



## King's Research Portal

[Link to publication record in King's Research Portal](#)

*Citation for published version (APA):*

Wang, J., Bhalerao, A., Yin, T., See, S., & He, Y. (2022). *CAMANet: Class Activation Map Guided Attention Network for Radiology Report Generation*.

### **Citing this paper**

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.


### **General rights**

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

### **Take down policy**

If you believe that this document breaches copyright please contact [librarypure@kcl.ac.uk](mailto:librarypure@kcl.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# CAMANet: Class Activation Map Guided Attention Network for Radiology Report Generation

Jun Wang<sup>a</sup>, Abhir Bhalerao<sup>a,\*</sup>, Terry Yin<sup>b</sup>, Simon See<sup>b</sup>, Yulan He<sup>a,c,d</sup>

<sup>a</sup>Department of Computer Science, University of Warwick, UK

<sup>b</sup>NVIDIA AI Tech Centre

<sup>c</sup>Department of Informatics, King's College London, UK

<sup>d</sup>The Alan Turing Institute, UK

## ARTICLE INFO

### Article history:

Received 1 May 2013

Received in final form 10 May 2013

Accepted 13 May 2013

Available online 15 May 2013

Communicated by S. Sarkar

### Keywords:

Radiology report generation,

Cross-modal alignment,

Class activation map

## ABSTRACT

Radiology report generation (RRG) has gained increasing research attention because of its huge potential to mitigate medical resource shortages and aid the process of disease decision making by radiologists. Recent advancements in Radiology Report Generation (RRG) are largely driven by improving models' capabilities in encoding single-modal feature representations, while few studies explore explicitly the cross-modal alignment between image regions and words. Radiologists typically focus first on abnormal image regions before they compose the corresponding text descriptions, thus cross-modal alignment is of great importance to learn an abnormality-aware RRG model. Motivated by this, we propose a **Class Activation Map** guided **Attention Network** (CAMANet) which explicitly promotes cross-modal alignment by employing the aggregated class activation maps to supervise the cross-modal attention learning, and simultaneously enriches the discriminative information. Experimental results demonstrate that CAMANet outperforms previous SOTA methods on two commonly used RRG benchmarks.<sup>1</sup>

© 2022 Elsevier B. V. All rights reserved.

## 1. Introduction

Radiology report generation (RRG) aims to automatically describes radiology images, e.g., X-Ray and MRI, by human-like language. Generating a coherent report requires expertise from radiologists, who are however among the most in-demand physician specialists in most countries. Moreover, it can take at least 5 minutes to describe a radiology image even for a professional radiologist (Jing et al., 2018). Consequently, there had been growing interest in automating RRG because of its huge potential to efficiently and effectively assist the diagnosis process.

RRG is a challenging task and with the availability of large-scale datasets and newly developed high-performance computer vision and language models, some valuable insights and improvements have been recently reported (Miura et al., 2021; Nishino et al., 2020; Liu et al., 2019). The performance however is still far from satisfactory for methods to be deployed in practice. This is

<sup>1</sup>Our source code and trained models will be made available upon paper acceptance.

\*Corresponding author:

e-mail: [abhir.bhalerao@warwick.ac.uk](mailto:abhir.bhalerao@warwick.ac.uk) (Abhir Bhalerao)

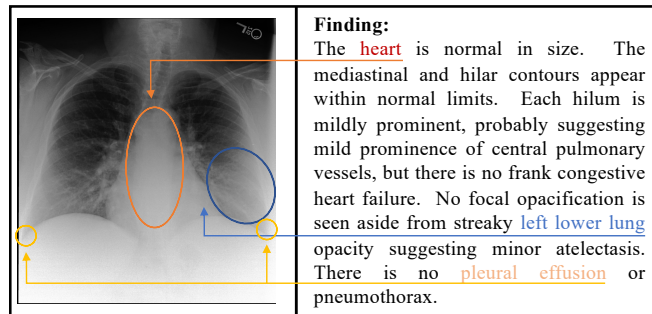


Fig. 1. A Chest X-Ray image with its report findings. The aligned visual and textual features are marked in the same color.

because, different from the traditional image captioning task which mostly requires generating 1 or 2 sentences, RRG requires 2-4 times *more* sentences to describe potentially abnormal regions in an image. Also, radiology reports tend to contain sentences with more sophisticated semantic relationships with their corresponding image regions, calling for a need for *more precise* cross-modal alignment, i.e., alignment between words and image regions. An example is shown in Figure 1 where the cross-modal alignments are marked in the same color. The problem is exacerbated as there exist data biases in commonly used datasets with significantly fewer radiology reports about X-ray images containing abnormal regions, making it hard for the RRG models to efficiently capture abnormal features which often exist only in small parts of pathological images. What is worse, even in pathological cases, most report statements may be associated with a description of normal findings, as shown in Figure 1.

Previous approaches (Zhang *et al.*, 2020; Jing *et al.*, 2018; Chen *et al.*, 2020, 2021) often focus on improving the visual representation capability or better learning the semantic patterns by utilizing a learnable memory, while few studies explicitly explore representation learning with information about image abnormality detection and cross-modal alignment. To bridge this gap, we propose a Class Activation Map guided Attention Network (CAMANet) which distills the discriminative information into the encoder and explicitly improve cross-modal alignment by leveraging class activation maps (Zhou *et al.*, 2016).

This work makes four principal contributions:

1. A novel end-to-end class activation map guided attention model where the class-activation map is utilized to explicitly promote the cross-modal alignment and discriminative representation learning. To the best of our limited knowledge, CAMANet is the first work to leverage class activation maps in this way.
2. A Visual Discriminative Map (VDM) Generation module designed to derive the visual discriminative map from the class activation maps based on the pseudo labels from an automatic labeler.
3. A VDM Assisted Encoder which enriches the discriminative information by utilizing the self-attention mechanism and the VDM.
4. A visual-textual attention consistency module which considers the VDM as the ground truth to supervise the cross-modal attention learning in the decoder, so as to explicitly promote the cross-modal alignment.

Experimental results demonstrate that CAMANet outperforms the previous state-of-the-art methods on two widely-used benchmarks. We conduct ablation studies to further verify the effectiveness of each proposed individual component in CAMANet. Discussion and proposals are given to inspire future work.




Fig. 2. The architecture of CAMANet: An image is fed into the Visual Extractor to obtain patch features which are then utilized to generate the VDM via the VDM Generation module. The proposed VDMAE leverages the VDM to derive a discriminative representation and enrich the discriminative information. Combined with the word embedding, the visual tokens are sent to the transformer to produce the report. After that, the VTAC module generates a TDM from the cross-modal attention scores and considers the VDM as the ground truth to supervise the cross-modal alignment learning.

## 2. Related Work

### 2.1. Image Captioning

Image captioning, aiming to describe an image with human-like sentences, is considered as a high-level visual understanding problem taking advantage of both the techniques of computer vision and natural language processing. Because of the great success in the machine translation and language generation, recent state-of-the-art approaches (Pei et al., 2019; Wang et al., 2020; Liu et al., 2018; Lu et al., 2017; You et al., 2016; Rennie et al., 2017) also follow an encoder-decoder architecture and have demonstrated a great improvement in some traditional image captioning benchmarks. In particular, the most successful models (Cornia et al., 2020; Guo et al., 2020; Pan et al., 2020; Ji et al., 2021) usually adopt the Transformer (Vaswani et al., 2017) as their backbone due to its self-attention mechanism and impressive capability of modelling long-range dependency. However, these methods are tailored for traditional scene images, and not suitable for the radiology images with long sentence report and fine-grained, abnormal regions. Though several works (Krause et al., 2017; Melas-Kyriazi et al., 2018) have been developed to tackle the long text generation, they often fail in capturing specific medical observations and tend to generate normal descriptions, resulting in unsatisfactory performance.

### 2.2. Radiology Report Generation

Following on from the success in image captioning (Cornia et al., 2020; Guo et al., 2020; Pan et al., 2020; Ji et al., 2021; Pei et al., 2019; Wang et al., 2020; Liu et al., 2018), recent state-of-the-art (SOTA) RRG studies all adopt similar architectures: an encoder-decoder to generate the report combined with a convolutional neural network as the visual extractor. For example, Jing et al. (2018) proposed a hierarchical LSTM model to mitigate the long sentence problem and developed a co-attention mechanism to localize abnormal regions. Other works have utilized disease topic information (Liu et al., 2019; Zhang et al., 2020) and obtain

better results. Liu et al. (2021b) extended these works by taking advantage of both the prior and posterior knowledge, i.e., knowledge from the similar images and a predefined topic bag for report generation. Chen et al. (2020) designed a relational memory and a memory-driven conditional layer normalization to better learn the report patterns.

The aforementioned studies often neglect the importance of the cross-modal alignment and the discriminative (disease-related) representation learning. As the only work we know exploring the cross-modal alignment in RRG, R2GenCMN (Chen et al., 2021), designed a shared memory matrix, expected to be able to implicitly learn the cross-modal alignment. However, learning the cross-modal alignment in this way proves difficult since the model is only trained under the cross entropy loss for report generation without any other forms of supervision/guidance for the cross-modal alignment learning. To address these problems, we propose a visual-textual attention consistency module to explicitly promote the cross-modal alignment learning by formulating a task that forces the visual and textual modalities to focus on the same image regions. In addition, we develop a VDM assisted encoder to enrich the discriminative information.

There are a few works (Zhou et al., 2020; Chen et al., 2018) exploring the image-word matching in scene images, which are similar to the cross-modal alignment. However, these methods normally require a well pre-trained fully-supervised detector, e.g. Fast RCNN (Girshick, 2015) and Faster RCNN (Ren et al., 2015) trained on ground truth labels to detect generate the object proposals offline. In addition, they have the valuable bounding box annotation and utilize it to form the mapping supervision. However, there is no large dataset on radiology report generation task containing the costly and valuable ground truth labels to train an accurate detector and form the extra supervision. As a result, they are unsuitable for radiology report generation tasks, and we need to integrate and train the visual extractor in the whole architecture and discover an alternative way to perform the cross-modal alignment. Another significant difference between our work and (Zhou et al., 2020) is that the supervision of (Zhou et al., 2020) is applied on each word-region pair thus making it difficult to learn, while our method improves the cross-modal alignment by ensuring the consistency between the block-level, derived visual and textual discriminative maps.

### 2.3. Class Activation Maps

Proposed by Zhou et al. (2016), a class activation map (CAM) was first used for weakly supervised object detection (Jiang et al., 2021; Bae et al., 2020; Xie et al., 2021), aiming to localize the objects with only image-level annotations. CAMs can identify the discriminative regions in an image by utilizing the *global average pooling layer* (GAP). Consequently, this technique is adopted in some weakly supervised semantic segmentation methods (Ru et al., 2022; Sun et al., 2021; Chen et al., 2022) to generate the pseudo labels and provides performance gains. Different from these works, here we use a CAM to enhance the discriminative representation capability and the cross-modal alignment. We have not seen CAMs being leveraged in this way for the RRG task.

## 3. Method

Before detailing our method, in Section 3.1 we give some background of RRG including the problem formulation and the general architecture. In Section 3.2, we elaborate each proposed component of CAMANet. The objective function is finally described in Section 3.3. An overview of the CAMANet architecture is illustrated in Figure 2.

### 3.1. Background

Given a radiology image  $I$ , such as a chest X-Ray image, the purpose of RRG is to generate the coherent findings (report)  $R$  from  $I$ . Recent state-of-the-art methods often adopt an encoder-decoder framework to generate the report. In particular, a visual extractor, e.g. ResNet101 (He et al., 2016) or DenseNet121 (Huang et al., 2017), is first employed to extract visual features  $\mathbf{v}^i \in \mathbb{R}^{H \times W \times C}$

Lung Opacity	Cardiomegaly	No Finding
Lung Lesion	Consolidation	Edema
Pneumothorax	Pneumonia	Atelectasis
Pleural Effusion	Pleural Other	Fracture
Support Devices	Enlarged Cardiomeastinum	

**Table 1. The 14 categories predefined and output by the automatic labeler.**

which are then flattened to a sequence of visual tokens  $\mathbf{v}^s \in \mathbb{R}^{HW \times C}$ .  $H$ ,  $W$ ,  $C$  are the height, width and the number of channels respectively. This process is formulated as:

$$\{v_1^s, v_2^s, \dots, v_k^s, \dots, v_{N^s-1}^s, v_{N^s}^s\} = f_i(\mathbf{I}), \quad (1)$$

where  $v_k^s$  denotes the patch feature in the  $k^{th}$  position in  $\mathbf{v}^s$ , and  $N^s = H \times W$ .  $f_i$  is the visual feature extractor.

These visual tokens are then fed into the transformer based encoder-decoder to generate a report  $\mathbf{R}$ . Specifically, at time step  $T$ , the encoder maps the visual tokens into an intermediate representations  $l^i \in \mathbb{R}^{1 \times D}$ . An embedding layer is applied to obtain the word embedding  $l^r \in \mathbb{R}^{1 \times D}$  of each word  $w$  in the report  $\mathbf{R}$ .  $D$  is the number of dimensions of hidden states. After this, the decoder takes these two sequence features as source inputs and predicts the current output (word). In general, we express the encoding and decoding processes as:

$$\{l_1^i, l_2^i, \dots, l_{N^s}^i\} = f_{en}(v_1^s, v_2^s, \dots, v_{N^s}^s), \quad (2)$$

$$\{l_1^r, l_2^r, \dots, l_{T-1}^r\} = f_{em}(w_1, w_2, \dots, w_{T-1}), \quad (3)$$

$$p_T = f_{de}(l_1^i, l_2^i, \dots, l_{N^s}^i; l_1^r, l_2^r, \dots, l_{T-1}^r), \quad (4)$$

Where the  $f_{en}$ ,  $f_{em}$  and  $f_{de}$  denotes the encoder, embedding layer and the decoder respectively.  $p_T$  is the word prediction at time step  $T$ .

The core of the encoder and decoder is the attention mechanism *Attn*, Vaswani et al. (2017), expressed as,

$$Attn(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = softmax\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right), \quad (5)$$

$$\mathbf{Q} = \mathbf{W}_q \cdot \mathbf{q}, \mathbf{K} = \mathbf{W}_k \cdot \mathbf{k}, \mathbf{V} = \mathbf{W}_v \cdot \mathbf{v}, \quad (6)$$

where  $\mathbf{W}_q$ ,  $\mathbf{W}_k$  and  $\mathbf{W}_v$  are weights and  $\mathbf{q}$ ,  $\mathbf{k}$  and  $\mathbf{v}$  refer to the query, key and value respectively.  $d$  is the number of dimensions of the hidden states. Note that different from the self-attention layer in other modules where the query, key and value are the same, e.g, visual token or textual token sequence, the second attention layer in the decoder takes the textual tokens as the query and the visual tokens as the key/value. To differentiate it from a typical attention layer, we name it as the *cross-modal attention* in this paper. We refer interested readers to (Vaswani et al., 2017) to understand details of the computation of self-attention in Transformers. We next present the details of the proposed CAMANet.

### 3.2. CAMANet

Learning the cross-modal alignment and discriminative representation is challenging but essential for RRG. To achieve this, CAMANet uses a Visual Discriminative Map Assisted Encoder (VDMAE) and a Visual-Textual Attention Consistency (VTAC) module which take advantage of the class activation map (CAM) (Zhou et al., 2016) technique.

### 3.2.1. Visual Discriminative Map Generation

The first step is to generate a visual discriminative map indicating the discriminative regions in an image. This discriminative map is then used in the VDM Assisted Encoder (VDMAE) and the Visual-Textual Attention Consistency (VTAC) modules to enable the discriminative representation learning and cross-modal alignment.

We propose to leverage the class activation map (CAM) technique to localize discriminative regions. CAM requires the image category labels to form a classification task, so as to obtain the patch contributions (importance) to each category. The category labels are often unavailable for RRG datasets because of the time and expense required to create them. To overcome this problem in CAMANet, we utilize CheXbert (Smit *et al.*, 2020), an automatic labeller, to generate a pseudo label  $\mathbf{y}$  for each image. We can form a multi-label classification task given that  $\mathbf{y}$  is a multi-hot vector encoding the presence of 14 common observations in Chest X-ray images is given in Table 1. As shown in the lower right red box in Figure 2, a classification head is added, taking the extracted visual tokens from the visual extractor as input, to predict the presence  $\tilde{y}_i \in \{0, 1\}$  for  $i^{\text{th}}$  disease formulated as:

$$\mathbf{v}_g = \frac{1}{N^s} \sum_{j=1}^{N^s} \mathbf{v}_j^s, \quad (7)$$

$$\{\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_{N^c}\} = \Phi(\text{Softmax}(\mathbf{W}_c \cdot \mathbf{v}_g)), \quad (8)$$

$$\Phi(o) = \begin{cases} 0 & \text{if } o \leq 0.5, \\ 1 & \text{if } o > 0.5, \end{cases} \quad (9)$$

where  $\mathbf{v}_g$  and  $\mathbf{W}_c$  denote the global visual feature and weights of the classification head.  $N^c$  is the number of disease categories.

After the forward process of the visual extractor, we can obtain the class activation map  $\mathbf{m}^i = \{m_1^i, m_2^i, \dots, m_{N^s}^i\}$  for  $i^{\text{th}}$  category with  $\tilde{y}_i = 1$ . Remember that the classification is for the multi-label scenario, hence each sample may show the presence of multiple observations. To this end, we aggregate the class activation maps for all the categories present and generate a visual discriminative map  $\mathbf{d}^v$  to target important regions. In particular, we first use the ReLU activation function to zero out negative contributions. Min-max normalization is then applied to each class activation map. Equation 10 formulates this process on the class activation map of  $i^{\text{th}}$  category  $\mathbf{m}^i$ :

$$\tilde{m}_j^i = \frac{\sigma(m_j^i) - \min(\sigma(\mathbf{m}^i))}{\sigma(\max(\mathbf{m}^i)) - \sigma(\min(\mathbf{m}^i))}, \quad (10)$$

where  $\tilde{m}_j^i$  is the normalized contribution/importance of  $j^{\text{th}}$  patch for  $i^{\text{th}}$  category and  $\sigma$  denotes the ReLU activation function. Then, the final visual discriminative map  $\mathbf{d}^v \in \mathbb{R}^{1 \times N^s}$  is obtained via maximum pooling over all the class activation maps:

$$\mathbf{d}^v = \text{MaxPool}(\{\mathbf{m}^k | \tilde{y}_k = 1\}). \quad (11)$$

Note that though we utilize the class activation map technique to generate the visual discriminative map, any weakly supervised techniques producing patch-level scores could be applied here. As discussed in subsection 2.2 weakly supervised detection techniques, e.g., Region Of Interest (ROI) based weakly supervised models will generate region-level scores and require a well-trained detector, so is unsuitable to RRG.

### 3.2.2. VDM Assisted Encoder

Since the dataset is dominated by normal samples, e.g., normal images and descriptions, the RRG model struggles to capture abnormal information and describe abnormal regions. To mitigate this problem, we propose a visual discriminative map (VDM) assisted encoder, the VDMAE, to enrich the discriminative information. This is achieved by learning a discriminative representation in terms of the VDM  $\mathbf{d}^v$ , which is distilled into the encoder. We show this procedure by the green rectangle in Figure 2. Specifically,

the discriminative representation  $\mathbf{r} \in \mathbb{R}^{1 \times C}$  is obtained by matrix multiplication between the visual discriminative map  $\mathbf{d}^v \in \mathbb{R}^{1 \times N^s}$  and the image patch features,

$$\mathbf{r} = \mathbf{d}^v \otimes \mathbf{v}^s, \quad (12)$$

where  $\mathbf{v}^s \in \mathbb{R}^{N^s \times C}$  is the extracted visual token sequence obtained by Equation 1 and  $\otimes$  denotes matrix multiplication. A normalization layer is further added to stabilize the distribution:

$$\mathbf{r}' = \text{LayerNorm}(\mathbf{r}). \quad (13)$$

After obtaining the discriminative representation, we add it as an extra visual token as input to the encoder and leverage the power of an attention mechanism to gradually incorporate the VDM features for representation learning. Then Equation 2 becomes:

$$\{\mathbf{r}^*, l_1^i, \dots, l_{N^i}^i\} = f_{en}(\mathbf{r}', v_1^s, \dots, v_{N^i}^s). \quad (14)$$

Through the self-attention mechanism in the encoder, other visual tokens can fully interact with this to learn the useful discriminative information. Note that the encoded discriminative representation  $\mathbf{r}^*$  is detached in the decoder for the sake of better cross-modal alignment in the visual-textual attention consistency module. We elaborate this further in the next section.

### 3.2.3. Visual-Textual Attention Consistency

Cross-modal alignment is essential to learn a coherent and disease-aware model. Previous work learns the alignment implicitly either by a cross-modal attention in the decoder or through a shared memory. However, such strategies are less efficient considering the severe data bias problem and more sophisticated sentence discourse relations in radiology reports. We therefore design a Visual-Textual Attention Consistency (VTAC) module to explicitly learn the cross-modal alignment using the visual discriminative map  $\mathbf{d}^v$  as the ground truth to supervise the cross-modal attention learning in the decoder, as demonstrated in the upper right blue rectangular in Figure 2. The intuition behind this is that when generating a report, the language model should attend to the same regions, e.g., those discriminative regions, as the visual classifier which aims to detect abnormal regions.

Nonetheless, each word in the report has its own attention to the image in the cross-modal attention. This raises a question of how to utilize a single discriminative map to supervise multiple attentions from different words. To deal with this problem, we also form a textual discriminative map  $\mathbf{d}^l$  by aggregating the attentions from *important* words. Then the problem turns to how best to select them and aggregate their attention maps.

One simple way to achieve this is the Term Frequency-Inverse Document Frequency (TF-IDF) (Ramos *et al.*, 2003) score, which however, is not suitable for radiology report corpus as the important words, e.g., medical name entities, in medical domain normally have low TF-IDF scores. For instance, most reports have the words “lung” and “heart” which leads to a low inverse document frequency value, and the frequency of these words in one report normally just one or two. Therefore, the TF-IDF value of these words is small. For example, the IDF value for “heart” on IU-Xray dataset is 1.51 which is similar to the word “is” (IDF 1.43), “are” (IDF 1.33) and “and” (IDF 1.52). Nonetheless, these medical terms are important and should be selected when performing the cross-modal alignment.

To this end, we propose to use the encoded visual discriminative representation  $\mathbf{r}^*$  which contains rich visual discriminative information, to select the attended words. More concretely, we expect that the embeddings of important words should have higher similarity to the visual discriminative representation in the latent space. Specifically, we first compute the cosine similarities between the embedding of each word and the visual discriminative representation via:

$$s_j = \frac{l_j^r \cdot \mathbf{r}^*}{\|l_j^r\| \|\mathbf{r}^*\|}, \quad (15)$$



where  $s_j$  denotes the similarity between  $j^{\text{th}}$  word and the discriminative representative and  $\cdot$  is the dot product operation.

Then, the top  $k\%$  words with the higher similarities are selected to generate a textual discriminative map. Given the cross-model attention scores  $\mathbf{A} = \{\mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^\gamma\}$  of important words from the cross-modal attention layer in the final decoder, the generation of the *textual* discriminative map  $\mathbf{d}^t \in \mathbb{R}^{1 \times N^s}$  is derived via Equation 16 and Equation 17, which is similar to the generation of the *visual* discriminative map,

$$\tilde{\mathbf{a}}_j^i = \frac{\sigma(\mathbf{a}_j^i) - \min(\sigma(\mathbf{a}^i))}{\sigma(\max(\mathbf{a}^i)) - \sigma(\min(\mathbf{a}^i))} \quad (16)$$

$$\mathbf{d}^t = \text{MaxPool}([\tilde{\mathbf{a}}^1, \tilde{\mathbf{a}}^2, \dots, \tilde{\mathbf{a}}^\gamma]), \quad (17)$$

where  $\tilde{\mathbf{a}}_j^i$  is the normalized attention score of  $i^{\text{th}}$  word for  $j^{\text{th}}$  image region.  $N^r$  denotes the number of words in the report.  $\gamma = \lceil k \times N^r \rceil$  is the number of selected words to generate the textual discriminative map.

However, since Equation 17 does not consider the relative importance of the selected words, we utilize the calculated similarity to generate the weights of each word via  $w_j = \sigma(s_j)$ . By considering the relative importance, Equation 17 then becomes:

$$\mathbf{d}^t = \text{MaxPool}([w_1 \tilde{\mathbf{a}}^1, w_2 \tilde{\mathbf{a}}^2, \dots, w_\gamma \tilde{\mathbf{a}}^\gamma]). \quad (18)$$

After obtaining the textual discriminative map,  $\mathbf{d}^t$ , the visual-textual attention consistency is achieved by making  $\mathbf{d}^t$  to be close to the *visual* discriminative map  $\mathbf{d}^v$  through a Mean Squared Error (MSE) loss:

$$\mathcal{L}_{mse} = \frac{1}{N^s} \sum_{j=1}^{N^s} \|\mathbf{d}_j^t - \mathbf{d}_j^v\|, \quad (19)$$

As mentioned earlier, subsection 3.2.2, we detach the encoded visual discriminative representation  $\mathbf{r}^*$  in the decoder. The reason is that adding it as an extra visual token into the decoder will result in a size difference between the visual discriminative map and textual discriminative map ( $N^s$  vs  $N^s + 1$ ), since there are no discriminative tokens available before generating the VDM.

### 3.3. Objective Function

Given the entire predicted word token sequence  $\{p_i\}$  as the generated report and the associated ground truth report  $\{w_i\}$ , CAMANet is jointly optimized with the cross-entropy loss  $\mathcal{L}_{ce}$ , the binary cross entropy loss  $\mathcal{L}_{bce}$  for the multi-label classification in visual discriminative map generation and the mean square error loss  $\mathcal{L}_{mse}$  in the visual-textual attention consistency module:

$$\mathcal{L}_{ce} = -\frac{1}{N^r} \sum_{i=1}^{N^r} w_i \cdot \log(p_i), \quad (20)$$

$$\mathcal{L} = \mathcal{L}_{ce} + \lambda \mathcal{L}_{bce} + \delta \mathcal{L}_{mse}. \quad (21)$$

Here  $\lambda$  and  $\delta$  are two hyper-parameters which balance the loss contributions.

## 4. Experiments

We validate the effectiveness of CAMANet on two commonly used RRG datasets, i.e., IU-Xray and MIMIC-CXR. Four widely used text generation evaluation metrics: BLEU{1-4} (Papineni et al., 2002), Rouge-L (Lin, 2004), METEOR (Denkowski and Lavie, 2011) and CIDEr (Vedantam et al., 2015) are employed to gauge the model performance.

**Table 2. Comparative results of CAMANet with previous studies. The best values are highlighted in bold and the second best are underlined. BL and RG, MTOR and CIDR are the abbreviations of BLEU, ROUGE, METEOR and CIDEr respectively.**

Dataset	Method	BL-1	BL-2	BL-3	BL-4	RG-L	MTOR	CIDR
IU-Xray	<i>ADAATT</i>	0.220	0.127	0.089	0.068	0.308	-	0.295
	<i>ATT2IN</i>	0.224	0.129	0.089	0.068	0.308	-	0.220
	<i>SentSAT + KG</i>	0.441	0.291	0.203	0.147	0.304	-	0.304
	<i>HRGR</i>	0.438	0.298	0.208	0.151	0.322	-	0.343
	<i>CoAT</i>	0.455	0.288	0.205	0.154	0.369	-	0.277
	<i>CMAS – RL</i>	0.464	0.301	0.210	0.154	0.362	-	0.275
	<i>R2Gen</i>	0.470	0.304	0.219	0.165	0.371	0.187	-
	<i>KERP</i>	0.482	0.325	0.226	0.162	0.339	-	0.280
	<i>CMCL</i>	0.473	0.305	0.217	0.162	0.378	0.186	-
	<i>R2GenCMN*</i>	0.475	0.309	0.222	0.170	0.375	0.191	-
	<i>XPRONet</i>	<b>0.525</b>	<u>0.357</u>	<u>0.262</u>	<u>0.199</u>	<b>0.411</b>	<b>0.220</b>	<u>0.359</u>
<i>CAMANet(Ours)</i>	<u>0.504</u>	<b>0.363</b>	<b>0.279</b>	<b>0.218</b>	<u>0.404</u>	<u>0.203</u>	<b>0.418</b>	
MIMIC-CXR	<i>RATCHET</i>	0.232	-	-	-	0.240	0.101	-
	<i>ST</i>	0.299	0.184	0.121	0.084	0.263	0.124	-
	<i>ADAATT</i>	0.299	0.185	0.124	0.088	0.266	0.118	-
	<i>ATT2IN</i>	0.325	0.203	0.136	0.096	0.276	0.134	-
	<i>TopDown</i>	0.317	0.195	0.130	0.092	0.267	0.128	-
	<i>CMCL</i>	0.344	0.217	0.140	0.097	<b>0.281</b>	0.133	-
	<i>R2Gen</i>	<u>0.353</u>	<u>0.218</u>	0.145	0.103	0.277	<u>0.142</u>	-
	<i>R2GenCMN</i>	<u>0.353</u>	<u>0.218</u>	<u>0.148</u>	<u>0.106</u>	0.278	<u>0.142</u>	-
	<i>XPRONet</i>	0.344	0.215	0.146	0.105	<u>0.279</u>	0.138	<u>0.154</u>
	<i>CAMANet(Ours)</i>	<b>0.374</b>	<b>0.230</b>	<b>0.155</b>	<b>0.112</b>	<u>0.279</u>	<b>0.145</b>	<b>0.161</b>

#### 4.1. Datasets

Proposed by the Indiana University, IU-Xray<sup>2</sup> (Demner-Fushman et al., 2016) contains 7,470 X-ray images and 3,955 corresponding reports. The majority of patients provided both the frontal and lateral radiology images. MIMIC-CXR<sup>3</sup> (Johnson et al., 2019) is a large chest X-ray dataset with 473,057 X-ray images and 206,563 reports provided by the Beth Israel Deaconess Medical Center. Both of these two datasets are publicly available. We follow the same data splits as (Chen et al., 2020), to divide the IU-Xray dataset into train (70%), validation (10%) and test (20%) sets and remove samples without both view of images. The official data split is adopted for the MIMIC-CXR dataset.

#### 4.2. Implementation Details

Following the same setting of previous work (Li et al., 2018; Chen et al., 2020, 2021), we utilize both images of a patient on IU-XRay by concatenating the visual tokens, and one view for MIMIC-CXR. Images are firstly resized to (256, 256) and then cropped to (224, 224) (random crop in training and center crop in inference). During the training, we randomly apply one of the operation from {*rotation, scaling*} to further augment the datasets. We employ the DenseNet121 (Huang et al., 2017) pre-trained on ImageNet (Deng et al., 2009) as our visual extractor and a randomly initialized memory driven Transformer (Chen et al., 2020) as the backbone for the encoder-decoder module with 3 layers, 8 attention heads and 512 dimensions for the hidden states. The visual extractor produces  $7 \times 7$  visual tokens, thus  $N^s = 49$ .

The Adam (Kingma and Ba, 2015) is used to optimize CAMANet. The learning rates are set to  $1e - 3$  and  $2e - 3$  for the visual extractor and encoder-decoder on IU-Xray, while MIMIC-CXR has a smaller learning rate with  $5e - 5$  and  $1e - 4$  respectively.  $\lambda$

<sup>2</sup><https://openi.nlm.nih.gov/>

<sup>3</sup><https://physionet.org/content/MIMIC-cxr-jpg/2.0.0/>

and  $\delta$  in Equation 20 are set to 1 and 0.15 on IU-Xray, and 1 and 0.5 on MIMIC-CXR. The proportion of important words  $k$  in Equation 18 is 0.25 and 0.3 on IU-Xray and MIMIC-CXR. Note that the optimal hyper-parameters are determined by evaluating the models on the validation sets.

The same as the most promising studies reported, we adopt a beam size of 3 to balance the effectiveness and efficiency. We implement our model via the PyTorch (Paszke et al., 2019) deep learning framework on Nvidia RTX6000 GPU cards.

#### 4.3. Comparisons to SOTA methods

We compare CAMANet to 13 previous studies, i.e., *ADAATT* (Lu et al., 2017), *ATT2IN* (Rennie et al., 2017), *SentSAT + KG* (Zhang et al., 2020), *HRGR* (Li et al., 2018), *CoAT* (Jing et al., 2018), *CMAS - RL* (Jing et al., 2019), *R2Gen* (Chen et al., 2020), *KERP* (Li et al., 2019), *R2GenCMN* (Chen et al., 2021), *RATCHET* (Hou et al., 2021), *TopDown* (Anderson et al., 2018), *CMCL* (Liu et al., 2021a), *XPRONet* (Wang et al., 2022).

As Table 2 demonstrates, our proposed method achieves the best performance on all the evaluation metrics MIMIC-CXR datasets except for the RG-L metric on which CAMANet scores slightly lower (-0.2%) than *CMCL*. Nevertheless, CAMANet surpasses *CMCL* on all the remaining six evaluation metrics by a notable margin, indicating that CAMANet can generate more semantic reports rather than bias to one indicator such as precision or recall. Similar pattern can be seen on IU-Xray dataset where CAMANet obtains best scores on four out of seven evaluation metrics and is the second best performed method on the remaining three metrics.

Moreover, CAMANet outperforms the previous state-of-the-art approaches by a notable margin. For instance, CAMANet surpasses the second-best BLEU-4 and CIDEr scores by 1.9% and 5.9% respectively on the IU-Xray dataset. Similarly, a 2.1% and 0.6% improvements can be seen on BLEU-1 and BLEU-4 scores of CAMANet compared to the second-best performing method on MIMIC-CXR. Note that the MIMIC-CXR dataset is much larger and challenging and making even a small improvement proves difficult.

We mainly attribute the superiority of CAMANet to the improved cross-modal alignment achieved by the visual-textual attention consistency module, and the enriched disease-related feature representations by the proposed VDM assisted encoder. Further visualizations are provided in Section 4.7. The following ablation studies confirm the effectiveness of each component in CAMANet.

**Use of Pseudo Labels** CAMANet utilizes the pseudo labels to generate the CAMs. Note that the visual extractor of CAMANet is IMAGENET pretrained. Methods *SentSAT + KG* (Zhang et al., 2020) and *RATCHET* (Hou et al., 2021) utilize pseudo labels to pre-train the visual extractor; *HRGR* adopts two auxiliary datasets with ground truth labels to pre-train the visual extractor; whereas *KERP* utilizes some manually labelled data and templates; *CMCL* (Liu et al., 2021a) uses the pseudo labels to fine-tune their visual extractor and determines the visual difficulty; *XPRONet* (Wang et al., 2022) utilizes the pseudo labels to form a cross-modal prototype-driven network. CAMANet still outperforms these methods. Moreover, the pseudo labels we used are provided with the MIMIC-CXR dataset.

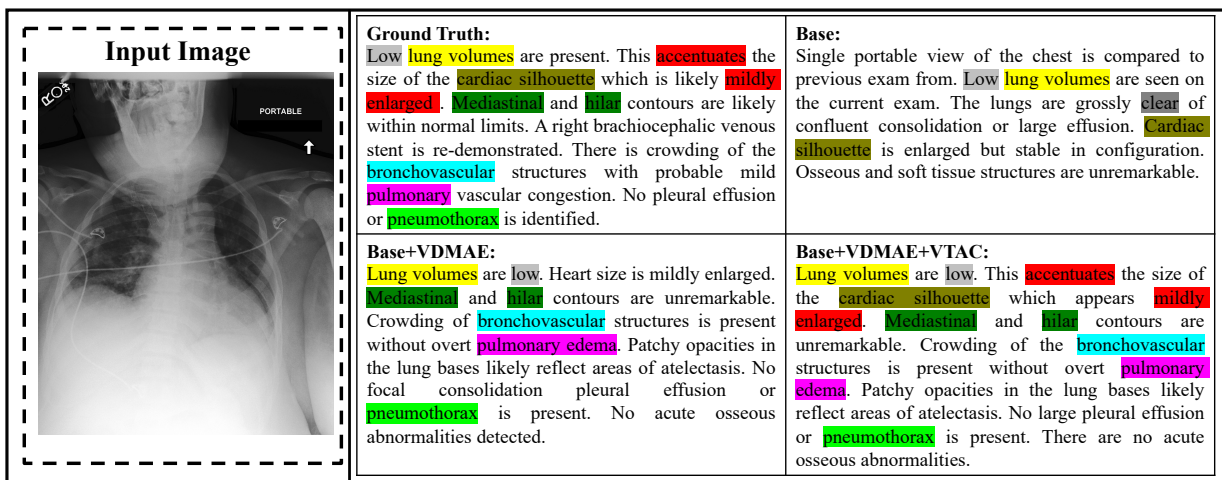
#### 4.4. Ablation Studies

Here, we firstly explore the influence of each component in CAMANet including the VDM Assisted Encoder (VDMAE) and Visual-Textual Attention Consistency (VTAC) module. The following models are used to conduct the ablation studies:

- **Base:** The base model consists only of the visual extractor and the base encoder-decoder without other extensions.

**Table 3. The experimental results of ablation studies on the IU-Xray and MIMIC-CXR datasets. The best values are highlighted in bold. BL and RG, MTOR and CIDR are the abbreviations of BLEU, ROUGE, METEOR and CIDEr respectively.**

IU-Xray	BL-1	BL-2	BL-3	BL-4	RG-L	MTOR	CIDR	AVG.Δ
Base	0.451	0.289	0.209	0.159	0.365	0.175	0.336	-
+VDMAE	0.473	0.321	0.238	0.177	0.402	0.189	0.415	11.8%
+VMAE+VTAC	<b>0.504</b>	<b>0.363</b>	<b>0.279</b>	<b>0.218</b>	<b>0.407</b>	<b>0.203</b>	<b>0.418</b>	<b>22.8%</b>
MIMIC-CXR	BL-1	BL-2	BL-3	BL-4	RG-L	MTOR	CIDR	AVG.Δ
Base	0.324	0.203	0.138	0.100	0.276	0.135	0.144	-
+VDMAE	0.357	0.227	0.148	0.107	0.277	0.142	<b>0.162</b>	7.7%
+VMAE+VTAC	<b>0.374</b>	<b>0.230</b>	<b>0.155</b>	<b>0.112</b>	<b>0.279</b>	<b>0.145</b>	0.161	<b>10.5%</b>



**Fig. 3. An example of the reports generated by different models. Most medical terms are highlighted in different colors to better differentiate the findings of report.**

- **Base+VDMAE:** The standard encoder in the base model is replicated with our proposed VDM assisted encoder by adding the VDM Generation and VDMAE module.
- **Base+VDMAE+VTAC:** This is the full CAMANet containing all our proposed components.

We present the main ablation study results in Table 3. A remarkable improvement can be seen by adding the VDMAE module, e.g., CIDEr scores increase from 0.336 to 0.415, and 0.144 to 0.162 on IU-Xray and MIMIC-CXR benchmarks respectively. Our method gains the most significant improvements when further integrating with the visual-textual attention consistency module. In particular, the full model obtains an average of **22.8%** improvement over all the evaluation metrics on IU-Xray, and **10.5%** on MIMIC-CXR. The improvements on MIMIC-CXR are less obvious than IU-Xray since it is harder to learn a robust discriminative representation and cross-modal alignments in such a large dataset (50 times larger than IU-Xray). Section 4.7 presents some qualitative results to further illustrate the model’s effectiveness.

#### 4.5. Sensitivity of the Proportion of Selected Words

To investigate the sensitivity of the hyper-parameter  $k$ , we vary the proportion of the selected important word tokens from 0.15 to 0.35. As illustrated in Figure 4 and Figure 5, CAMANet is not overly sensitive to this proportion. However, it is still important to strike a balance when setting the value of  $k$  as a small  $k$  may not encompass all the important words, while too large a value may introduce irrelevant words. The fluctuation on IU-Xray is relatively greater than MIMIC-CXR, which is expected since IU-Xray is significantly smaller and the reports contain larger word variance compared to MIMIC-CXR.

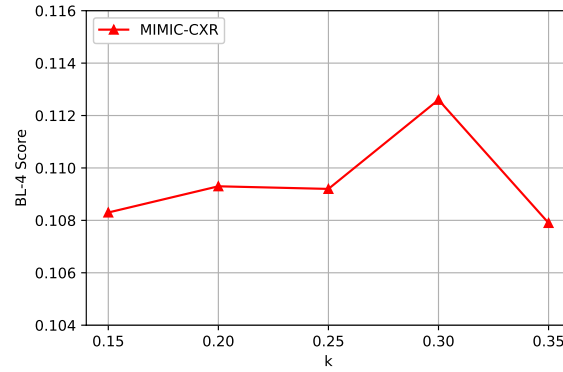


Fig. 4. Effect of varying  $k$ , proportion of important tokens on MIMIC-CXR. (BLEU-4 score).

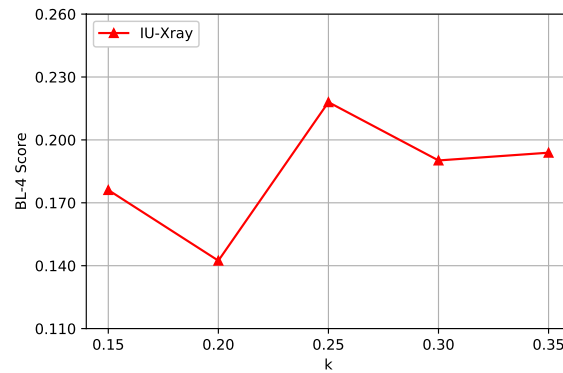


Fig. 5. Effect of varying  $k$ , proportion of important tokens on IU-Xray. (BLEU-4 score).

#### 4.6. Computational Costs

*Training Time.* The batch size is set to 32 for IU-Xray and 64 for MIMIC-CXR. We train all the models for 30 epochs on both datasets. An early stop strategy is adopted, i.e., training stops when no improvement is seen on the validation set for more than 10 epochs. Table 4 shows the average training time of one epoch for the *Base* and *CAMANet* on IU-Xray and MIMIC-CXR. *CAMANet* only slightly increases the training time on MIMIC-CXR, while the increased training time is negligible on IU-Xray.

Dataset	Method	
	Base	CAMANet
IU-Xray	0.41	0.42
MIMIC-CXR	56.35	60.06

Table 4. The average training time of one epoch for different models. The unit of measure is minutes.

*The number of parameters.* Here, we compare the number of parameters of the *Base* and *CAMANet* models in Table 5. The increase in number of learnable parameters is negligible compared to the base model.

Method	Base	CAMANet
#Param	46.07M	46.09M

Table 5. The number of parameters of different models.




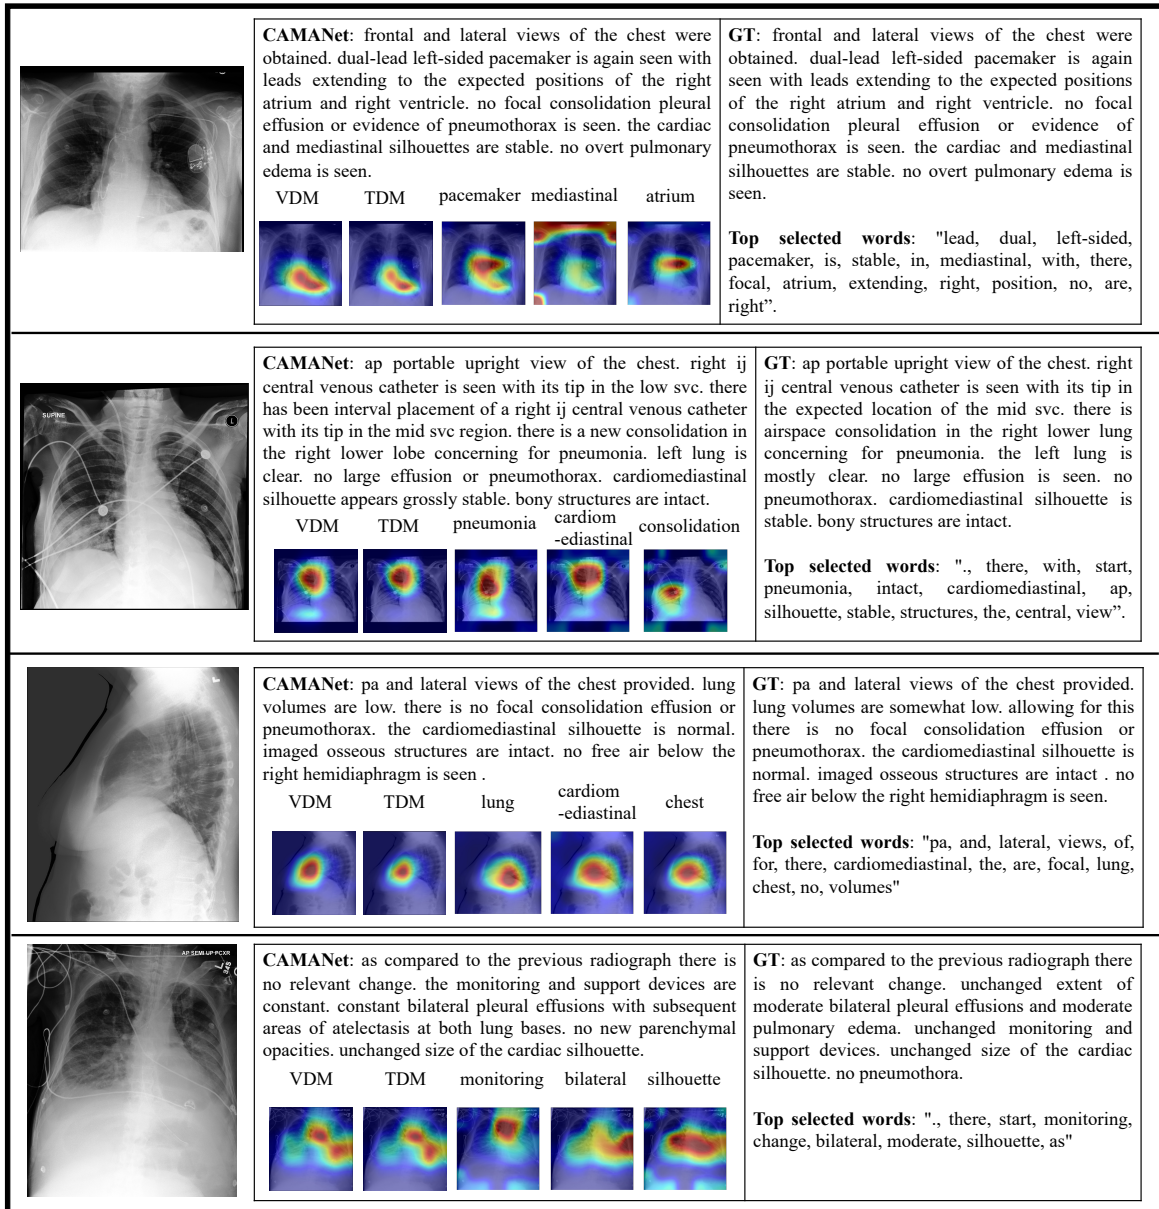
Fig. 6. The visualizations of cross-modal attention maps for some medical terms by different models. The VDM and TDM are also shown.

#### 4.7. Qualitative Results

To further verify the proposed method, we provide an example in Figure 3 which shows the reports generated by different models. CAMANet appears to better capture the disease information and generate the abnormal description, e.g., “...*This accentuates the size of the cardiac silhouette which appears mildly enlarged...*”, compared to the base model. Moreover, we visualize its generated visual discriminative map (VDM), textual discriminative map (TDM) and the cross-modal attention map from the last decoder in Figure 6 to explore whether the cross-modal alignment is truly enhanced by the proposed VTAC module. It can be seen that the *Base* model cannot learn any useful cross-modal alignments when generating the reports. Similar to the *Base* model, *Base+VDMAE* also fails in the cross-modal attention. This is expected since the discriminative representation is detached from the decoder. Nonetheless, owing to the enriched disease information, the *Base+VDMAE* model still can capture some abnormal single-model patterns, and therefore achieves improved performance compared to the *Base* model. It is clear that after adding the VTAC module, CAMANet demonstrates potent cross-modal alignment capability where the medical terms, e.g., ‘*lung*’, focus on the associated discriminative image regions, while trivial words, e.g., ‘*without*’, pay attention only to the background. Visualizations of the VDM and TDM further verify the effectiveness of employing a VDM as a form of pseudo label to supervise the cross-modal attention learning.

During the experiments, we found that even if a sample is wrongly classified, the model can still produce a class activation map focusing on the foreground, e.g, chest or lung, rather than the background, reducing the reliance on pseudo labels.

Finally we present further visualization results in Figure 7. The visualizations show the reports predicted by CAMANet and the associated ground truth. We also provide the visual discriminative map (VDM), textual discriminative map (TDM), the cross-attention maps of some medical terms in the inference mode and include important words in the training mode.



**Fig. 7.** Some examples of generated reports by CAMANet and the associated VDM, TDM, Cross-modal attention maps and selected important words. Duplicated words in the selected important words have been removed.

### 5. Conclusions

In this work, we propose a novel class activation map guided RRG framework, CAMANet, which explicitly leverages cross-modal alignment and disease-related representation learning. Our visual discriminative map (VDM) assisted encoder distills the discriminative information into the model via a derived discriminative representation and the self-attention mechanism. The generated VDM is then regarded as the ground truth in the proposed visual-textual attention consistency (VTAC) module to supervise the cross-modal attention learning, aimed at explicitly promoting cross-modal alignment. Experimental results on two widely used RRG benchmarks prove the superiority of CAMANet over previous studies. The ablation studies further verify the effectiveness of individual components of CAMANet. Moreover, a number of illustrative visualizations and discussions are provided to inspire future research in this area.

**Declaration of competing interest**

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

**Acknowledgments**

Jun Wang's PhD studentship is jointly funded by the University of Warwick and China Scholarship Council. We would like to thank the University of Warwick's Scientific Computing Group for their support of this work and acknowledge the resources provided by UKRI/EPSRC HPC platform, Sulis and Avon.

**Source code**

We will be providing public access to Python code and relevant model data in due course.



## References

- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L., 2018. Bottom-up and top-down attention for image captioning and visual question answering, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6077–6086.
- Bae, W., Noh, J., Kim, G., 2020. Rethinking class activation mapping for weakly supervised object localization, in: European Conference on Computer Vision, Springer. pp. 618–634.
- Chen, K., Gao, J., Nevatia, R., 2018. Knowledge aided consistency for weakly supervised phrase grounding, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4042–4050.
- Chen, Z., Shen, Y., Song, Y., Wan, X., 2021. Cross-modal memory networks for radiology report generation, in: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 5904–5914.
- Chen, Z., Song, Y., Chang, T.H., Wan, X., 2020. Generating radiology reports via memory-driven transformer, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1439–1449.
- Chen, Z., Wang, T., Wu, X., Hua, X.S., Zhang, H., Sun, Q., 2022. Class re-activation maps for weakly-supervised semantic segmentation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 969–978.
- Cornia, M., Stefanini, M., Baraldi, L., Cucchiara, R., 2020. Meshed-memory transformer for image captioning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 10578–10587.
- Demner-Fushman, D., Kohli, M.D., Rosenman, M.B., Shooshan, S.E., Rodriguez, L., Antani, S., Thoma, G.R., McDonald, C.J., 2016. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association* 23, 304–310.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition, Ieee. pp. 248–255.
- Denkowski, M., Lavie, A., 2011. Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems, in: Proceedings of the sixth workshop on Statistical Machine Translation, pp. 85–91.
- Girshick, R., 2015. Fast r-cnn, in: Proceedings of the IEEE International Conference on Computer Vision, pp. 1440–1448.
- Guo, L., Liu, J., Zhu, X., Yao, P., Lu, S., Lu, H., 2020. Normalized and geometry-aware self-attention network for image captioning, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10327–10336.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778.
- Hou, B., Kaissis, G., Summers, R.M., Kainz, B., 2021. Ratchet: Medical transformer for chest x-ray diagnosis and reporting, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 293–303.
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4700–4708.
- Ji, J., Luo, Y., Sun, X., Chen, F., Luo, G., Wu, Y., Gao, Y., Ji, R., 2021. Improving image captioning by leveraging intra-and inter-layer global representation in transformer network, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 1655–1663.
- Jiang, P.T., Zhang, C.B., Hou, Q., Cheng, M.M., Wei, Y., 2021. Layercam: Exploring hierarchical class activation maps for localization. *IEEE Transactions on Image Processing* 30, 5875–5888.
- Jing, B., Wang, Z., Xing, E., 2019. Show, describe and conclude: On exploiting the structure information of chest x-ray reports, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 6570–6580.
- Jing, B., Xie, P., Xing, E., 2018. 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), pp. 2577–2586.
- Johnson, A.E., Pollard, T.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Peng, Y., Lu, Z., Mark, R.G., Berkowitz, S.J., Horng, S., 2019. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization, in: International Conference on Learning Representations, (Poster).
- Krause, J., Johnson, J., Krishna, R., Fei-Fei, L., 2017. A hierarchical approach for generating descriptive image paragraphs, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 317–325.
- Li, C.Y., Liang, X., Hu, Z., Xing, E.P., 2019. Knowledge-driven encode, retrieve, paraphrase for medical image report generation, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 6666–6673.
- Li, Y., Liang, X., Hu, Z., Xing, E.P., 2018. Hybrid retrieval-generation reinforced agent for medical image report generation. *Advances in Neural Information Processing Systems* 31.
- Lin, C.Y., 2004. Rouge: A package for automatic evaluation of summaries, in: Text Summarization Branches Out, pp. 74–81.
- Liu, F., Ge, S., Wu, X., 2021a. Competence-based multimodal curriculum learning for medical report generation, in: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Association for Computational Linguistics, Online. pp. 3001–3012.
- Liu, F., Ren, X., Liu, Y., Wang, H., Sun, X., 2018. Simnet: Stepwise image-topic merging network for generating detailed and comprehensive image captions, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 137–149.
- Liu, F., Wu, X., Ge, S., Fan, W., Zou, Y., 2021b. Exploring and distilling posterior and prior knowledge for radiology report generation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 13753–13762.
- Liu, G., Hsu, T.M.H., McDermott, M., Boag, W., Weng, W.H., Szolovits, P., Ghassemi, M., 2019. Clinically accurate chest x-ray report generation, in: Machine Learning for Healthcare Conference, PMLR. pp. 249–269.
- Lu, J., Xiong, C., Parikh, D., Socher, R., 2017. Knowing when to look: Adaptive attention via a visual sentinel for image captioning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 375–383.
- Melas-Kyriazi, L., Rush, A.M., Han, G., 2018. Training for diversity in image paragraph captioning, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 757–761.
- Miura, Y., Zhang, Y., Tsai, E., Langlotz, C., Jurafsky, D., 2021. 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, pp. 5288–5304.
- Nishino, T., Ozaki, R., Momoki, Y., Taniguchi, T., Kano, R., Nakano, N., Tagawa, Y., Taniguchi, M., Ohkuma, T., Nakamura, K., 2020. Reinforcement learning with imbalanced dataset for data-to-text medical report generation, in: Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 2223–2236.
- Pan, Y., Yao, T., Li, Y., Mei, T., 2020. X-linear attention networks for image captioning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 10971–10980.
- Papineni, K., Roukos, S., Ward, T., Zhu, W.J., 2002. 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.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* 32.

- Pei, W., Zhang, J., Wang, X., Ke, L., Shen, X., Tai, Y.W., 2019. Memory-attended recurrent network for video captioning, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8347–8356.
- Ramos, J., et al., 2003. Using tf-idf to determine word relevance in document queries, in: Proceedings of the first instructional conference on machine learning, Citeseer. pp. 29–48.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems* 28.
- Rennie, S.J., Marcheret, E., Mroueh, Y., Ross, J., Goel, V., 2017. Self-critical sequence training for image captioning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7008–7024.
- Ru, L., Zhan, Y., Yu, B., Du, B., 2022. Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16846–16855.
- Smit, A., Jain, S., Rajpurkar, P., Pareek, A., Ng, A.Y., Lungren, M., 2020. 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), pp. 1500–1519.
- Sun, K., Shi, H., Zhang, Z., Huang, Y., 2021. Ecs-net: Improving weakly supervised semantic segmentation by using connections between class activation maps, in: Proceedings of the IEEE International Conference on Computer Vision, pp. 7283–7292.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30.
- Vedantam, R., Lawrence Zitnick, C., Parikh, D., 2015. Cider: Consensus-based image description evaluation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4566–4575.
- Wang, J., Bhalerao, A., He, Y., 2022. Cross-modal prototype driven network for radiology report generation, in: *European Conference on Computer Vision*, Springer.
- Wang, J., Tang, J., Luo, J., 2020. Multimodal attention with image text spatial relationship for ocr-based image captioning, in: Proceedings of the 28th ACM International Conference on Multimedia, pp. 4337–4345.
- Xie, J., Luo, C., Zhu, X., Jin, Z., Lu, W., Shen, L., 2021. Online refinement of low-level feature based activation map for weakly supervised object localization, in: Proceedings of the IEEE International Conference on Computer Vision, pp. 132–141.
- You, Q., Jin, H., Wang, Z., Fang, C., Luo, J., 2016. Image captioning with semantic attention, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4651–4659.
- Zhang, Y., Wang, X., Xu, Z., Yu, Q., Yuille, A., Xu, D., 2020. When radiology report generation meets knowledge graph, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 12910–12917.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2921–2929.
- Zhou, Y., Wang, M., Liu, D., Hu, Z., Zhang, H., 2020. More grounded image captioning by distilling image-text matching model, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4777–4786.