

SLAKE: A SEMANTICALLY-LABELED KNOWLEDGE-ENHANCED DATASET FOR MEDICAL VISUAL QUESTION ANSWERING

Bo Liu* Li-Ming Zhan** Li Xu** Lin Ma † Yan Yang ‡ Xiao-Ming Wu* †

* Department of Computing, The Hong Kong Polytechnic University, Hong Kong

† Department of Ultrasound, West China Hospital of Sichuan University, China

‡ Sichuan Academy of Medical Sciences, Sichuan Provincial People’s Hospital, China

ABSTRACT

Medical visual question answering (Med-VQA) has tremendous potential in healthcare. However, the development of this technology is hindered by the lacking of publicly-available and high-quality labeled datasets for training and evaluation. In this paper, we present a large bilingual dataset, SLAKE, with comprehensive semantic labels annotated by experienced physicians and a new structural medical knowledge base for Med-VQA. Besides, SLAKE includes richer modalities and covers more human body parts than the currently available dataset. We show that SLAKE can be used to facilitate the development and evaluation of Med-VQA systems. The dataset can be downloaded from <http://www.med-vqa.com/slake>.

Index Terms— Dataset, medical visual question answering, multi-modality fusion.

1. INTRODUCTION

Developing machines that can understand visual content and answer questions like humans is a long-standing goal of AI research. In recent years, visual question answering (VQA) has become an active field of research. Medical visual question answering (Med-VQA) is a domain-specific branch of VQA, where a clinical question comes with a radiology image and the goal is to design a system that can correctly answer the question based on the visual information of the image.

Med-VQA has a wide range of application prospects in healthcare sectors and a broad impact on the wellness of the general public. With a reliable Med-VQA system, patients can easily acquire information about their health and be more engaged in the process of decision making. For doctors, Med-VQA systems can be used to assist diagnosis by providing them a second medical opinion. The systems can also be used in clinical education to train medical professionals. Besides, Med-VQA technology can be potentially integrated into many conversational AI platforms to bring enormous benefits to healthcare industry.

*Equal contribution.

†Corresponding author.




Fig. 1. Exemplar image and questions of our SLAKE dataset.

However, the research of Med-VQA is at an early stage. Unlike VQA in the general domain, where large-scale high-quality datasets [2, 3] are available, there is a lack of publicly-available and well-annotated datasets for training and evaluating Med-VQA systems. To correctly answer a clinical question about a radiology image, it requires clinical expertise and domain-specific medical knowledge, which makes it difficult to construct a realistic and accurate dataset for Med-VQA. VQA-RAD [1] is a first step in this direction. To our knowledge, it is the only available dataset with manual annotation, based on which several Med-VQA models have been proposed [4, 5]. VQA-RAD is a diverse dataset containing a variety of different types of clinical questions, with each question type sufficiently represented. But it does not provide semantic labels, e.g., labeled segmentations of organs and tumors or bounding boxes on objects, which are essential for training a Med-VQA model to find the region of interest in an image to answer complex clinical questions. Moreover, a

Table 1. Comparison of SLAKE with VQA-RAD.

Dataset	# Images	# QA Pairs	Question Type	Language	Knowledge Graph
VQA-RAD [1]	315	3.5K	Vision-only	EN	✗
SLAKE (Ours)	642	14K	Knowledge-based & Vision-only	Bilingual (EN & ZH)	✓

practical Med-VQA system needs to exploit external knowledge apart from visual content to answer complex compositional questions involving inquires such as “the functionality of an organ”, “the cause of a disease”, or “the treatment of a disease”, which is also not supported in VQA-RAD.

To fill these gaps, we construct a **semantically-labeled knowledge-enhanced (SLAKE)** dataset with accurate visual and textual annotations and an extendable knowledge base for Med-VQA. It takes our team more than half of a year to complete all the tasks, including building the annotation system, constructing the medical knowledge graph (KG), selecting and labeling images, generating questions, and analyzing the dataset. As shown in Figure 1, for each radiology image, we provide two kinds of visual annotations: *masks* for semantic segmentation and *bounding boxes* for object detection. Besides basic clinical questions, we also design compositional questions that require multiple reasoning steps, and knowledge-based questions like [6] that involve external medical knowledge. In general, questions in SLAKE can be categorized as vision-only questions and knowledge-based questions. We provide detailed annotations to distinguish the two types of questions and guide the Med-VQA model to search for answers on the knowledge graph. Besides these new features, SLAKE is designed to be an English-Chinese bilingual dataset to broaden its application range. Further, SLAKE covers more body parts (e.g., neck and pelvic cavity) and more types of questions (e.g., shape and KG-related) than VQA-RAD. A comparison between our SLAKE and VQA-RAD is provided in Table 1.

In summary, our contributions are two-fold:

- We create SLAKE, a large-scale, semantically annotated, and knowledge-enhanced bilingual dataset for training and testing Med-VQA systems.
- We experiment with representative Med-VQA methods to show that SLAKE can be used as a benchmark to train systems to solve practical and complex tasks.

2. THE SLAKE DATASET

In this section, we elaborate on the construction of our SLAKE dataset. In general, we ensure the diversity of the dataset in terms of modalities (e.g., CT, MRI, and X-Ray), covered body parts (e.g., head, neck, and chest), and question types (e.g., vision-only, knowledge-based, and bilingual).

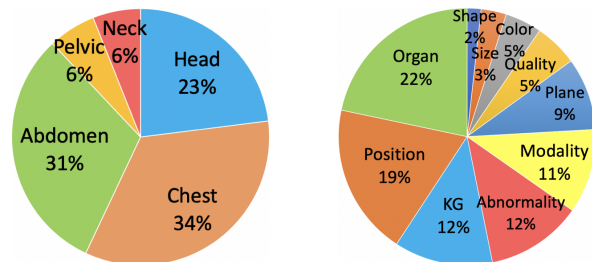


Fig. 2. Left: proportions of images of five body parts. Right: distribution of the content types of questions.

2.1. Image Acquisition and Annotation

We select radiology images, covering healthy and unhealthy cases, from three open source datasets [7]¹ [8]² [9]³. From [8], we randomly select 179 chest X-Ray images and keep the original disease labels. From [7] and [9], we randomly choose 463 single-slice images from 3D volume cases. Then, experienced physicians label organs and diseases as detailed as possible with ITK-SNAP [10]⁴ as shown in Figure 1.

In total, we annotate 642 images, including 12 diseases and 39 organs of the whole body. The diseases mainly include cancer (e.g., brain, liver, kidney, lung, etc.), and thoracic diseases (e.g., atelectasis, effusion, mass, pneumothorax, etc.). The images include 140 head CTs or MRIs, 41 neck CTs, 219 chest X-Rays or CTs, 201 abdomen CTs or MRIs, and 41 pelvic cavity CTs. The distribution is shown in Figure 2 (Left). Among these images, there are 282 CTs, 181 MRIs, and 179 X-Rays. All CTs and MRIs are axial single-slice. The number of images for each body part is set based on the complexity of the body part. For example, the number of diseases and organs in abdomen is much more than that in neck, so there are more images of abdomen than neck in the dataset.

2.2. Knowledge Graph Construction

To answer questions that require external medical knowledge, we construct a medical knowledge graph centered on organs and related diseases, which are the main objects of radiology images. We extract a set of 52.6K triplets $\langle head, relation, tail \rangle$ with medical knowledge from OwnThink⁵, a large-scale

¹<http://medicaldecathlon.com>

²<https://nihcc.app.box.com/v/ChestXray-NIHCC>

³<https://doi.org/10.5281/zenodo.3431873>

⁴<http://www.itksnap.org>

⁵<https://www.ownthink.com>

Table 2. Statistics of questions in our SLAKE dataset.

	Training set	Validation set	Test set
Plane	931	173	176
Quality	535	109	118
Modality	1072	203	217
Position	1876	412	390
Organ	2125	462	454
KG	1202	278	260
Abnormal	1230	245	221
Color	424	108	115
Shape	157	42	46
Size	297	77	73
Total	9849	2109	2070

knowledge base built on Wikipedia. Here, *head* and *tail* are entities such as organ, disease, etc., and *relation* represents the relationship between entities, such as function or treatment. Then, we traverse the set to retrieve triplets related to organs and the corresponding diseases. We further clean the data by manually filtering out some entities that are not presented in medical images such as gastritis and nephritis.

Next, in order to extensively cover frequently referenced knowledge, we refine the filtered triplets with the following rules: (1) The triplets about an organ must describe its function or body system; (2) The triplets about a disease must describe the symptoms, locations, causes, treatment or prevention methodologies. Some examples are shown in Table 3.

Finally, we make the triplets bilingual and obtain 2603 triplets in English and 2629 triplets in Chinese.

2.3. Question Generation

Questions are proposed by experienced doctors. To accelerate this process, we develop an annotation system. In this system, we first pre-define a question template for each body part (i.e., head, neck, chest, abdomen, and pelvic cavity). Then, we define ten different content types (e.g., *modality*, *position*, *color*) for the questions, as shown in Table 2 and Figure 2 (Right). In each template, we provide many candidate questions for each content type. For example, the candidate question for a head image with the content type *organ* may be “Is this a study of the head?” or “What organ system is imaged?”. Physicians could choose those candidate questions or amend or even rewrite them entirely based on their personal clinical experience. The flexibility of our annotation system ensures the question diversity of SLAKE. Note that because we provide different candidates for bilingual questions, the number and content of them in our dataset are not the same.

Moreover, we provide semantic label for each question. Specifically, we use $\langle vhead, _ , _ \rangle$ (*vhead* is a placeholder) to denote vision-only questions. For a knowledge-based question like “Which organs in this image belong to the digestive system?”, we denote it as $\langle vhead, belong\ to, digestive\ sys-$

Table 3. Examples of our medical knowledge graph.

	Examples
Organ	$\langle Heart, Function, Promote\ blood\ flow \rangle$
	$\langle Kidney, Belong\ to, Urinary\ System \rangle$
	$\langle Duodenum, Length, 20-25cm \rangle$
Disease	$\langle Pneumonia, Location, Lung \rangle$
	$\langle Lung\ Cancer, Cause, Smoke \rangle$
	$\langle Brain\ Tumor, Symptom, Visual\ impairment \rangle$
	$\langle Cardiomegaly, Treatment, Medication \rangle$
	$\langle Atelectasis, Prevention, Exercise \rangle$

tem⟩. Such labeling helps to distinguish question type and identify the part of the question involving external knowledge.

Besides, recent studies [11, 12] have shown that VQA models may be susceptible to the statistical bias of answer distribution of the datasets. To mitigate the inherent bias of SLAKE, we make the answers balanced in general such that the VQA model will not be biased to the most popular answer in the dataset. For example, for the question “Is this a study of the abdomen?”, we make sure this question is asked with abdomen images and non-abdomen images with 50–50 chance, thereby keeping the numbers of “Yes” and “No” balanced.

2.4. Dataset Splitting

Here, we describe how to divide the obtained 642 images with 14,028 question-answer pairs and 5232 medical knowledge triplets for the training and evaluation of Med-VQA models.

In general, the splitting aims to provide a reliable measure of the generalization ability of the model trained on our dataset. Specifically, we split the dataset into training (70%), validation (15%), and test (15%) sets at the image level. The images in our dataset are split with the 75:15:15 ratio in each of the 8 categories: “head CT”, “head MRI”, “neck CT”, and “chest X-Ray”, “chest CT”, “abdomen CT”, “abdomen MRI”, and “pelvic cavity CT”. Note that we only divide the images but the questions associated with each image are not split.

Besides, since VQA is usually formulated as a classification task [4, 5, 13], we follow the convention and make sure answers in the test set must appear in the training set. Finally, the images are split into 450 for training, 96 for validating, and 96 for testing. The number of questions of different type in each set is shown in Table 2.

3. EXPERIMENTS

In this section, we conduct extensive experiments to comprehensively evaluate our SLAKE dataset. To be elaborated later, Table 4 demonstrates the usefulness and the challenge of SLAKE with commonly used Med-VQA methods. To show the effectiveness of the constructed medical knowledge graph, we conduct an ablation study presented in Table 5.

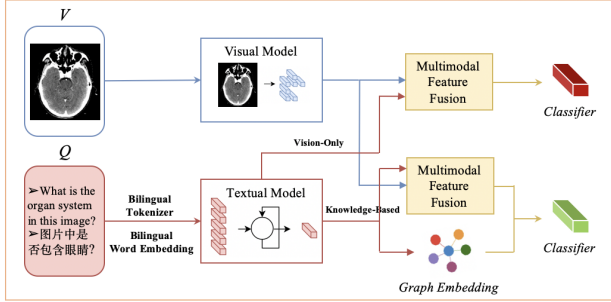


Fig. 3. The Med-VQA framework on our SLAKE dataset.

Table 4. Accuracy for vision-only questions (%).

Language	Models	Overall	Open-ended	Closed-ended
English	VGG+SAN	72.73	70.34	76.13
	VGG _{seg} +SAN	75.36	72.20	79.84
Chinese	VGG+SAN	74.27	73.64	75.20

3.1. Experiment Setup

The pipeline of our experiments is illustrated in Figure 3. We experiment with a commonly used Med-VQA framework, stacked attention network (SAN) [13], on SLAKE. We use VGG16 [14] to extract visual features from radiology images. For bilingual questions, we first design a bilingual tokenizer to create bilingual word embeddings for the English questions and Chinese questions respectively. Then, a 1024D-LSTM is applied to extract textual semantics from these embeddings and classify types of questions. There are two sub-pipelines in Figure 3. Given the extracted visual and textual features, vision-only tasks will be directed to the multimodal fusion module of SAN to create fused features for classification. For knowledge-based tasks, question-related embeddings extracted from the knowledge graph will be combined with the multimodal fused features for classification.

3.2. Dataset Analysis

We report the results for vision-only and knowledge-based questions in Table 4 and Table 5 respectively. Answers of “closed-ended” questions are limited multiple-choice options, while answers of “open-ended” questions are free-form texts. Open-ended questions are generally harder to answer than closed-ended ones.

Vision-only questions. In Table 4, we report the results in accuracy for vision-only questions in both English and Chinese. Compared with VQA in the general domain, clinical questions in Med-VQA need to be answered as accurate as possible because they relate to health and safety. It can be seen that the baseline models achieve accuracy of around 73% which is still far away from practical use in the medical do-

Table 5. Accuracy for knowledge-based questions (%).

Language	Models	Overall
English	VGG+SAN	70.27
	VGG+SAN+KG	72.30
Chinese	VGG+SAN+KG	75.01

main. There is a wide gap between this and clinical standard, which shows that SLAKE is challenging. Moreover, it can be seen that the overall accuracy is roughly the average of those of open-ended and closed-ended questions, proving that the question distribution of SLAKE is balanced.

Besides, to demonstrate the usefulness of the semantic visual annotations elaborated in Section 2.1, we design another model, VGG_{seg}+SAN. First, we pretrain a fully convolutional network (FCN) with VGG backbone by the segmentation task of radiology images with respect to the *mask* labels in the training set. Then, we initialize the VGG backbone in the Med-VQA model with the pretrained parameters. The overall accuracy increases from 72.73% to 75.36% with a 2.6% improvement, which shows that our semantic visual annotations could improve the reasoning abilities of the model.

Knowledge-based questions. We leverage the self-built medical knowledge graph to answer knowledge-based questions. First, we randomly initialize an embedding for each entity in the knowledge graph and use the TransE [15] method to enforce the embeddings of the entities in each triplet, $\langle head, relation, tail \rangle$, to satisfy: $head + relation \approx tail$. Then, based on the semantic textual annotations (Section 2.3), we train two LSTMs to predict the words for the “relation” and “tail” of a question separately. Next, we find the corresponding entity embeddings of the relation and tail from the graph and use them to obtain the head entity embedding based on the above approximate equation, which is then combined with the fused multimodal features for final prediction. The result is reported in Table 5. For comparison, we also try to predict answers without using the knowledge graph. The result is 2.0% lower, indicating that the constructed knowledge graph is informative and it is helpful to leverage external structural knowledge to tackle knowledge-based questions.

4. CONCLUSION

We have introduced SLAKE, a new large bilingual dataset to facilitate the training and evaluation of Med-VQA systems. SLAKE is a diverse and balanced dataset containing rich visual and textual annotations and a unique medical knowledge graph, which allows the development of more powerful Med-VQA systems. Remarkably, our experiments show that the semantic annotations and external knowledge can significantly improve the performance of standard Med-VQA models. We hope SLAKE will serve as a stepping stone to push forward the research of Med-VQA.

5. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access by [7, 9, 8]. Ethical approval was not required as confirmed by the license attached with the open access data.

6. ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their helpful comments. Thanks to Lau et al [1] for their pioneering work in Med-VQA, NIH Clinical Center for sharing their open access dataset [8], and all the doctors and medical students who helped with this research. This research was supported by the grant of P0030935 (ZVPY) funded by PolyU (UGC).

7. REFERENCES

- [1] Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman, “A dataset of clinically generated visual questions and answers about radiology images,” *Scientific data*, vol. 5, no. 1, pp. 1–10, 2018.
- [2] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick, “Clevr: A diagnostic dataset for compositional language and elementary visual reasoning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2901–2910.
- [3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh, “Vqa: Visual question answering,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2425–2433.
- [4] Binh D Nguyen, Thanh-Toan Do, Binh X Nguyen, Tuong Do, Erman Tjiputra, and Quang D Tran, “Overcoming data limitation in medical visual question answering,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 522–530.
- [5] Li-Ming Zhan, Bo Liu, Lu Fan, Jiabin Chen, and Xiaoming Wu, “Medical visual question answering via conditional reasoning,” in *Proceedings of the 28th ACM International Conference on Multimedia*, New York, NY, USA, 2020, MM ’20, p. 2345–2354, Association for Computing Machinery.
- [6] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel, “Fvqa: Fact-based visual question answering,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 10, pp. 2413–2427, 2018.
- [7] Amber L Simpson, Michela Antonelli, Spyridon Bakas, Michel Bilello, Keyvan Farahani, Bram Van Ginneken, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, et al., “A large annotated medical image dataset for the development and evaluation of segmentation algorithms,” *arXiv preprint arXiv:1902.09063*, 2019.
- [8] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers, “Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2097–2106.
- [9] Ali Emre Kavur, M. Alper Selver, Oğuz Dicle, Mustafa Barış, and N. Sinem Gezer, “CHAOS - Combined (CT-MR) Healthy Abdominal Organ Segmentation Challenge Data,” Apr. 2019.
- [10] Paul A. Yushkevich, Joseph Piven, Heather Cody Hazlett, Rachel Gimpel Smith, Sean Ho, James C. Gee, and Guido Gerig, “User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability,” *Neuroimage*, vol. 31, no. 3, pp. 1116–1128, 2006.
- [11] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh, “Analyzing the behavior of visual question answering models,” *arXiv preprint arXiv:1606.07356*, 2016.
- [12] Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh, “Yin and yang: Balancing and answering binary visual questions,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5014–5022.
- [13] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola, “Stacked attention networks for image question answering,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 21–29.
- [14] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [15] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko, “Translating embeddings for modeling multi-relational data,” in *Advances in neural information processing systems*, 2013, pp. 2787–2795.