

# Fact-Checking of AI-Generated Reports

Razi Mahmood<sup>1</sup>, Diego Machado Reyes<sup>1</sup>, Ge Wang<sup>1</sup>, Mannudeep Kalra<sup>2</sup>, and Pingkun Yan<sup>1</sup>

<sup>1</sup> Department of Biomedical Engineering and Center for Biotechnology and Interdisciplinary Studies, Rensselaer Polytechnic Institute, Troy, NY 12180, USA

<sup>2</sup> Department of Radiology, Massachusetts General Hospital, Harvard Medical School, Boston MA 02114, USA

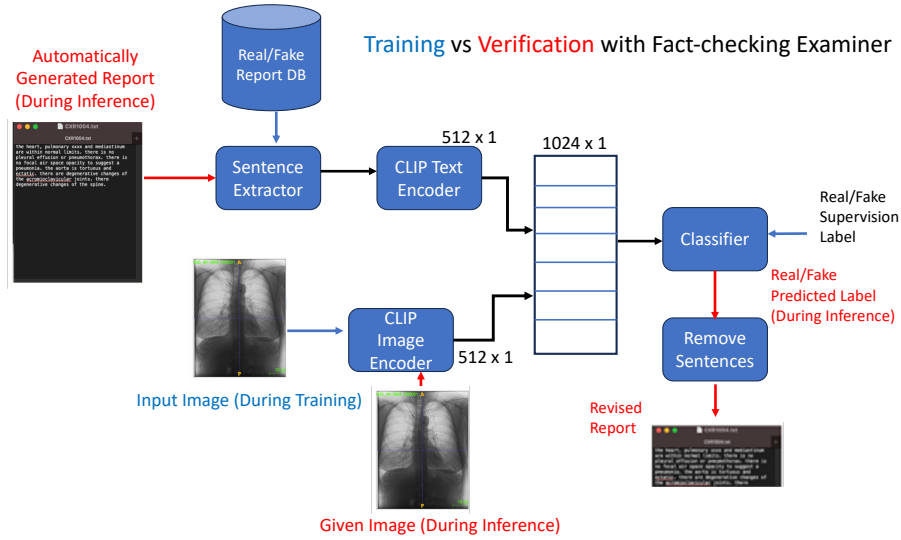
**Abstract.** With advances in generative artificial intelligence (AI), it is now possible to produce realistic-looking automated reports for preliminary reads of radiology images. However, it is also well-known that such models often hallucinate, leading to false findings in the generated reports. In this paper, we propose a new method of fact-checking of AI-generated reports using their associated images. Specifically, the developed examiner differentiates real and fake sentences in reports by learning the association between an image and sentences describing real or potentially fake findings. To train such an examiner, we first created a new dataset of fake reports by perturbing the findings in the original ground truth radiology reports associated with images. Text encodings of real and fake sentences drawn from these reports are then paired with image encodings to learn the mapping to real/fake labels. The examiner is then demonstrated for verifying automatically generated reports.

**Keywords:** Generative AI · Chest X-rays · Fact-checking · Radiology Reports.

## 1 Introduction

With the developments in radiology artificial intelligence (AI), many researchers have turned to the problem of automated reporting of imaging studies [3, 5, 11, 13, 15, 16, 22, 26]. This can significantly reduce the dictation workload of radiologists, leading to more consistent reports with improved accuracy and lower overall costs. While the previous work has largely used image captioning [23, 27] or image-to-text generation methods for report generation, more recent works have been using large language models (LLMs) such as GPT-4 [6, 12]. These newly emerged LLMs can generate longer and more natural sentences when prompted with good radiology-specific linguistic cues [4, 7].

However, with powerful language generation capabilities, hallucinations or false sentences are prevalent as it is difficult for those methods to identify their own errors. This has led to fact-checking methods for output generated by LLMs and large vision models (LVMs) [10, 18, 21]. Those methods detect errors either through patterns of phrases found repeatedly in text or by consulting other external textual sources for the veracity of information [10, 18, 21]. In radiology



**Fig. 1.** Illustration of the training (blue) and inference (red) phases of the image-driven fact-checking examiner. The operations common to both phases are in black.

report generation, however, we have a potentially good source for fact checking, namely, the associated images, as findings reported in textual data must be verifiable through visual detection in the associated imaging. Since most methods of report generation already examine the images in order to detect findings and generate the sentences, bootstrapping them with an independent source of verification is needed in order to identify their own errors.

In this paper, we propose a new imaging-driven method of fact-checking of AI-generated reports. Specifically, we develop a fact-checking examiner to differentiate between real and fake sentences in reports by learning the association between an image and sentences describing real or potentially fake findings. To train such an examiner, we first create a new dataset of fake reports by perturbing the findings in the original ground truth radiology reports associated with images. Text encodings of real and fake sentences drawn from these reports are then paired with image encodings to learn the mapping to real or fake labels via a classifier. The utility of such an examiner is demonstrated for verifying automatically generated reports by detecting and removing fake sentences. Future generative AI approaches can use the examiner to bootstrap their report generation leading to potentially more reliable reports.

Our overall approach to training and inference using the examiner is illustrated in Figure 1. To create a robust examiner that is not attuned for any particular automated reporting software, it is critical to create a dataset for training that encompasses a wide array of authentic and fabricated samples. Hence we first synthesize a dataset of real and fake reports using a carefully controlled process of perturbation of actual radiology reports associated with

the images. We then pair each image with sentences from its corresponding actual report as real sentences with real label, and the perturbed sentences from fake reports as fake sentences with fake label. Both textual sentence and images are then encoded by projecting in a joint image-text embedding space using the CLIP model [19]. The encoded vectors of image and the paired sentence are then concatenated to form the feature vector for classification. A binary classifier is then trained on this dataset to produce a discriminator for real/fake sentences associated with a given image.

The fact-checker can be used for report verification in inference mode. Given an automatically produced radiology report, and the corresponding input imaging study, the examiner extracts sentences from the report, and the image-sentence pair is then subjected to the same encoding process as used in training. The combined feature vector is then given to the classifier for determination of the sentence as real or fake. A revised report is assembled by removing those sentences that are deemed fake by the classifier to produce the new report.

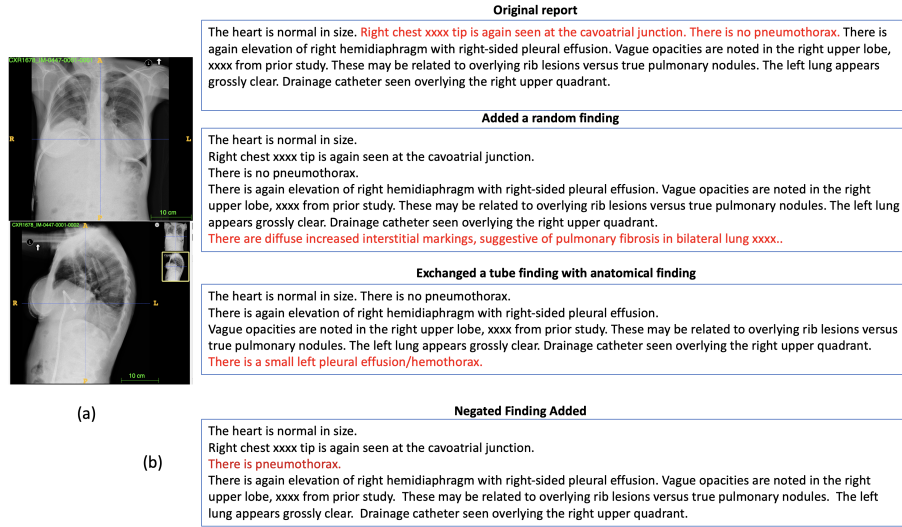
## 2 Generation of a synthetic report dataset

The key idea in synthetic report generation is to center the perturbation operations around findings described in the finding sections of reports, as these are critical to preliminary reads of imaging studies.

### 2.1 Modeling finding-related errors in automated reports

The typical errors seen in the finding sections of reports can be due to (a) addition of incorrect findings not seen in the accompanying image, (b) exchange errors, where certain findings are missed and others added, (c) reverse findings reported i.e. positive instance reported when negative instances of them are seen in image and vice versa, (d) spurious or unnecessary findings not relevant for reporting, and finally (e) incorrect description of findings in terms of fine-grained appearance, such as extent of severity, location correctness, etc.

From the point of real/fake detection, we focus on the first 3 classes of errors for synthesis as they are the most common. Let  $R = \{S_i\}$  be a ground-truthed report corresponding to an image  $I$  consisting of sentences  $\{S_i\}$  describing corresponding findings  $\{F_i\}$ . Then we can simulate a random addition of a new finding by extending the report  $R$  as  $R_a = \{S_i\} \cup \{S_a\}$  where  $S_a$  describes a new finding  $F_a \notin \{F_i\}$ . Similarly, we simulate condition (b) through an exchange of finding where one finding sentence  $S_r$  is removed to be replaced by another finding sentence  $S_a$  as  $R_e = \{S_i\} - \{S_r\} \cup \{S_a\}$ . Finally, we can simulate the replacement of positive with negative findings and vice versa to form a revised report  $R_r = \{S_i\} - \{S_p\} \cup \{S_{p'}\}$  where  $S_p$  is a sentence corresponding to a finding  $F_p$  and  $S_{p'}$  is a sentence corresponding to the finding  $F_{p'}$  which is in opposite sense of the meaning. For example, a sentence "There is pneumothorax", could be replaced by "There is no pneumothorax" to represent a reversal of polarity of the finding. Figure 2 shows examples of each of the type of operations of add, exchange and reverse findings respectively.



**Fig. 2.** Illustration of the fake reports drawn from actual reports. (a) Frontal and lateral views of a chest X-ray. (b) Corresponding original and fake radiology reports. The affected sentences during the synthesis operation are shown in red.

## 2.2 Detecting findings in sentences

Since detecting findings is key to our approach, our synthetic dataset generation focused on chest X-ray datasets, as finding detectors are well-developed for these datasets. Further, the majority of work on automated reporting has been done on chest X-rays and finding-labeled datasets are publicly available [8,9,24]. However, most of the existing approaches summarize findings at the report level. To locate findings at the sentence level, we used NLP tools such as Spacy [1] to separate sentences. We then combined ChexPert [8] labeler and NegSpacy [1] parser to extract positive and negative findings from sentences. Table 1 shows examples of findings detected in sentences. The detected findings were then validated against the ground truth labels provided at the report level in the datasets. All unique findings across reports were then aggregated into a pool  $\{F_{pool}\}$  and all unique sentences in the original reports were aggregated and mapped to their findings (positive or negative) to create the pool of sentences  $\{S_{pool}\}$ .

## 2.3 Fake report creation

For each original report  $R$  associated with an image  $I$ , we create three instances of fake reports  $R_a$ ,  $R_e$ , and  $R_r$  corresponding to the operations of addition, exchange and reversal of findings, respectively. Specifically, for creating  $R_a$  type of reports, we randomly draw from  $S_{pool}$  a sentence that contains a randomly selected finding  $F_a \notin \{F_i\}$ , where  $\{F_i\}$  are the set of findings in  $R$  (positive or negative). Similarly, to create  $R_e$ , we randomly select a finding pair  $(F_{ei}, F_{eo})$ ,

**Table 1.** Illustration of extracting findings from reports. Negated findings are shown within square brackets.

Sentences	Detected findings
There is effusion and pneumothorax.	['effusion', 'pneumothorax']
No pneumothorax, pleural effusion, but there is lobar air space consolidation.	['consolidation', [['pneumothorax', 'pleural effusion']]]
No visible pneumothorax or large pleural effusion.	[['pneumothorax', 'pleural effusion']]]
Specifically, no evidence of focal consolidation, pneumothorax, or pleural effusion.	[['focal consolidation'], [['pneumothorax', 'pleural effusion']]]
No definite focal alveolar consolidation, no pleural effusion demonstrated.	[['alveolar consolidation'], [['pleural effusion']]]

**Table 2.** Details of the fake report dataset distribution. 2557 frontal views were retained for images. 64 negative findings were retained and 114 positive findings.

Dataset	Patients	Images /Views	Reports	Pos/Neg Findings	Unique Sentences	Image-Sent Pairs	Fake Reports
Original	1786	7470/2557	2557	119/64	3850	25535	7671
Training	1071	2037	2037	68	2661	20326	4074
Testing	357	254	254	68	919	2550	508

where  $F_{ei} \in \{F_i\}$  and  $F_{eo} \in \{F_{pool}\} - \{F_i\}$ . We then remove the associated sentence with  $F_{ei}$  in  $R$  and replace it with a randomly chosen sentence associated with  $F_{eo}$  in  $\{S_{pool}\}$ . Finally, to create the reversed findings reports,  $R_r$ , we randomly select a positive or negative finding  $F_p \in \{F_i\}$ , remove its corresponding sentence and swap it with a randomly chosen sentence  $S_{p'} \in \{S_{pool}\}$ , containing findings  $F_{p'}$  that is reversed in polarity. The images, their perturbed finding, and associated sentences were recorded in each case of fake reports so that they could be used to form the pairing dataset for training the fact-checking examiner described next.

### 3 Fact-checking of AI-generated reports

#### 3.1 Fact-checking Examiner

The fact-checking examiner is a classifier using deep-learned features derived from joint image-text encodings. Specifically, since we combine images with textual sentences, we chose a feature encoding that is already trained on joint image and text pairs. In particular, we chose the CLIP joint image-text embedding model [19] to project the image and textual sentences into a common 512-length encoding space. While other joint image-text encoders could potentially work, we chose CLIP as our encoder because it was pre-trained on natural

image-text pair and subsequently tuned on radiology report-image pairs [3]. We then concatenate the image and textual embedding into a 1024-length feature vector to train a binary classifier. In our splits, the real/fake incidence distribution was relatively balanced (2:1) so that the accuracy could be used as a reliable measure of performance. We experimented with several classifiers ranging from support vector machines (SVM) to neural net classifiers. As we observed similar performance, we retained a simple linear SVM as sufficient for the task.

### 3.2 Improving the quality of reports through verification

We apply the fact-checking examiner to filter our incorrect/irrelevant sentences in automatically produced reports as shown in Figure 1 (colored in red). Specifically, given an automatically generated report for an image, we pair the image with each sentence of the report. We then use the same CLIP encoder used in training the examiner, to encode each pair of image and sentence to form a concatenated feature vector. The examiner predicted fake sentences are then removed to produce the revised report.

We develop a new measure to judge the improvement in the quality of the automatic report after applying the fact-checking examiner. Unlike popular report comparison measures such as BLEU [17], ROUGE [14] scores that perform lexical comparisons, we use a semantic similarity measure formed from encoding the reports through large language models such as SentenceBERT [20]. Specifically, let  $R = \{S_i\}$ ,  $R_{auto} = \{S_{auto}\}$ ,  $R_{corrected} = \{S_{corrected}\}$  be the original, automated, and corrected reports with their sentences respectively. To judge the improvement in quality of the report, we adopt SentenceBERT [20] to encode the individual sentences of the respective reports to produce an average encoding per report as  $E_R, E_{auto}, E_{corrected}$  respectively. Then the quality improvement score,  $QI(R)$  per triple of reports  $(R, R_{auto}, R_{corrected})$  is given by the difference in the cosine similarity between the pairwise encodings as

$$QI(R, R_{auto}, R_{corrected}) = d(E_R, E_{corrected}) - d(E_R, R_{auto}), \quad (1)$$

where  $d$  is the cosine similarity between the average encodings. This measure allows for unequal lengths of reports. A positive value indicates an improvement while a negative value indicates a worsening of the performance. The overall improvement in the quality of automatically generated reports is then given by

$$QI = (n_{positive} + n_{same} - n_{negative})/n_R \quad (2)$$


where

$$\begin{aligned} n_{positive} &= |\arg_R(d(E_R, E_{corrected}) > d(E_R, R_{auto}))| \\ n_{same} &= |\arg_R(d(E_R, E_{corrected}) = d(E_R, R_{auto}))| \\ n_{negative} &= |\arg_R(d(E_R, E_{corrected}) < d(E_R, R_{auto}))| \end{aligned} \quad (3)$$

are the number of times the corrected reports are closer to original reports, same similarity as the AI report, or worse than the AI report by applying the examiner respectively, and  $n_R$  is the total number of automated reports evaluated.

**Table 3.** Performance of real/fake discriminator.

	Prec.	Recall	F1	Sup.
Real class	0.86	0.93	0.90	3648
Accuracy	—	—	0.84	5044
Macro Avg	0.82	0.77	0.79	5044
Weighted Avg	0.84	0.84	0.84	5044

**Fig. 3.** performance of real/fake report sentence differentiation.

## 4 Results

To test our approach for fact-checking of radiology reports, we selected an open access dataset of chest X-rays from Indiana University [2] provided on Kaggle, which contains 7,470 chest X-Ray (frontal and lateral views) images with corresponding 2,557 non-duplicate reports from 1786 patients. The dataset also came with annotations documenting important findings at the report level. Of the 1786 patients, we used a (60-20-20)% patient split for training the examiner, testing the examiner, and evaluating its effectiveness in report correction respectively, thus ensuring no patient overlap between the partitions.

### 4.1 Fake report dataset created

By applying NLP methods of sentence extraction, we extracted 3850 unique sentences from radiology reports. By applying the finding extractor at the sentence level as described in Section 2.2, we catalogued a total of 119 distinct positive and 64 negative findings as shown Table 2. Using these findings and their sentences in the 2557 unique reports, and the 3 types of single perturbation operations described in Section 2.1, we generated 7,671 fake reports as shown in Table 2.

The training and test dataset for the fact-checking examiner was generated by randomly drawing sentences from sentence pool  $\{S_{pool}\}$ . Each image was first paired with each sentence from its original report and the pair was given the "Real" label. The perturbed sentence drawn from  $\{S_{pool}\}$  from the fake reports was then retrieved from each fake report and paired with the image and given the "Fake" label. The list of pairs produced were processed to remove duplicate pairings. By this process, we generated 20,326 pairs of images with real/fake sentences for training, and 2,550 pairs for testing as shown in Table 2 using 80% of the 1,786 patients.

### 4.2 Fact-checking examiner accuracy

Using the train-test splits shown in Table 2, we trained fact-checking examiner with encodings of image-sentence pairs shown in Table 2. The resulting classifier

**Table 4.** Report quality evaluation on two automatically generated report datasets.

<b>Dataset</b>	<b>Patients</b>	<b>Reports</b>	$n_{positive}$	$n_{same}$	$n_{negative}$	<b>QI score</b>
Synthetic Reports	358	3661	1105	1008	1548	15.63%
NIH Reports	198	198	60	55	83	16.1%

achieved an average accuracy of 84.2% and the AUC was 0.87 as shown in Figure 3b. The precision, recall, F-score, macro accuracy by average and weighted average methods are shown in Figure 3a. By using 10 fold cross-validation in the generation of the (60-20-20) splits for the image-report dataset, and using different classifiers provided in the Sklearn library (decision tree, logistic regression, etc.) the average accuracy lay in the range  $0.84 \pm 0.02$ . The errors seen in the classification were primarily for complex description of findings containing modifiers and negations.

### 4.3 Overall report quality improvement evaluation

We evaluated the efficacy of the fact-checking examiner on two report datasets, one synthetic with controlled "fakeness" and another dataset generated by a published algorithm described in [22]. Specifically, using the 20% partition of patients from the Indiana reports that was not used to train or test the examiner, we selected 3089 of the fake reports shown in Table 2. We evaluated the improvement in report quality using the method described in Section 3.2. These results are summarized in Table 4. Since our fake reports had only one fake sentence added, the performance improvement while still present, is modest around 15.63%.

To test the performance on automated reports generated by existing algorithms, we obtained a reference dataset consisting of freshly created reports on the NIH image dataset [24] created by radiologists as described in [25]. We retained the output of an automated report generation algorithm for the same images described in [22]. A total of 198 pairs of original and automatically created reports along with their associated imaging from the NIH dataset was used for this experiment. The results of quality improvement is shown in Table 4. As it can be seen, the quality improvement is slightly greater for reports produced by automated report extraction methods.

## 5 Conclusion

In this paper, we have proposed for the first time, an image-driven verification of automatically produced radiology reports. A dataset was carefully constructed to elicit the different types of errors produced by such methods. A Novel fact-checking examiner was developed using pairs of real and fake sentences with their corresponding imaging. The work will be extended in future to cover larger variety of defects and extended evaluation on a larger number automated reports.



## References

1. NegSpacy Parser, "<https://spacy.io/universe/project/negspacy>"
2. Demmer-Fushma, D., et al.: Preparing a collection of radiology examinations for distribution and retrieval. *Journal of American Medical Informatics Association (JAMIA)* **23**(2), 304–310 (2014)
3. Endo, M., Krishnan, R., Krishna, V., Ng, A.Y., Rajpurkar, P.: Retrieval-based chest x-ray report generation using a pre-trained contrastive language-image model. *Proceedings of Machine Learning Research* **158**, 209–219 (2021)
4. Fei-Fei, J.K.J.J.R.K.L.: A hierarchical approach for generating descriptive image paragraphs. In: *IEEE CVPR* (2017)
5. Gale, W., Oakden-Rayner, L., Carneiro, G., Bradley, A.P., Palmer, L.J.: Producing radiologist-quality reports for interpretable artificial intelligence. *arXiv preprint arXiv:1806.00340* (2018)
6. Grewal, H., Dhillon, G., Monga, V., Sharma, P., Buddhavarapu, V.S., Sidhu, G., Kashyap, R., Grewal, H., MD, M.G.D., Monga, V., Sharma, P., Buddhavarapu, V.S., Sidhu, G., MBBS, M.R.K.: Radiology gets chatty: The chatgpt saga unfolds. *Cureus* **15** (6 2023). <https://doi.org/10.7759/CUREUS.40135>, <https://www.cureus.com/articles/161200-radiology-gets-chatty-the-chatgpt-saga-unfolds>
7. Guo, J., Lu, S., Cai, H., Zhang, W., Yu, Y., Wang, J.: Long text generation via adversarial training with leaked information. In: *AAAI-2018*. pp. 5141–5148 (2018)
8. 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: *Thirty-Third AAAI Conference on Artificial Intelligence* (2019)
9. Johnson, A.E.W., Pollard, T.J., Berkowitz, S.J., Greenbaum, N.R., Lungren, M.P., y. Deng, C., Mark, R.G., Horng, S.: Mimic-cxr: A large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042* (2019)
10. Lab, N.J.: Ai will start fact-checking. we may not like the results., <https://www.niemanlab.org/2022/12/ai-will-start-fact-checking-we-may-not-like-the-results/>
11. Li, C.Y., Liang, X., Hu, Z., Xing, E.P.: Knowledge-driven encode, retrieve, paraphrase for medical image report generation. *arXiv preprint arXiv:1903.10122* (2019)
12. Li, X., Zhang, L., Wu, Z., Liu, Z., Zhao, L., Yuan, Y., Liu, J., Li, G., Zhu, D., Yan, P., et al.: Artificial general intelligence for medical imaging. *arXiv preprint arXiv:2306.05480* (2023)
13. Li, Y., Liang, X., Hu, Z., Xing, E.P.: Hybrid retrieval-generation reinforced agent for medical image report generation. In: *Advances in Neural Information Processing Systems*. pp. 1530–1540 (2018)
14. Lin, C.Y.: Rouge: a package for automatic evaluation of summaries. In: *Workshop on Text Summarization Branches Out* (2004)
15. Liu, G., et al.: Clinically accurate chest x-ray report generation. *arXiv:1904.02633v* (2019)
16. Pang, T., Li, P., Zhao, L.: A survey on automatic generation of medical imaging reports based on deep learning. *BioMedical Engineering OnLine* **22**(1), 48 (2023). <https://doi.org/10.1186/s12938-023-01113-y>, <https://doi.org/10.1186/s12938-023-01113-y>

17. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics. pp. 311–318. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA (Jul 2002). <https://doi.org/10.3115/1073083.1073135>, <https://aclanthology.org/P02-1040>
18. Passi, K., Shah, A.: Distinguishing fake and real news of twitter data with the help of machine learning techniques. ACM International Conference Proceeding Series pp. 1–8 (8 2022). <https://doi.org/10.1145/3548785.3548811>
19. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. Proceedings of Machine Learning Research **139**, 8748–8763 (2 2021), <https://arxiv.org/abs/2103.00020v1>
20. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. CoRR **abs/1908.10084** (2019), <http://arxiv.org/abs/1908.10084>
21. Suprem, A., Pu, C.: Midas: Multi-integrated domain adaptive supervision for fake news detection (2022), <https://arxiv.org/pdf/2205.09817.pdf>
22. Syeda-Mahmood, T., Wong, K.C.L., Gur, Y., Wu, J.T., Jadhav, A., Kashyap, S., Karargyris, A., Pillai, A., Sharma, A., Syed, A.B., Boyko, O., Moradi, M.: Chest x-ray report generation through fine-grained label learning. In: MICCAI-2020 (2020)
23. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3156–3164 (2015)
24. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases (2017), <https://uts.nlm.nih.gov/metathesaurus.html>
25. Wu, J.T., Wong, K.C., Gur, Y., Ansari, N., Karargyris, A., Sharma, A., Morris, M., Saboury, B., Ahmad, H., Boyko, O., Syed, A., Jadhav, A., Wang, H., Pillai, A., Kashyap, S., Moradi, M., Syeda-Mahmood, T.: Comparison of chest radiograph interpretations by artificial intelligence algorithm vs radiology residents. JAMA Network Open **3**, e2022779–e2022779 (10 2020). <https://doi.org/10.1001/JAMANETWORKOPEN.2020.22779>, <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2771528>
26. Xiong, Y., Du, B., Yan, P.: Reinforced transformer for medical image captioning. In: Machine Learning in Medical Imaging: 10th International Workshop, MLMI 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13, 2019, Proceedings 10. pp. 673–680. Springer (2019)
27. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: International conference on machine learning. pp. 2048–2057 (2015)