

RepsNet: Combining Vision with Language for Automated Medical Reports

Ajay K. Tanwani, Joelle Barral, and Daniel Freedman

Verily & Google Research, USA

{ajaytanwani, jbaral, danielfreedman}@google.com

Abstract. Writing reports by analyzing medical images is error-prone for inexperienced practitioners and time consuming for experienced ones. In this work, we present RepsNet that adapts pre-trained vision and language models to interpret medical images and generate automated reports in natural language. RepsNet consists of an encoder-decoder model: the encoder aligns the images with natural language descriptions via contrastive learning, while the decoder predicts answers by conditioning on encoded images and prior context of descriptions retrieved by nearest neighbour search. We formulate the problem in a visual question answering setting to handle both categorical and descriptive natural language answers. We perform experiments on two challenging tasks of medical visual question answering (VQA-Rad) and report generation (IU-Xray) on radiology image datasets. Results show that RepsNet outperforms state-of-the-art methods with 81.08% classification accuracy on VQA-Rad 2018 and 0.58 BLEU-1 score on IU-Xray. Supplementary details are available at: <https://sites.google.com/view/repsnet>

Keywords: vision and language · visual question answering · report generation.

1 Introduction

A long standing goal in artificial intelligence is to seamlessly interpret and describe medical images/videos with natural language. In this paper, we combine both vision and language modalities to interpret medical images in a visual question answering (VQA) setting, whereby we predict the answer to a given image and question using a novel encoder-decoder model (see Fig. 1). We present RepsNet that fuses the encoded image and question features by contrastive alignment, while the decoder learns the conditional probability distribution to generate descriptions with: 1) encoded image and question features, and 2) prior context retrieved from nearest neighbouring reports of the image. We leverage publicly available ResNeXt [37] and BERT [10] for warm-starting the encoder, and GPT-2 [28] as the base model for the natural language decoder.

We present its application to assist practitioners with automatic report generation from medical images [16,7,20,22,25]. Existing methods using hand-written

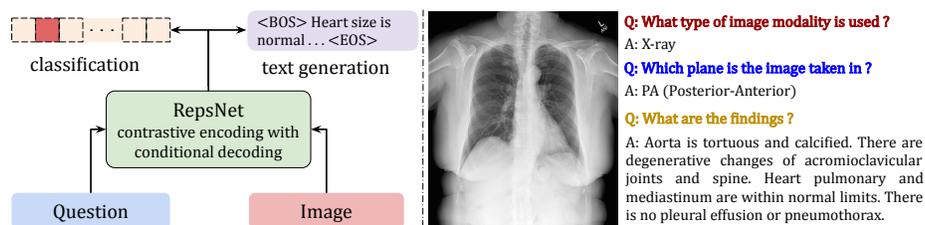


Fig. 1. (left) RepsNet analyzes medical images and automates report writing by providing answers to questions via classifying among known answer categories or generating natural language descriptions, (right) radiology report generation example with top two categorical answers and bottom one natural language descriptive answer.

notes, dictation services or electronic medical record templates are widely perceived to be time-consuming and cumbersome. To this end, we parse the medical report into a set of questions and handle both categorical (yes/no, multiple choices) and natural language descriptive answers (open-ended) in a visual question answering setting. We evaluate the proposed approach on two publicly available benchmark datasets: 1) visual question answering radiology (VQA-Rad) datasets in the span of 2018 – 2021 [19], 2) Indiana University x-ray (IU-Xray) dataset containing chest x-ray images paired with reports describing findings and impressions [8]. RepsNet outperforms state-of-the-art models across both VQA-Rad and IU-Xray datasets.

Contributions: This paper makes three contributions:

- We present RepsNet, an encoder-decoder model for writing reports that adapts pretrained models by contrastive alignment of images with answers in the encoding phase, and generates natural language descriptions by conditional decoding on images and prior context of retrieved reports.
- A visual question answering formulation to handle both categorical and natural language descriptive answers in generating automated reports.
- Experiments on publicly available VQA-Rad and IU-Xray datasets with 81.08% classification accuracy and 0.58 BLEU-1 score respectively, showing significant performance improvement over state-of-the-art methods.

2 Related Work

Vision and Language Pretraining: Self-supervised pre-training of language models such as BERT [10], GPT/GPT-2 [28], XLNet have shown promising results in transferring knowledge across related tasks [12]. This has led to combining both visual and language modalities by cross-alignment of domains in a joint embedding space [9,30]. Examples include LXMERT [35], ViLBERT [24], PixelBERT [14], VideoBERT [34] and VisualGPT [36]. Authors in [43,27] use contrastive learning to pair images with textual descriptions as a whole, in contrast to grounding the masked works in the image locally in [33,13]. To incorporate prior knowledge in pretrained language generation models [39], Ziegler *et al.* [44] adapt a pretrained model for arbitrary source conditioning. Despite a few

promising approaches, cross-domain conditioning of a pretrained model remains a challenge and can degrade the pretrained model representations.

Visual Question Answering and Image Captioning: Describing medical images with visual question answering [4] or natural language [38,3] is difficult due to rare and diverse nature of abnormalities, weak association of image features with text in reports, lack of prior domain knowledge, case-based reasoning, and long descriptions of findings. Medical VQA has recently received attention with small scale datasets such as VQA-Rad to categorize answers by classification [5,11,21]. Several works have followed the image captioning line of work with an emphasis on generating long descriptions [16,2], incorporating domain-specific medical knowledge [42,22], retrieving description from a template [20,23], question answering [29,32], among others [40,25].

In this paper, we investigate the automated report writing under a novel visual question answering framework to handle both categorical and descriptive answers for a given image and a question. We use contrastive learning to align the paired images and report answers in an embedding space, and retrieve nearest neighbour report answers to incorporate prior knowledge in generating medical descriptions.

3 RepsNet: Proposed Approach

Problem Formulation: Given an image or a set of images $\mathbf{x} \in \mathbf{X}$, we are interested in generating a report comprising of s answers $\mathbf{y} = \{\mathbf{y}_1 \dots \mathbf{y}_s\} \in \mathbf{Y}$, corresponding to the natural language questions $\mathbf{q} = \{\mathbf{q}_1 \dots \mathbf{q}_s\} \in \mathbf{Q}$. Each answer \mathbf{y}_i may be **close-ended** belonging to a fixed possible set of categories or **open-ended** comprising of multiple natural language sentences. Each word $\mathbf{w} \in \mathbf{V}$ in the open-ended answer belongs to a known natural language vocabulary. We seek to learn the model parameters Θ to maximize the conditional likelihood $\mathcal{P}_\Theta(\mathbf{y}_i | \mathbf{x}, \mathbf{q}_i)$ of predicting the answers for a given image and a set of questions,

$$\Theta = \arg \max_{\Theta} \sum_{i=1}^s \log \mathcal{P}_\Theta(\mathbf{y}_i | \mathbf{x}, \mathbf{q}_i). \quad (1)$$

We formulate the problem with an encoder decoder model. The encoder $f_{\theta_{\text{enc}}} : \{\mathbf{X}, \mathbf{Q}\} \rightarrow \{\bar{\mathbf{X}}, \bar{\mathbf{Q}}\} \in \mathbb{R}^{\{n_x, n_q\} \times \{d_x, d_q\}}$ transforms the image and the input text sequence to a joint cross-aligned visual and language representation space with n_x image pixels/regions, n_q text tokens, and $\{d_x, d_q\}$ hidden space dimensions of image and text embeddings respectively. The decoder $h_{\theta_{\text{dec}}} : \{\bar{\mathbf{X}}, \bar{\mathbf{Q}}, \bar{\mathbf{C}}\} \rightarrow \mathcal{P}(\mathbf{Y})$ models the conditional probability distribution of predicting the target answer \mathbf{Y} given the encoded hidden states $\{\bar{\mathbf{X}}, \bar{\mathbf{Q}}\}$, and the **prior context** $\bar{\mathbf{C}} \in \mathbb{R}^{n_c \times d_c}$ of n_c tokens with dimension d_c that represents the domain specific knowledge for controlled text generation (we discuss prior context further in the next section). Note that we only use the prior context for generating open-ended answers.

In this paper, we leverage large-scale pretrained models for warm-starting the encoder and the decoder model parameters. For close-ended answers, we

map the combined image and question features to the output layer of all possible close-ended answers for classification. For open-ended answers, the decoder retrieves the prior context $\bar{\mathbf{C}}$ as the nearest neighbouring answers of the encoded image features, and greedily maximizes the learned conditional distribution $\mathcal{P}_{\theta_{\text{dec}}}(\mathbf{Y}_t | \mathbf{Y}_{0:t-1}, \bar{\mathbf{X}}, \bar{\mathbf{Q}}, \bar{\mathbf{C}})$ to generate the answer sequence $\mathbf{Y}_{1:t}$ in an auto-regressive manner (see Fig. 2).

3.1 Contrastive Image-Text Encoder

The encoder has four constituent parts: 1) **image encoder** to extract visual features, 2) **text encoder** to tokenize and contextualize natural language questions and answers features, 3) **bilinear attention network** to fuse the image and the question, and 4) **contrastive alignment** of visual features and textual answers.

Image Encoder: We use the ResNeXt-101 [37] architecture as the base image encoder. We remove the last linear and pooling layer and add a 2D adaptive average pooling layer to resize the input image to a fixed feature space of $14 \times 14 \times 2048$ that preserves the correspondence between the visual features and the input image ($n_x = 196, d_x = 2048$). Moreover, we add image transformations, namely *color jittering, normalization, random erasing*, to augment the training data distribution within each batch before extracting the visual features.

Text Encoder: We adapt the BERT [10] model for the text encoder that is pre-trained to predict masked words locally based on the context provided by other non-masked words in the sequence. We filter out the punctuation marks and tokenize the text using WordPiece algorithm [10], before extracting the textual features.

Bilinear Attention Network (BAN): We use a BAN to fuse the cross-modal encoded question and image features [18]. The outer product or the bilinear product exhaustively combines the multi-modal features at the cost of higher computational complexity; in comparison to naive concatenation or inner product between the features. Compared to other co-attention mechanisms, BAN exploits bilinear interaction maps where each feature is pooled by low-rank bilinear approximations. Residual learning on top combines multiple bilinear attention maps for effective joint representation of question and image features.

For the sake of brevity and a slight abuse of notation, we denote $\bar{\mathbf{X}}$ for both the image and the combined (image and question) features in describing the rest of the encoder and the decoder sections.

Contrastive Vision and Text Learning: We align images with natural language descriptions via bidirectional contrastive learning [6], that pulls together a given image-answer pair, while pushing away observations that correspond to different image-answer pairs.

Given the encoded image (and question) $\bar{\mathbf{X}}$ and the natural language answer features $\bar{\mathbf{Y}} \in \mathbb{R}^{n_y \times d_y}$ with n_y tokens of dimension d_y , we first project them to a d -dimensional space with a linear transformation to $\hat{\mathbf{X}} \in \mathbb{R}^d$ and $\hat{\mathbf{Y}} \in \mathbb{R}^d$. During training, the loss operates on a mini-batch of N_T image-text pairs $\{\hat{\mathbf{x}}_i, \hat{\mathbf{y}}_i\}_{i=1}^{N_T}$.

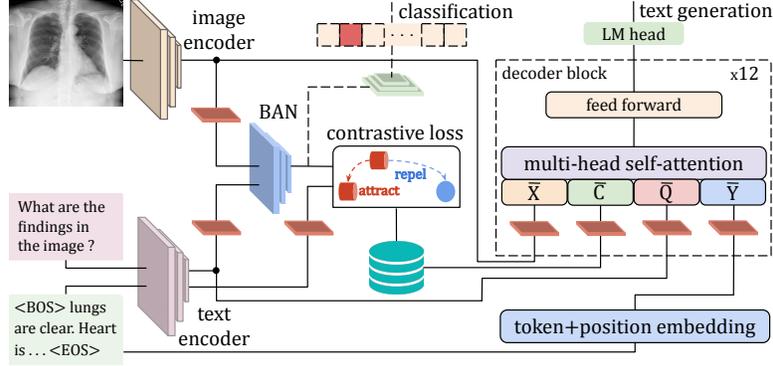


Fig. 2. RepNet encoded image and question features are fused via bilinear attention network (BAN), before self-supervised contrastive alignment with natural language descriptions. The answer is categorized via classification among fixed answer categories or generated by conditional language decoding on image, question and prior context of answers retrieved by nearest neighbour search. Note that we omit the question features \bar{Q} in describing the conditional language decoder below for brevity.

where each pair is in turn taken as a positive sample to maximize agreement against all other negative samples, i.e.,

$$\mathcal{L}_{\hat{\mathbf{x}} \rightarrow \hat{\mathbf{y}}} = -\frac{1}{N_T} \sum_{i=1}^{N_T} \log \frac{\exp(\langle \hat{\mathbf{x}}_i, \hat{\mathbf{y}}_i \rangle / \tau)}{\sum_{j=1}^{N_T} \exp(\langle \hat{\mathbf{x}}_i, \hat{\mathbf{y}}_j \rangle / \tau)}, \quad (2)$$

where $\langle \hat{\mathbf{x}}, \hat{\mathbf{y}} \rangle = \frac{\hat{\mathbf{x}}^\top \hat{\mathbf{y}}}{\|\hat{\mathbf{x}}\| \|\hat{\mathbf{y}}\|}$ represents the cosine similarity distance and $\tau \in \mathbb{R}^+$ represents the temperature parameter to scale the similarity metric. Similar to the image-to-text loss in Eq. (2), we also define the text-to-image loss $\mathcal{L}_{\hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}}$ to account for the asymmetry with respect to each input modality as in [43,27]. Overall bidirectional encoder loss \mathcal{L}_{enc} is the sum of the two constituent contrastive losses weighted by constant $\alpha_l \in \mathbb{R}^+$,

$$\mathcal{L}_{\text{enc}} = \alpha_l (\mathcal{L}_{\hat{\mathbf{x}} \rightarrow \hat{\mathbf{y}}} + \mathcal{L}_{\hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}}). \quad (3)$$

Prior Context Knowledge: We store the normalized natural language answers of the train set $\hat{\mathbf{Y}}_{\text{train}}$ during model training. We then compute topk nearest neighbours $\bar{\mathbf{C}}$ that maximize the cosine similarity between a given encoded image $\hat{\mathbf{X}}$ and the stored natural language answers $\hat{\mathbf{Y}}_{\text{train}}$. We use the FAISS library for scalable nearest neighbour search [17]. The prior context aids the decoder to attend to longer horizon dependencies and get additional case-based details for controlled text generation. This is particularly relevant in describing medical images with specific terminologies, writing style and class imbalanced abnormalities, i.e.,

$$\bar{\mathbf{C}} = \text{topk} \left[\max_{i \in \hat{\mathbf{Y}}_{\text{train}}} \langle \hat{\mathbf{X}}, \hat{\mathbf{Y}}_{\text{train}}^{(i)} \rangle \right]. \quad (4)$$

3.2 Conditional Language Decoder

The probability distribution of generating the output text sequence $\mathbf{Y}_{1:t}$ conditioned on the contextualized encoding sequence $\mathcal{P}_{\theta_{\text{dec}}}(\mathbf{Y}_{1:t}|\bar{\mathbf{X}}, \bar{\mathbf{C}})$ can be decomposed into a product of conditional distributions using Bayes’ rule,

$$\mathcal{P}_{\theta_{\text{dec}}}(\mathbf{Y}_{1:t}|\bar{\mathbf{X}}, \bar{\mathbf{C}}) = \prod_{i=1}^t \mathcal{P}_{\theta_{\text{dec}}}(y_i|y_{0:i-1}, \bar{\mathbf{X}}, \bar{\mathbf{C}}), \quad (5)$$

where $\mathbf{y}_0 = \langle \mathbf{BOS} \rangle$ is a special token reserved for the beginning of a sentence. We model the conditional language generation with a stack of transformer-based blocks, using the GPT-2 model as the base pretrained language decoder [28]. We introduce modifications to the GPT-2 model for conditioning on image and prior context features by directly adding their attention outputs to the pretrained self-attention layers of the model, similar to [44,2], thereby adding the attention outputs for different conditional inputs with a parsimonious increase in the number of parameters only (see supplementary materials for details). The conditional probability distribution in Eq. (5) is maximized by optimizing the cross-entropy loss on the ground-truth and the predicted sequences.

Overall Approach: During training, we adapt the pretrained language and vision models in an end-to-end manner for contrastive encoding and conditional decoding with a small amount of image-text pairs. The overall training loss function comprises of the contrastive loss and the cross-entropy loss. During natural language generation, we predict the output sequence in an auto-regressive manner with greedy or beam search decoding, and stop generating the sequence once we predict a special end of text ($\langle \mathbf{EOS} \rangle$) token.

4 Experiments, Results and Discussion

We evaluate the performance of RepNet in interpreting visual concepts on publicly available VQA-Rad [19] for classification and IU-Xray [8] for natural language generation. We are interested in evaluating: 1) how feasible it is to adapt the pretrained language and vision models in describing small set of medical images, 2) what is the role of contrastive encoding in learning joint visual linguistic representations, 3) does conditional decoding on image features and prior context help with generating medical language descriptions, and 4) how does RepNet fare in performance among the state-of-the-art approaches.

4.1 Visual Question Answering

VQA-Rad Dataset: We use the VQA-Rad datasets [19] from 2018 – 2019 and also introduce an aggregated dataset, VQA-Rad All, that combines all the VQA-Rad datasets from 2018 – 2021. Radiology images in the datasets are taken from open-access MedPix database, and the questions are predominantly posed from categories, such as image plane, imaging modality, organ system involved

Table 1. Classification accuracy on the **Table 2.** BLEU scores (B1 - B4) for medical report generation on IU-Xray dataset. Bottom three rows increase minimum occurrence threshold RepsNet yields better scores than other from 0 to 5 to 10 instances. RepsNet outperforms all other competing methods.

	2018	2019	All		B1	B2	B3	B4
MEVF [5]	66.10	–	–	Co-Att [16]	0.45	0.29	0.20	0.15
MMQ [11]	67.00	–	–	HRGR [20]	0.44	0.30	0.21	0.15
QCR [41]	69.65	–	–	CMAS [15]	0.46	0.30	0.21	0.15
CLEF [1]	–	62.40	–	Mem-T [7]	0.47	0.30	0.22	0.16
CRPD [21]	72.70	–	–	VTI [25]	0.49	0.36	0.29	0.15
RepsNet-0	81.08	67.57	63.69	PPKED [22]	0.48	0.31	0.22	0.17
RepsNet-5	83.55	79.83	71.93	RepsNet	0.58	0.44	0.32	0.27
RepsNet-10	87.05	81.17	80.37					

and image abnormalities. The VQA problem is posed as a multi-class classification over all possible set of answers, and classification accuracy on evaluation set is used as the performance metric. We use the standard training and evaluation splits provided with the datasets (see summary in supplementary materials).

Results: Table 1 shows that RepsNet outperforms all other competing methods across all the datasets. Similar to other methods, RepsNet uses bilinear attention mechanism to fuse the image and the question features. Contrary to other methods, RepsNet does not use fixed Glove word embeddings [26] or RNNs for sentence level representations; instead it learns the entire contextual embeddings using BERT-style transformer with WordPiece tokenization. Ablation study in Table 3 also shows that the performance increases the most with the use of pre-trained models. We observe from Table 1 that simply filtering out instances and class categories with less than 5 and 10 instances per class category $M_o = \{5, 10\}$ proportionally increases the classification accuracy across all datasets, at the cost of reducing the overall number of instances and class categories to mitigate class imbalance in the datasets. Note that we do not take into account unseen class category instances of the evaluation set in computing the classification accuracy.

4.2 Medical Report Generation

IU-XRay: The Indiana University x-ray dataset [8] comprises of frontal and lateral views of chest x-ray images that are associated with radiology report sections, namely *impressions*, *findings* and *manual tags*. For brevity, we only report results for populating the findings question in this work, i.e., we associate the same question for all the answers. After omitting the reports without findings section, we randomly split the remaining 3607 reports into 80% training and 20% evaluation sets. On average, each report instance has 5.7 sentences, while each sentence has 6.5 words that describe the image findings. Note that no classification labels are available to detect the anomalies. Max number of tokens for a report section is set to 200 and the report findings are zero-padded in case

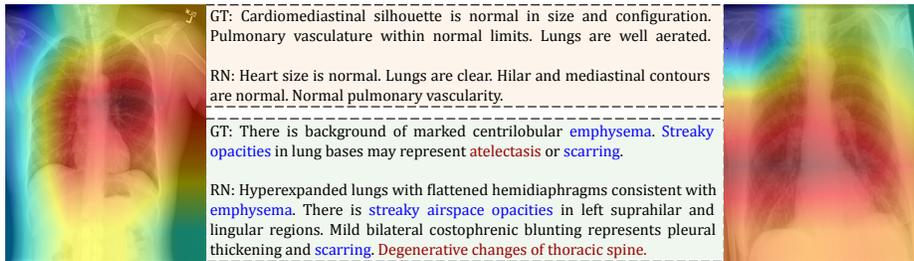


Fig. 3. Heatmap visualization and comparison between ground-truth (GT) and RepsNet generated (RN) report of: (left) normal case, (right) abnormal case. RepsNet shows strong alignment with ground-truth in describing medical findings. Text in blue shows abnormalities, text in red represents misalignment.

Table 3. Ablation study on VQA-Rad dataset to quantify the effect of pre-training, pre-processing and contrastive learning. Classification accuracy increases the most with pre-training, while pre-processing and contrastive learning stage further improve the performance.

pretraining		✓	✓	✓
preprocess			✓	✓
contrastive				✓
	74.47	79.12	80.09	81.08

Table 4. Ablation study on IU-Xray dataset with: (top) visual features - Vis, (middle) Vis with contrastive encoding - Vis + CE, and (bottom) Vis with CE and prior context - Vis + CE + PC. BLEU scores improve with contrastive learning and prior context.

	B1	B2	B3	B4
Vis	0.48	0.38	0.30	0.26
Vis + CE	0.55	0.42	0.32	0.27
Vis + CE + PC	0.58	0.44	0.32	0.27

its length is less than max number of tokens. We use the sentence level BLEU scores as the performance metric computed using the `nltk` library that compares n -gram similarity between the ground-truth and the generated report where n varies from 1 to 4 (whereas for classification accuracy evaluation in VQA-Rad datasets, we compare the predicted and the ground-truth indices of the class categories).

Results: Results are summarized in Table 2. It can be seen that RepsNet performs significantly better than the state-of-the-art report generation methods across all the BLEU scores, suggesting the feasibility of adapting large-scale pretrained language and vision models on a small set of domain-specific medical data. Ablation study in Table 4 reveals that adding visual features, contrastive learning and prior context subsequently boosts the performance of the GPT2-decoder. Fig. 3 provides a qualitative comparison between the ground-truth and the generated report findings, along with the heatmap visualizations using grad-cam [31] for an intuitive understanding of the approach. We observe a strong alignment in generating normal report findings, whereas part of the findings sometimes get omitted and/or added in describing the abnormalities, especially for rare cases (see supplementary materials for video demonstration and other examples). Systematic dealing of rare cases with external domain knowledge and past medical history of patients is a promising direction of our future work. We are also interested in incorporating attention mechanisms for conditional

visualization of generated text on image patches as a measure of uncertainty in the prediction. Making these reports self-explainable is critical for its wider adoption. Other areas of interest include reducing the liability of generated report errors, as well as working with medical experts to evaluate the generated reports.

5 Conclusion

In this paper, we have presented RepsNet that adapts pre-trained vision and language models for describing a small set of domain-specific medical images. We take a unified visual question answering approach to predict class categories or generate descriptive answers for writing automated medical reports. RepsNet is specifically tailored for contrastive alignment of images and text in the encoding phase, and combining visual and prior context of nearest neighboring reports with natural language generator in the decoding phase. This has enabled RepsNet to provide state-of-the-art results on challenging tasks of visual question answering and medical report generation on radiology images. In our future work, we plan to extend our approach to summarizing reports from videos, and transfer the developed methodology to clinical sites for automated reporting in gastroenterology.

References

1. Abacha, A.B., Hasan, S.A., Datla, V., Liu, J., Demner-Fushman, D., Müller, H.: Vqa-med: Overview of the medical visual question answering task at imageclef 2019. In: CLEF (2019)
2. Alfarghaly, O., Khaled, R., Elkorany, A., Helal, M., Fahmy, A.: Automated radiology report generation using conditioned transformers. *Informatics in Medicine Unlocked* **24**, 100557 (2021)
3. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and VQA. *CoRR* **abs/1707.07998** (2017)
4. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: VQA: visual question answering. *CoRR* **abs/1505.00468** (2015)
5. Binh Nguyen, Thanh-Toan Do, B.N.T.D.E.T.Q.T.: Overcoming data limitation in medical visual question answering. In: MICCAI (2019)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. *CoRR* **abs/2002.05709** (2020)
7. Chen, Z., Song, Y., Chang, T., Wan, X.: Generating radiology reports via memory-driven transformer. *CoRR* **abs/2010.16056** (2020)
8. Demner-Fushman, D., et al.: Preparing a collection of radiology examinations for distribution and retrieval. *J. Am. Medical Informatics Assoc.* **23**(2), 304–310 (2016)
9. Desai, K., Johnson, J.: Virtex: Learning visual representations from textual annotations. *CoRR* **abs/2006.06666** (2020)
10. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR* **abs/1810.04805** (2018)
11. Do, T., Nguyen, B.X., Tjiputra, E., Tran, M., Tran, Q.D., Nguyen, A.: Multiple meta-model quantifying for medical visual question answering. In: MICCAI (2021)

12. Dong, L., et al.: Unified language model pre-training for natural language understanding and generation. In: *NeurIPS*. vol. 32 (2019)
13. Gupta, T., Vahdat, A., Chechik, G., Yang, X., Kautz, J., Hoiem, D.: Contrastive learning for weakly supervised phrase grounding. *CoRR* **abs/2006.09920** (2020)
14. Huang, Z., Zeng, Z., Liu, B., Fu, D., Fu, J.: Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *CoRR* **abs/2004.00849** (2020)
15. Jing, B., Wang, Z., Xing, E.P.: Show, describe and conclude: On exploiting the structure information of chest x-ray reports. *CoRR* **abs/2004.12274** (2020)
16. Jing, B., Xie, P., Xing, E.P.: On the automatic generation of medical imaging reports. *CoRR* **abs/1711.08195** (2017), <http://arxiv.org/abs/1711.08195>
17. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734* (2017)
18. Kim, J.H., Jun, J., Zhang, B.T.: Bilinear attention networks. In: *Advances in Neural Information Processing Systems*. vol. 31 (2018)
19. Lau, J.J., Gayen, S., Ben Abacha, A., Demner-Fushman, D.: A dataset of clinically generated visual questions and answers about radiology images. *Nature Scientific Data* **5** (2018)
20. Li, C.Y., Liang, X., Hu, Z., Xing, E.P.: Hybrid retrieval-generation reinforced agent for medical image report generation. *CoRR* **abs/1805.08298** (2018)
21. Liu, B., Zhan, L.M., Wu, X.M.: Contrastive pre-training and representation distillation for medical visual question answering based on radiology images. In: *MICCAI*. pp. 210–220 (2021)
22. Liu, F., Wu, X., Ge, S., Fan, W., Zou, Y.: Exploring and distilling posterior and prior knowledge for radiology report generation. In: *CVPR*. pp. 13753–13762 (2021)
23. Liu, G., et al.: Clinically accurate chest x-ray report generation. *CoRR* **abs/1904.02633** (2019)
24. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *CoRR* **abs/1908.02265** (2019)
25. Najdenkoska, I., Zhen, X., Worring, M., Shao, L.: Variational topic inference for chest x-ray report generation. *CoRR* **abs/2107.07314** (2021)
26. Pennington, J., Socher, R., Manning, C.: GloVe: Global vectors for word representation. In: *EMNLP*. pp. 1532–1543 (2014)
27. 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. *CoRR* **abs/2103.00020** (2021)
28. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners (2019)
29. Ren, F., Zhou, Y.: Cgmvcqa: A new classification and generative model for medical visual question answering. *IEEE Access* **8**, 50626–50636 (2020)
30. Sariyildiz, M.B., Perez, J., Larlus, D.: Learning visual representations with caption annotations. *CoRR* **abs/2008.01392** (2020)
31. Selvaraju, R.R., et al.: Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization. *CoRR* **abs/1610.02391** (2016)
32. Sharma, D., Purushotham, S., Reddy, C.K.: MedFuseNet: an attention-based multimodal deep learning model for visual question answering in the medical domain. *Scientific Reports* **11**(1) (2021). <https://doi.org/10.1038/s41598-021-98390-1>
33. Sun, C., Baradel, F., Murphy, K., Schmid, C.: Contrastive bidirectional transformer for temporal representation learning. *CoRR* **abs/1906.05743** (2019)
34. Sun, C., Myers, A., Vondrick, C., Murphy, K., Schmid, C.: Videobert: A joint model for video and language representation learning. *CoRR* **abs/1904.01766** (2019)

35. Tan, H., Bansal, M.: LXMERT: learning cross-modality encoder representations from transformers. CoRR **abs/1908.07490** (2019)
36. Xia, Q., et al.: XGPT: cross-modal generative pre-training for image captioning. CoRR **abs/2003.01473** (2020)
37. Xie, S., Girshick, R.B., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. CoRR **abs/1611.05431** (2016)
38. 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: ICML. vol. 37, pp. 2048–2057 (2015)
39. Yu, W., Zhu, C., Li, Z., Hu, Z., Wang, Q., Ji, H., Jiang, M.: A survey of knowledge-enhanced text generation. CoRR **abs/2010.04389** (2020)
40. Yuan, J., et al.: Automatic radiology report generation based on multi-view image fusion and medical concept enrichment. In: MICCAI. pp. 721–729 (2019)
41. Zhan, L.M., Liu, B., Fan, L., Chen, J., Wu, X.M.: Medical visual question answering via conditional reasoning. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 2345–2354 (2020)
42. Zhang, Y., Wang, X., Xu, Z., Yu, Q., Yuille, A., Xu, D.: When radiology report generation meets knowledge graph. In: AAAI. vol. 34, pp. 12910–12917 (2020)
43. Zhang, Y., Jiang, H., Miura, Y., Manning, C.D., Langlotz, C.P.: Contrastive learning of medical visual representations from paired images and text. CoRR **abs/2010.00747** (2020)
44. Ziegler, Z.M., Melas-Kyriazi, L., Gehrmann, S., Rush, A.M.: Encoder-agnostic adaptation for conditional language generation. CoRR **abs/1908.06938** (2019)