scoring: Each report is broken down into lines. The radiologist is asked to score the error of each line of the report. The radiologists select from five possible error categories (No error, Not actionable, Actionable nonurgent error, Urgent error, or Emergent error), which we map to error scores 0, 1, 2, 3, 4, respectively, in order of severity. We collapse severity per line to a single severity score per annotation using the following metrics: (1) maximum error severity (MES) across all lines in the report, which measures the worst error

	MES				AES			
	0	≤ 1	≤ 2	≤ 3	0	≤ 1	≤ 2	≤ 3
X-REM (N=117)	0.17	0.35	0.48	0.87	0.23	0.46	0.68	0.91
Baseline (N=69)	0.10	0.33	0.45	0.86	0.11	0.34	0.52	0.82
Human Benchmark (N=53)	0.34	0.49	0.64	0.94	0.35	0.56	0.69	0.94

Table 2: Human Evaluation Study Results. In each row, we show the CDF for each report source across MES and AES scoring. For example, under MES column 3, we have the proportion of reports which have a max error severity of 3 or less.

in the report, and (2) average error severity (AES) per line in the original report. AES is the sum of error severity across lines standardized by the number of lines, which avoids punishing longer reports.

Our analysis finds that X-REM increases the number of zero-error reports by 70% and provides statistically strong improvements over the baseline. We present the distribution of MES and AES across annotations for each source in Table 2. According to the radiologists, 17% of reports generated by X-REM have zero mistakes, which surpasses 10% from the baseline but lags behind the human benchmark's 35%.

Paired Comparison: In order to directly assess the improvement of X-REM over the baseline, we analyze the 40 studies that have at least one annotation from both CXR-RePaiR and the baseline. We find that in 55% of studies, X-REM strictly improves MES against the baseline, and 62.5% have the same or lower MES compared to the baseline. On average across shared cases, our method reduces MES from 2.38 to 2.00 compared to the baseline. On the AES metric, 57.5% of shared studies have a strictly lower AES with X-REM than the baseline, while 65% have an equal or lower AES compared to the baseline. On average across shared cases, X-REM reduces AES from 2.47 to 1.82 compared to the baseline.

4.3. Ablation Experiments

In row 1, 2, and 4 of Table 3, we evaluate the effect of appending a post-processing filter under the following three settings: appending an NLI-based filter (row 1), appending a BERTScore-based filter (row 2), and not appending a filter (row 4). For the BERTScore-based filter, we select the candidate text r when $\text{BERTScore}_{\text{F}_1}(B, r) < 0.5$. While preliminary results suggested that applying the NLI filter yielded the best RadCliQ score, after Yu et al. (2022) revamped the RadCliQ metric to compose solely of BLEU and RadGraph F_1, not applying any filter gives the best RadCliQ score. Moreover, we observe that the presence of a filter increases the natural language metrics while decreasing the clinical metrics.

Row 3 and 4 of Table 3 compare the performance of cosine similarity and image-text matching score as the final similarity metric used in the retrieval. When we compare the retrieved reports (without a filter at the end), we observe that the image-text matching score outperforms cosine similarity across every metric. Moreover, the vanilla image-text retrieval module with cosine similarity (row 3) still outperforms CXR-RePaiR from Table 1, which can be attributed to the performance gap of the base language-image models.

	Metric	Filter	i	k	P-T	RadCliQ (↓)	RadGraph F_1	CheXbert	BERTScore	BLEU2
1	ITM	NLI	50	2	✓	3.054	0.133	0.384	0.287	0.084
2	ITM	BERTScore	50	2	✓	3.052	0.130	0.387	0.281	0.088
3	Cosine	N/A	50	2	✓	3.079	0.122	0.336	0.244	0.081
4	ITM	N/A	50	2	✓	3.047	0.134	0.391	0.279	0.086
5	ITM	NLI	10^3	2	✓	3.055	0.131	0.383	0.287	0.085
6	ITM	NLI	50	1	✓	3.117	0.109	0.420	0.279	0.072
7	ITM	NLI	50	3	✓	3.055	0.134	0.358	0.273	0.083
8	ITM	NLI	50	2		3.214	0.072	0.250	0.195	0.055

Table 3: Ablation of models with varying choices for the final filter (row 1,2), similarity metric for retrieval (row 3,4), parameter values for i and k (row 5,6,7), and pre-training the model (row 8). Row 1 reports the original version of X-REM.

Row 5,6, and 7 of Table 3 evaluate the model performance under various values for i and k. The small gap between row 1 and 5 suggests that cosine similarity is useful for pruning irrelevant reports in the earlier stages. A comparison of rows 1, 6, and 7 indicates that the model benefits from retrieving multiple reports as opposed to a single report.

In row 8, there is a significant drop in the performance of a model that was only fine-tuned from the published pre-trained checkpoint compared to row 1, which has been both pre-trained and fine-tuned on MIMIC-CXR. The difference suggests that radiology-related images and texts are vastly different from the data ALBEF was initially pre-trained on (Conceptual Captions (Sharma et al., 2018), SBU Captions (Ordonez et al., 2011), COCO (Lin et al., 2014), and Visual Genome (Krishna et al., 2016)) and that additionally pre-training the language-image model on domain-relevant data is necessary.

5. Limitations

As X-REM generates an impression section of a radiology report, we only referred to the impression sections of MIMIC-CXR. However, we expect that including other components of a radiology report such as the findings and the indications section can improve the model performance by providing more information about the patient. Looking at multiple chest X-rays taken from different angles may be helpful as well. Finally, X-REM is still below the theoretical upper-bound performance of a retrieval-based approach calculated by Yu et al. (2022), indicating the high likelihood of a presence of a better similarity metric for the retrieval step.

Given the radiology report generation model's potential application in a medical setting, the models should be evaluated with a high standard. While our model has been optimized on clinical and natural language metrics, there may exist other important factors to reflect during the evaluation process. We focused on training and testing models on a single medical dataset that may contain bias, and the model's out-of-domain performance has not been empirically shown yet.

References

Mauro Annarumma, Samuel J. Withey, Robert J. Bakewell, Emanuele Pesce, Vicky Goh, and Giovanni Montana. Automated triaging of adult chest radiographs with deep artificial neural networks. *Radiology*, 291(1):196–202, 2019. doi: 10.1148/radiol.2018180921. URL https://doi.org/10.1148/radiol.2018180921. PMID: 30667333.

William Boag, Tzu-Ming Harry Hsu, Matthew Mcdermott, Gabriela Berner, Emily Alesentzer, and Peter Szolovits. Baselines for Chest X-Ray Report Generation. In Adrian V. Dalca, Matthew B.A. McDermott, Emily Alsentzer, Samuel G. Finlayson, Michael Oberst, Fabian Falck, and Brett Beaulieu-Jones, editors, *Proceedings of the Machine Learning for Health NeurIPS Workshop*, volume 116 of *Proceedings of Machine Learning Research*, pages 126–140. PMLR, 13 Dec 2020. URL https://proceedings.mlr.press/v116/boag20a.html.

Zhihong Chen, Yan Song, Tsung-Hui Chang, and Xiang Wan. Generating radiology reports via memory-driven transformer. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1439–1449, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.112. URL https://aclanthology.org/2020.emnlp-main.112.

Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. *CoRR*, abs/1912.08226, 2019. URL http://arxiv.org/abs/1912.08226.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL https://aclanthology.org/N19-1423.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *CoRR*, abs/2010.11929, 2020. URL https://arxiv.org/abs/2010.11929.

T. Dyer, L. Dillard, M. Harrison, T. Naunton Morgan, R. Tappouni, Q. Malik, and S. Rasalingham. Diagnosis of normal chest radiographs using an autonomous deep-learning algorithm. *Clinical Radiology*, 76(6):473.e9–473.e15, 2021. ISSN 0009-9260. doi: https://doi.org/10.1016/j.crad.2021.01.015. URL https://www.sciencedirect.com/science/article/pii/S0009926021000763.

Mark Endo, Rayan Krishnan, Viswesh Krishna, Andrew Y. Ng, and Pranav Rajpurkar. Retrieval-based chest x-ray report generation using a pre-trained contrastive language-image model. In Subhrajit Roy, Stephen Pfohl, Emma Rocheteau, Girmaw Abebe Tadesse, Luis Oala, Fabian Falck, Yuyin Zhou, Liyue Shen, Ghada Zamzmi, Purity

Mugambi, Ayah Zirikly, Matthew B. A. McDermott, and Emily Alsentzer, editors, *Proceedings of Machine Learning for Health*, volume 158 of *Proceedings of Machine Learning Research*, pages 209–219. PMLR, 04 Dec 2021. URL https://proceedings.mlr.press/v158/endo21a.html.

Michael P. Hartung, Ian C. Bickle, Frank Gaillard, and Jeffrey P. Kanne. How to create a great radiology report. *RadioGraphics*, 40(6):1658–1670, 2020. doi: 10.1148/rg.2020200020. URL https://doi.org/10.1148/rg.2020200020. PMID: 33001790.

Saahil Jain, Ashwin Agrawal, Adriel Saporta, Steven Q. H. Truong, Du Nguyen Duong, Tan Bui, Pierre Chambon, Yuhao Zhang, Matthew P. Lungren, Andrew Y. Ng, Curtis P. Langlotz, and Pranav Rajpurkar. Radgraph: Extracting clinical entities and relations from radiology reports. *CoRR*, abs/2106.14463, 2021. URL https://arxiv.org/abs/2106.14463.

Baoyu Jing, Pengtao Xie, and Eric Xing. On the automatic generation of medical imaging reports. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2577–2586, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1240. URL https://aclanthology.org/P18-1240.

Alistair E. W. Johnson, Tom J. Pollard, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih-ying Deng, Yifan Peng, Zhiyong Lu, Roger G. Mark, Seth J. Berkowitz, and Steven Horng. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs, 2019. URL https://arxiv.org/abs/1901.07042.

Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *CoRR*, abs/1602.07332, 2016. URL http://arxiv.org/abs/1602.07332.

Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven C. H. Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *CoRR*, abs/2107.07651, 2021. URL https://arxiv.org/abs/2107.07651.

Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. *CoRR*, abs/2201.12086, 2022. URL https://arxiv.org/abs/2201.12086.

Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. URL http://arxiv.org/abs/1405.0312.

Yasuhide Miura, Yuhao Zhang, Emily Tsai, Curtis Langlotz, and Dan Jurafsky. Improving factual completeness and consistency of image-to-text radiology report generation.

In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5288–5304, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.416. URL https://aclanthology.org/2021.naacl-main.416.

Aaron Nicolson, Jason Dowling, and Bevan Koopman. Improving chest x-ray report generation by leveraging warm-starting. *CoRR*, abs/2201.09405, 2022. URL https://arxiv.org/abs/2201.09405.

Ntwali Placide Nsengiyumva, Hamidah Hussain, Olivia Oxlade, Arman Majidulla, Ahsana Nazish, Aamir J Khan, Dick Menzies, Faiz Ahmad Khan, and Kevin Schwartzman. Triage of Persons With Tuberculosis Symptoms Using Artificial Intelligence–Based Chest Radiograph Interpretation: A Cost-Effectiveness Analysis. *Open Forum Infectious Diseases*, 8(12), 12 2021. ISSN 2328-8957. doi: 10.1093/ofid/ofab567. URL https://doi.org/10.1093/ofid/ofab567. ofab567.

Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL https://proceedings.neurips.cc/paper/2011/file/5dd9db5e033da9c6fb5ba83c7a7ebea9-Paper.pdf.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL https://aclanthology.org/P02-1040.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *CoRR*, abs/2103.00020, 2021. URL https://arxiv.org/abs/2103.00020.

Alexey Romanov and Chaitanya Shivade. Lessons from natural language inference in the clinical domain. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1187. URL https://aclanthology.org/D18-1187.

Jarrel Seah, Cyril Tang, Quinlan Buchlak, Xavier Holt, Jeffrey Wardman, Anuar Aimoldin, Nazanin Esmaili, Hassan Ahmad, Hung Pham, John Lambert, Ben Hachey, Stephen Hogg, Benjamin Johnston, Christine Bennett, Luke Oakden-Rayner, Peter Brotchie, and Catherine Jones. Effect of a comprehensive deep-learning model on the accuracy of chest x-ray interpretation by radiologists: a retrospective, multireader multicase study. *The Lancet Digital Health*, 3, 07 2021. doi: 10.1016/S2589-7500(21)00106-0.

Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational*

Linguistics (Volume 1: Long Papers), pages 2556–2565, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1238. URL https://aclanthology.org/P18-1238.

Yongsik Sim, Myung Jin Chung, Elmar Kotter, Sehyo Yune, Myeongchan Kim, Synho Do, Kyunghwa Han, Hanmyoung Kim, Seungwook Yang, Dong-Jae Lee, and Byoung Wook Choi. Deep convolutional neural network–based software improves radiologist detection of malignant lung nodules on chest radiographs. *Radiology*, 294(1):199–209, 2020. doi: 10.1148/radiol.2019182465. URL https://doi.org/10.1148/radiol.2019182465. PMID: 31714194.

Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Ng, and Matthew Lungren. Combining automatic labelers and expert annotations for accurate radiology report labeling using BERT. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1500–1519, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.117. URL https://aclanthology.org/2020.emnlp-main.117.

An Yan, Zexue He, Xing Lu, Jiang Du, Eric Chang, Amilcare Gentili, Julian McAuley, and Chun-Nan Hsu. Weakly supervised contrastive learning for chest X-ray report generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4009–4015, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-emnlp.336. URL https://aclanthology.org/2021.findings-emnlp.336.

Feiyang Yu, Mark Endo, Rayan Krishnan, Ian Pan, Andy Tsai, Eduardo Pontes Reis, Eduardo Kaiser Ururahy Nunes Fonseca, Henrique Min Ho Lee, Zahra Shakeri Hossein Abad, Andrew Y. Ng, Curtis P. Langlotz, Vasantha Kumar Venugopal, and Pranav Rajpurkar. Evaluating progress in automatic chest x-ray radiology report generation. *medRxiv*, 2022. doi: 10.1101/2022.08.30.22279318. URL https://www.medrxiv.org/content/early/2022/08/31/2022.08.30.22279318.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with BERT. *CoRR*, abs/1904.09675, 2019. URL http://arxiv.org/abs/1904.09675.