

¹Fuhai Chen
²Fufeng Chen
³Xiaojing Ma
^{45*}Xuri Ge

Factored Attention and Embedding for Unstructured-view Topic-related Ultrasound Report Generation



Abstract: - Echocardiography is widely used to clinical practice for diagnosis and treatment, e.g., on the common congenital heart defects. The traditional manual manipulation is error-prone due to the staff shortage, excess workload, and less experience, leading to the urgent requirement of an automated computer-aided reporting system to lighten the workload of ultrasonologists considerably and assist them in decision making. Despite some recent successful attempts in automatic medical report generation, they are trapped in the ultrasound report generation, which involves unstructured-view images and topic-related descriptions. To this end, we investigate the task of the unstructured-view topic-related ultrasound report generation, and propose a novel factored attention and embedding model (termed FAE-Gen). The proposed FAE-Gen mainly consists of two modules, i.e., view-guided factored attention and topic-oriented factored embedding, which 1) capture the homogeneous and heterogeneous morphological characteristic across different views, and 2) generate the descriptions with different syntactic patterns and different emphatic contents for different topics. Experimental evaluations are conducted on a large-scale clinical cardiovascular ultrasound dataset (CardUltData). Both quantitative comparisons and qualitative analysis demonstrate the effectiveness and the superiority of FAE-Gen over seven commonly-used metrics.

Keywords: Medical Image Analysis, Ultrasound Image, Report Generation, Attention, Deep Learning, Pattern Recognition

I. INTRODUCTION

Echocardiography is widely used in hospitals for the diagnosis of common congenital heart defects in both children and adults, such as ventricular septal defect (VSD) [1] and atrial septal defect (ASD) [2]. An ultrasonologist completes an ultrasonic diagnostic report ahead of ultrasound scanner by analyzing the ultrasound images in different sections (views), where the images of each view record the different blood flow movements. However, this process of medical image interpretation and reporting can be error-prone due to staff shortage, excess workload, and less experience [3–6]. Therefore, an automated computer-aided reporting system is urgently required to reduce workload and error occurrences, where the ultrasound images are taken as inputs and the diagnostic report is automatically generated.

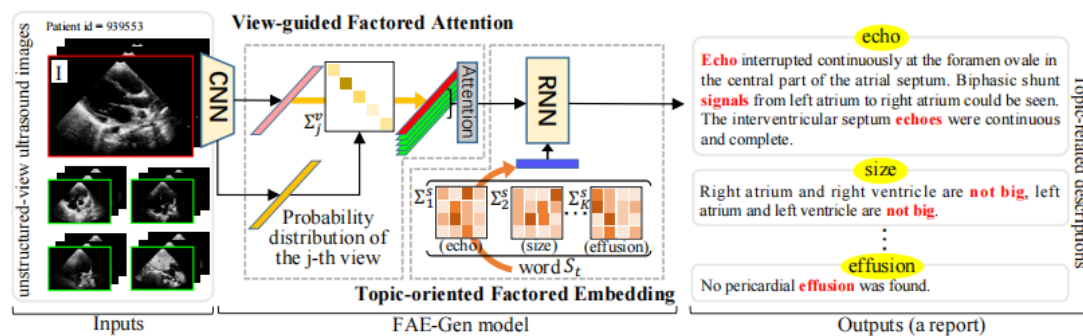


Fig. 1. The overview of the proposed FAE-Gen for automatic unstructured-view topic-related ultrasound report generation.

¹ College of Computer and Data Science, Fuzhou University, Fuzhou, 350116, Fujian, China.

² Fuzhou Institute for Data Technology, Fuzhou, Fujian, China.

³ Wuhan Asia Heart Hospital, Wuhan, Hubei, China.

⁴ School of Computing Science, University of Glasgow, Sir Alwyn Williams Building (SAWB), Glasgow, G12 8RZ, United Kingdom.

⁵ Shandong University, Jinan, 250100, Shandong, China.

*Corresponding author: Xuri Ge; email: xurige@fuzhou.edu.cn

Recently, there are increasing works attracted in the medical report generation problem [7–10]. Typically, inspired by image captioning [11], they achieve the system on the pairs of the medical images and the corresponding reports via a deep learning technology based encoder-decoder architecture, where the features of a medical image are first extracted by a convolutional neural network (CNN) [12,13] and then fed into recurrent neural networks (RNNs) [14] to generate the word sequences as a diagnostic report. For example, Jing et al. [10] proposed a co-attention model, where the visual (morphological) features and the tag features are extracted from CNN and jointly attended into hierarchical RNNs [15] to describe the chest radiology images. Most of these works focused on generating a generally-described report upon one or two structured-view (frontal or lateral) images, like radiology report generation task [8,10], which, however, cannot handle the case when the views of the medical images are unstructured⁶ and the descriptions are topic-related in ultrasound report generation as shown in Fig. 1.

In the automatical ultrasound report generation, there are two key issues for the unstructured-view inputs and the topic-related outputs. On the one hand, the ultrasound images of different views have their homogeneous morphological characteristic on the same pathology as well as their heterogeneous morphological characteristic over different views. Thus the morphological relevance and diversity of the across-view ultrasound images should be captured during the visual feature representation for the accurate and full-sided text generation. On the other hand, different topic-related diagnostic descriptions have different syntactic patterns and different emphatic contents. To this end, different topic related word embedding should be imported into RNN to generate the discriminative descriptions.

Driven by the above insights, we propose a novel factored attention and embedding model for automatical unstructured-view topic-related ultrasound report generation (termed FAE-Gen) as detailedly illustrated in Fig. 1. To capture the morphological relevance and diversity of the across-view ultrasound images, we develop a view-guided factored attention module to disentangle the view factors from the chaotic feature transformation. To import the topic-related word embedding into discriminative diagnostic descriptions, we design a topic-oriented factored embedding module to decompose the word embedding matrix into a topic-specific matrix and two shared transformation matrices. Specially, given unstructured-view ultrasound images of a patient, for each ultrasound image, we first extract its morphological feature (in pink) and the corresponding probability distribution of the recognized view (in yellow) from a pretrained CNN model. Then, in the view-guided factored attention module, the view distribution is diagonalized as a view-specific factor matrix, and its morphological feature is transformed into a view-related feature by composing of the factor matrix and the shared parameterized matrices. The new features of all images are attended into the state of RNN (Sec. II.A). Meanwhile, in each state of RNN, the word vector is transformed into a topic-related word feature by composing of the topic-related parameterized factor matrix and the other two shared parameterized matrices in topic-oriented factored embedding module. (Sec. II.B). In this way, the topic-related word features are fed into the states of RNN to generate a topic-related description. Finally, the overall model is trained upon the supervision of the real report descriptions in an end-to-end manner.

The contributions of our work are: 1) we are the first to deal with a new medical report generation problem, i.e., automatical ultrasound report generation, with unstructured-view image inputs and topic-related description outputs. 2) we propose a novel factored attention and embedding model (FAE-Gen) to capture the homogeneous and heterogeneous morphological characteristic across different views in the view-guided factored attention module, as well as to generate the topic-related diagnostic descriptions in the topic-oriented factored embedding module.

II. METHODOLOGY

The proposed FAE-Gen is based on the fundamental encoder-decoder architecture of image captioning [11], where an image I is first encoded into a visual feature \mathbf{v} by utilizing the convolution neural network (CNN) and then the visual feature \mathbf{h}^v is fed into a recurrent neural network (RNN) to generate the sequential words $\hat{S}_{1:T}$ (T denotes the textual length of the description). In the unstructured-view topic-related ultrasound report generation, suppose there are M ultrasound images $I_{1:M}$ in different views and K

⁶ **Defined as:** 1) the identifications of views are unknowable, 2) some views are missing, and 3) the ultrasound images (video frames) of each view vary a lot due to blood flow movements).

real descriptions $S^{1:K}$ with different topics, the objective of FAE-Gen is to maximize the log likelihood as below:

$$\log P(S^k | I_{1:M}; \theta) = \sum_{t=0}^T \log p(S_t^k | \mathbf{h}_{1:M}^v, \mathbf{y}_{1:M}, S_{0:t-1}^k), \quad (1)$$

where $k \in \{1, \dots, K\}$ denotes one of the topics. θ denotes the model parameter set that needs to be learned during model training, i.e., the optimal θ^* can be obtained by $\arg \max_{\theta} \sum_{(I_{1:M}, S^k)} \log P(S^k = S^k | I_{1:M}; \theta)$. To simplify, we omit θ in the right side. $\mathbf{h}_{1:M}^v$ and $\mathbf{y}_{1:M}$ denote the morphological features and the probability distributions of different views, respectively. To extract $\mathbf{h}_{1:M}^v$ and $\mathbf{y}_{1:M}$, we modify a CNN model, ResNet-50 [13], with two softmax outputs corresponding to the category spaces of the cardiovascular diseases (ASD, VSD, and normal) and five views (see Sec. III for details), respectively. We pretrain the ResNet-50 under the supervision of the disease and view categories, after which $\mathbf{h}_{1:M}^v$ can be obtained from the final fully-connected layer and $\mathbf{y}_{1:M}$ can be predicted from the corresponding softmax output. We transform \mathbf{h}^v into a view-related feature by using \mathbf{y} (Sec. II.A). To obtain the t -th word \hat{S}_t , we use the Bi-directional Long Short-Term Memory (Bi-LSTM)⁷ [16] to generate S_t depending on $\mathbf{h}_{1:M}^v$, $\mathbf{y}_{1:M}$ and the preceding words $S_{0:t-1}$ (S_0 is a start sign) according to the chain rule. Commonly, S_t is represented with a one-hot vector \mathbf{x}_t according to its index in the dictionary. The word feature \mathbf{h}_t^s is then obtained via $\mathbf{h}_t^s = \mathbf{E}\mathbf{x}_t$, where \mathbf{E} is an embedding matrix [17]. In our paper, \mathbf{E} decomposes into a topic-related factor matrix and other embedding matrices for topic-guided description generations (Sec. II.B).

A. View-guided Factored Attention

In the proposed FAE-Gen, there are multiple ultrasound images of different views as the inputs, leading to two key technical points. First, how to transform the morphological feature \mathbf{h}^v into view-related features for the enhanced and discriminative representation. Second, how to weight the importance of these view-specific features. To this end, we design a view-guided factored attention module as shown in Fig. 1, where the traditional transformation $\hat{\mathbf{h}}^v = \mathbf{W}\mathbf{h}^v$ is advanced as follows:

$$\hat{\mathbf{h}}_j^v = \mathbf{U}\Sigma_j^v\mathbf{V}\mathbf{h}_j^v, \quad (2)$$

where j denotes the j -th view. \mathbf{U} and \mathbf{V} are the shared parameterized matrices among all the samples. Σ_j^v is a view-related factor matrix obtained by $\text{diag}(\mathbf{y}_j)$, which plays a role of view guidance for the transformation. To weight these view-specific visual features and feed them into the RNN, we calculate their relevance to the hidden feature in each state of RNN via an attention mechanism. For the t -th RNN state, formulated as:

$$\alpha_{j,t} = \mathbf{W}^a \sigma(\mathbf{W}^v \hat{\mathbf{h}}_j^v + \mathbf{W}^z \mathbf{h}_{t-1}^s), \alpha_t = \text{softmax}(\alpha_t), \mathbf{h}_t^a = \sum_{j=1}^M \alpha_{j,t} \hat{\mathbf{h}}_j^v, \quad (3)$$

where \mathbf{W}^a , \mathbf{W}^v , and \mathbf{W}^z are the shared parameter matrices of linear transformation. \mathbf{h}_{t-1}^s denotes the preceding hidden feature in the t -th RNN state (detailed in Sec. II.B). σ is a non-linear function (we use hyperbolic tangent). α_t is the relevance vector. Finally, a weight-sum attention feature \mathbf{h}_t^a is obtained as Eq. 3.

B. Topic-oriented Factored Embedding

The traditional word embedding $\mathbf{h}_t^s = \mathbf{E}\mathbf{x}_t |_{t=0}^T$ causes the single syntactic pattern and emphatic content due to its uniform embedding matrix \mathbf{E} for the word representation. To generate the topic-related description, we design a topic-oriented factored embedding module as shown in Fig. 1. Specifically, we use the Bi-LSTM as our RNN model. The formulations in the t -th RNN state are:

$$\tilde{\mathbf{h}}_t^s = f(\overrightarrow{\mathbf{A}}\overrightarrow{\Sigma}_k^s \overrightarrow{\mathbf{B}}\mathbf{x}_{t-1}, \overrightarrow{\mathbf{W}}^s \mathbf{h}_t^a), \tilde{\mathbf{h}}_t^s = f(\overleftarrow{\mathbf{A}}\overleftarrow{\Sigma}_k^s \overleftarrow{\mathbf{B}}\mathbf{x}_{t-1}, \overleftarrow{\mathbf{W}}^s \mathbf{h}_t^a), \mathbf{h}_t^s = g([\tilde{\mathbf{h}}_t^s, \tilde{\mathbf{h}}_t^s]), \quad (4)$$

$$S_t \sim p_t = \text{softmax}(\mathbf{h}_t^s), \quad (5)$$

⁷ Bi-LSTM is an advanced version of RNN, which is widely used in long sentences or paragraphs to better capture the textual contexts. In this paper, we still call it RNN for readability.

where \rightarrow and \leftarrow denote the forward and the backward directions. f is a generic function in LSTM, which includes input, forget, output, and cell-related functions [14]. $\vec{\mathbf{A}}, \vec{\mathbf{B}}, \overleftarrow{\mathbf{A}},$ and $\overleftarrow{\mathbf{B}}$ are the shared parameter matrices. $\vec{\Sigma}_k^s$ and $\overleftarrow{\Sigma}_k^s$ are the k -th topic related parameter matrices, which are optimized with the topic-related descriptions. All the above matrices are initialized randomly. Bidi-rectional $\vec{\mathbf{h}}_t^s$ and $\overleftarrow{\mathbf{h}}_t^s$ are then mapped into \mathbf{h}_t^s via a non-linear function g . Finally, the current word S_t can be sampled according to the word probability distribution \mathbf{p}_t .

III. EXPERIMENTS

In this section, we first introduce the dataset we used on cardiovascular ultrasound images. Then we describe the experimental settings. Finally, we evaluate and discuss the performance of the proposed method on the cardiovascular ultrasound report generation task.

Dataset. We obtain the original cardiovascular ultrasound images from the specialized hospital. The patients' heart dynamic images are recorded through the Philips ultrasound machine. Comprehensive evaluation of the cardiac structure and function in patients by two dimensional echocardiography, color Doppler echocardiography and M-mode echocardiography. And Images and data are stored in Philips IE33 and 7C ultrasound machines. Images and data are diagnosed by an experienced ultrasound physician who has worked for more than 2 years analyzing images and data. All the private information of the patients have been processed. For cardiovascular diseases, doctors typically diagnose patients with observed data of echocardiography, and obtain three types of diagnostic results: normal, VSD, ASD. VSD refers to ventricular septal defect, and ASD refers to atrial septal defect. Doctors photographed 10 types of data for each patient, corresponding to 10 sections: (1, 2) parasternal left ventricular long axis 2D + color, (3, 4) parasternal artery short axis 2D + color, (5, 6) apical four-chamber section 2D + color, (7, 8) apical five-chamber section 2D + color, (9, 10) subxiphoid two-chamber section 2D + color. Each section data is a video sequence, where the section is treat as a view of image in our paper. And we select five views (1,3,5,7,9) among them, resulting in 120,179 ultrasound images. We further randomly select five images (in the video) of the five views for each sample, where some views may be missing, and some views may be repeated but with different blood flow motions. Doctors also generate a corresponding diagnosis report (in English) for each patient, which belongs to either normal, VSD or ASD. Thereout, we assign the reports to the samples according to their patient id, where each report is split into different descriptions according to the predefined topics, like echo and motion. Finally, we split the image-description pairs into 19,324 training samples and 4,593 testing samples. The dataset⁸ (named CardUltData) has been publicly released to promote the research.

Settings. Common evaluation metrics for image captioning are used to provide a quantitative comparisons. Specially, we use Bleu (B), Meteor (M), Rouge-L (R), and CIDEr (C) as [18]. Each of them evaluates the consistency (include co-occurrence) between the generated/candidate and real/reference descriptions. For the model architecture, we set the dimensions of all the features in RNN as 512. The sizes of the view-related and topic-related factor matrices are 5×5 and 10×10 , respectively. We employ PyTorch⁹ to implement our model with maximal 60 training epochs.

Table 1. Performance comparisons of cardiovascular ultrasound report generation on CardUltData. “B-n” denotes the n-gram co-occurrence in Bleu metric.

Method	B-1	B-2	B-3	B-4	C	M	R
CNN-RNN	0.859	0.826	0.800	0.779	5.966	0.567	0.850
SCST-ATT	0.867	0.832	0.806	0.779	6.027	0.570	0.862
FA-Gen	0.862	0.830	0.807	0.783	6.154	0.573	0.864

⁸ <http://mac.xmu.edu.cn/challenge/MICCAI2019-AutoGen-CDR19/index.html>

⁹ <http://pytorch.org>

FE-Gen	0.868	0.837	0.815	0.786	6.196	0.580	0.871
F AE-Gen	0.873	0.839	0.820	0.786	6.217	0.581	0.877

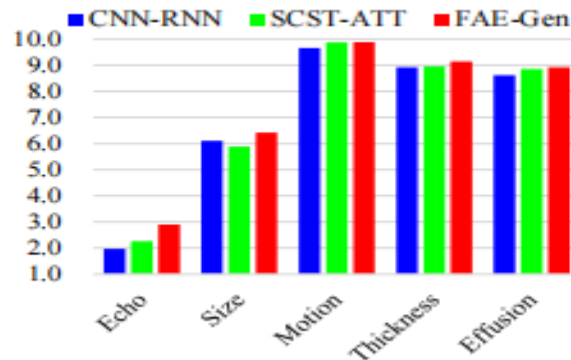


Fig. 2. CIDEr evaluations on different topics.

Evaluation and Discussion. We conduct the quantitative comparisons and qualitative analysis for the evaluation on cardiovascular ultrasound report generation. **First**, we compare the proposed FAE-Gen to some respective and state-of-the-art image captioning methods, including 1) CNN-RNN [11], the vanilla encoder-decoder method, 2) SCST-ATT [19], a state-of-the-art image captioning model with visual attention, 3) FE-Gen, an alternative version of FAE-Gen without view-guided factored attention, 4) FA-Gen, an alternative version of FAE-Gen without topic-oriented factored embedding.




Fig. 3. Example results on CardUitData. Left: input VSD images with multiple views. Middle: the generated descriptions, Right: the ground truth (real) descriptions. Key words are marked in red.

The results are shown in Tab. 1, from which we find that the proposed FAE-Gen outperforms the others over most of the metrics. Specially, FAE-Gen achieves significant improvements over FE-Gen and FA-Gen, which indicates the effectiveness and superiority of both the view-guided factored attention and topic-oriented factored embedding in the unstructured-view topic-related cardiovascular ultrasound report generation. Additionally, we compare the performances on CIDEr over different topics in Fig. 2, where FAE-Gen achieves significant improvements, especially on the topic Echo. This is probably due to the long and various syntaxes in the dataset. **Second**, we provide the generated descriptions by FAE-Gen in Fig. 3, where the descriptions generated by FAE-Gen are more consistent with the ground truths, including the key pathological descriptions, like increase, normal, and wide. We also find that the descriptions of some topics are consistent with the ground truth, which is probably due to the similar syntaxes of these topics in the reports.

IV. CONCLUSION

In this paper, we investigate a new computer-aided medical imaging task, i.e., ultrasound report generation, which involves unstructured-view ultrasound images and topic-related descriptions. To this end, we propose a novel factored attention and embedding model (FAE-Gen), which mainly consists of view-guided factored

attention and topic-oriented factored embedding modules. On the one hand, the view-guided factored attention module captures the homogeneous and heterogeneous morphological characteristic across different views. On the other hand, the topic-oriented factored embedding module generates the descriptions with different syntactic patterns and different emphatic contents for different topics. Experimental evaluations are conducted on a large-scale clinical cardiovascular ultrasound dataset (CardUltData). Both quantitative comparisons and qualitative analysis demonstrate the effectiveness and the superiority of FAE-Gen over seven commonly-used evaluation metrics.

FUNDING

This document is the results of the research project funded by the Fujian Provincial Department of Education Youth Project, China (grant JZ230006) and the Engineering Research Center of Big Data Intelligence, Ministry of Education, China, and Fujian Key Laboratory of Network Computing and Intelligent Information Processing (Fuzhou University), China, and the Fujian Provincial Health Commission science and technology plan project (2023CXA040).

COMPETING INTERESTS

The authors declared no potential conflicts of interest with respect to the research, author-ship, and publication of this article.

REFERENCES

- [1] Mary S Minette and David J Sahn. Ventricular septal defects. *Circulation*, 114(20):2190–2197, 2006.
- [2] Gary Webb and Michael A Gatzoulis. Atrial septal defects in the adult: recent progress and overview. *Circulation*, 114(15):1645–1653, 2006.
- [3] James B Seward, Pamela S Douglas, Raimund Erbel, et al. Hand-carried cardiac ultrasound (hcu) device: recommendations regarding new technology. *Journal of the American Society of Echocardiography*, 15(4):369–373, 2002.
- [4] Daniel F Niendorff, Athos J Rassias, Robert Palac, et al. Rapid cardiac ultrasound of inpatients suffering pea arrest performed by nonexpert sonographers. *Resuscitation*, 67(1):81–87, 2005.
- [5] Brian P Lucas, Carolina Candotti, Bosko Margeta, et al. Diagnostic accuracy of hospitalist-performed hand-carried ultrasound echocardiography after a brief training program. *Journal of hospital medicine: an official publication of the Society of Hospital Medicine*, 4(6):340–349, 2009.
- [6] Thomas M Stokke, Vidar Ruddox, Sebastian I Sarvari, et al. Brief group training of medical students in focused cardiac ultrasound may improve diagnostic accuracy of physical examination. *Journal of the American Society of Echocardiography*, 27(11):1238–1246, 2014.
- [7] Zizhao Zhang, Yuanpu Xie, Fuyong Xing, et al. Mdnnet: A semantically and visually interpretable medical image diagnosis network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6428–6436, 2017.
- [8] Yuan Xue, Tao Xu, L Rodney Long, et al. Multimodal recurrent model with attention for automated radiology report generation. In *MICCAI*, pages 457–466. Springer, 2018.
- [9] Yuan Li, Xiaodan Liang, Zhiting Hu, et al. Hybrid retrieval-generation reinforced agent for medical image report generation. In *NeuIPS*, pages 1537–1547, 2018.
- [10] Baoyu Jing, Pengtao Xie, and Eric Xing. On the automatic generation of medical imaging reports. In *ACL*, pages 2577–2586, 2018.
- [11] Oriol Vinyals, Alexander Toshev, Samy Bengio, et al. Show and tell: A neural image caption generator. In *CVPR*, pages 3156–3164, 2015.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeuIPS*, pages 1097–1105, 2012.
- [13] Kaiming He, Xiangyu Zhang, ShaoqingRen, et al. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [14] Sepp Hochreiter and Schmidhuber J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [15] Jonathan Krause, Justin Johnson, Ranjay Krishna, et al. A hierarchical approach for generating descriptive image paragraphs. In *CVPR*, pages 317–325, 2017.
- [16] Alex Graves and J. Schmidhuber. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks*, 18(5-6):602–610, 2005.
- [17] Tomas Mikolov, Kai Chen, Greg Corrado, et al. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

- [18] Xinlei Chen, Hao Fang, Tsung-Yi Lin, et al. Microsoft coco captions: Data collection and evaluation server. arXiv preprint arXiv:1504.00325, 2015.
- [19] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, et al. Self-critical sequence training for image captioning. In CVPR, pages 7008–7024, 2017.