

# Leveraging Medical Visual Question Answering with Supporting Facts

Tomasz Kornuta, Deepta Rajan, Chaitanya Shivade,  
Alexis Asseman, and Ahmet S. Ozcan

IBM Research AI, Almaden Research Center, San Jose, USA  
{tkornut, drajan, cshivade, asozcan}@us.ibm.com, alexis.asseman@ibm.com

**Abstract.** In this working notes paper, we describe IBM Research AI (Almaden) team’s participation in the ImageCLEF 2019 VQA-Med competition. The challenge consists of four question-answering tasks based on radiology images. The diversity of imaging modalities, organs and disease types combined with a small imbalanced training set made this a highly complex problem. To overcome these difficulties, we implemented a modular pipeline architecture that utilized transfer learning and multi-task learning. Our findings led to the development of a novel model called Supporting Facts Network (SFN). The main idea behind SFN is to cross-utilize information from upstream tasks to improve the accuracy on harder downstream ones. This approach significantly improved the scores achieved in the validation set (18 point improvement in F-1 score). Finally, we submitted four runs to the competition and were ranked seventh.

**Keywords:** ImageCLEF 2019 · VQA-Med · Visual Question Answering · Supporting Facts Network · Multi-Task Learning · Transfer Learning

## 1 Introduction

In the era of data deluge and powerful computing systems, deriving meaningful insights from heterogeneous information has shown to have tremendous value across industries. In particular, the promise of deep learning-based computational models [15] in accurately predicting diseases has further stirred great interest in adopting automated learning systems in healthcare [2]. A daunting challenge within the realm of healthcare is to efficiently sieve through vast amounts of multi-modal information and reason over them to arrive at a differential diagnosis. Longitudinal patient records including time-series measurements, text reports and imaging volumes form the basis for doctors to draw conclusive insights. In practice, radiologists are tasked with reviewing thousands of imaging studies each day, with an average of about three seconds to mark them as

---

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). CLEF 2019, 9-12 September 2019, Lugano, Switzerland.

anomalous or not, leading to severe eye fatigue [24]. Moreover, clinical workflows have a sequential nature tending to cause delays in triage situations, where the existence of answers to key questions about a patient’s holistic conditions can potentially expedite treatment. Thus, building effective question-answering systems for the medical domain by bringing advancements in machine learning research will be a game changer towards improving patient care.

Visual Question Answering (VQA) [17, 1] is a new exciting problem domain, where the system is expected to answer questions expressed in natural language by taking into account the content of the image. In this paper, we present results of our research on the VQA-Med 2019 dataset [3], an open challenge associated with the ImageCLEF 2019 initiative [10]. The main issue here, in comparison to the other recent VQA datasets such as TextVQA [23] or GQA [9], is dealing with scattered, noisy and heavily biased data. Hence, the dataset serves as a great use-case to study challenges encountered in practical clinical scenarios.

In order to address the data issues, we designed a new model called **Supporting Facts Network** (SFN) that efficiently shares knowledge between upstream and downstream tasks through the use of a pre-trained multi-task solver in combination with task-specific solvers. Note that posing the VQA-Med challenge as a multi-task learning problem [4] allowed the model to effectively leverage and encode relevant domain knowledge. Our multi-task SFN model outperforms the single task baseline by better adapting to label distribution shifts.

## 2 The VQA-Med dataset

The VQA-Med 2019 [3] is a Visual Question Answering (VQA) dataset embedded in the medical domain, with a focus on radiology images. It consists of:

- a training set of 3,200 images with 12,792 Question-Answer (QA) pairs,
- a validation set of 500 images with 2,000 QA pairs, and
- a test set of 500 images with 500 questions (answers were released after the end of the VQA-Med 2019 challenge).

In all splits the samples were divided into four categories, depending on the main task to be solved:

- **C1**: determine the modality of the image,
- **C2**: determine the plane of the image,
- **C3**: identify the organ/anatomy of interest in the image, and
- **C4**: identify the abnormality in the image.

Our analysis of the dataset (distribution of questions, answers, word vocabularies, categories and image sizes) led to the following findings and system-design related decisions:

- merge of the original training and validation sets, shuffle and re-sample new training and validation sets with a proportion of 19:1,
- use of weighted random sampling during batch preparation,

- addition of a fifth **Binary** category for samples with Y/N type questions,
- focus on accuracy-related metrics instead of the BLEU score,
- avoid label (answer classes) unification and cleansing,
- consider **C4** as a downstream task and exclude it from the pre-training of input fusion modules,
- utilization of image size as an additional input cue to the system.

### 3 Supporting Facts Network

Typical VQA systems process two types of input, visual (image) and language (question), that need to undergo various transformations to produce the answer. Fig. 1 presents a general architecture of such systems, indicating four major modules: two encoders responsible for encoding raw inputs to more useful representations, followed by a reasoning module that combines them and finally, an answer decoder that produces the answer.





Fig. 1. General architecture of Visual Question Answering systems

In the early prototypes of VQA systems, reasoning modules were rather simple and relied mainly on multi-modal fusion mechanisms. These fusion techniques varied from concatenation of image and question representations, to more complex pooling mechanisms such as Multi-modal Compact Bilinear pooling (MCB) [6] and Multi-modal Low-rank Bilinear pooling (MLB) [12]. Further, diverse attention mechanisms such as question-driven attention over image features [11] were also used. More recently, researchers have focused on complex multi-step reasoning mechanisms such as Relational Networks [21, 5] and Memory, Attention and Composition (MAC) networks [8, 18]. Despite that, certain empirical studies indicate early fusion of language and vision signals significantly boosts the overall performance of VQA systems [16]. Therefore, we explored the finding of an "optimal" module for early fusion of multi-modal inputs.

#### 3.1 Architecture of the Input Fusion Module

One of our findings from analyzing the dataset was to use the image size as additional input cue to the system. This insight triggered an extensive architecture search that included, among others, comparison and training of models with:

- different methods for question encoding, from 1-hot encoding with Bag-of-Words to different word embeddings combined with various types of recurrent neural networks,
- different image encoders, from simple networks containing few convolutional layers trained from scratch to fine-tuning of selected state-of-the-art models pre-trained on ImageNet,
- various data fusion techniques as mentioned in the previous section.



**Fig. 2.** Architectures of two modules used in the final system

The final architecture of our model is presented in Fig. 2a. We used GloVe word embeddings [20] followed by Long Short-Term Memory (LSTM) [7]. The LSTM outputs along with feature maps extracted from images using VGG-16 [22] were passed to the **Fusion I** module, implementing question-driven attention over image features [11]. Next, the output of that module was concatenated in the **Fusion II** module with image size representation created by passing image width and height through a fully connected (FC) layer.

Note that the green colored modules were initially pre-trained on external datasets (ImageNet and 6B tokens from Wikipedia 2014 and Gigaword 5 datasets for VGG-16 and GloVe models respectively) and later fine-tuned during training on the VQA-Med dataset.

### 3.2 Architectures of the Reasoning Modules

During the architecture search of the **Input Fusion** module we used the model presented in Fig. 3, with a simple classifier with two FC layers. These models were trained and validated on **C1**, **C2** and **C3** categories separately, while excluding **C4**. In fact, to test our hypothesis we trained some early prototypes only on samples from **C4** and the models failed to converge.

After establishing the **Input Fusion** module we trained it on samples from **C1**, **C2** and **C3** categories. This served as a starting point for training more complex reasoning modules. At first, we worked on a model that exploited information about 5 categories of questions by employing 5 separate classifiers which used data produced by the **Input Fusion** module. Each of these classifiers essentially specialized in one question category and had its own answer




Fig. 3. Architecture with a single classifier (IF-1C)

label dictionary and associated loss function. The predictions were then fed to the **Answer Fusion** module, which selected answers from the right classifier based on the question category predicted by the **Question Categorizer** module, whose architecture is shown in Fig. 2b. Please note that we pre-trained the module in advance on all samples from all categories and froze its weights during the training of classifiers.

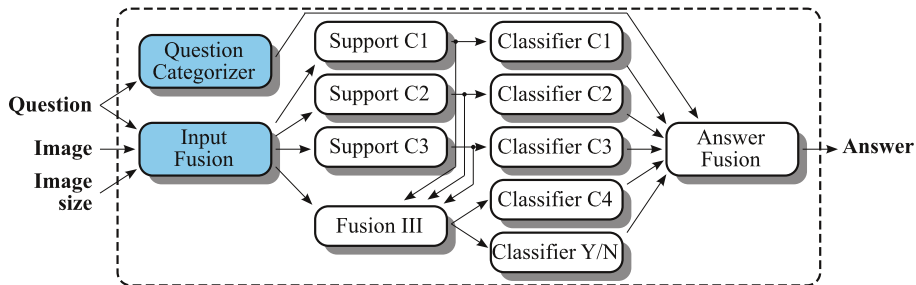


Fig. 4. Final architecture of the Supporting Facts Network (SFN)

The architecture of our final model, **Supporting Facts Network** is presented in Fig. 4. The main idea here resulted from the analysis of questions about the presence of abnormalities – to answer which the system required knowledge on image modality and/or organ type. Therefore, we divided the classification modules into two networks: Support networks (consisting of two FC layers) and final classifiers (being single FC layers). We added Plane (**C2**) as an additional supporting fact. The supporting facts were then concatenated with output from **Input Fusion** module in **Fusion III** and passed as input to the classifier specialized on **C4** questions. In addition, since **Binary Y/N** questions were present in both **C1** and **C4** categories, we followed a similar approach for that classifier.

## 4 Experimental Results

All experiments were conducted using PyTorchPipe [14], a framework that facilitates development of multi-modal pipelines built on top of PyTorch [19]. Our models were trained using relatively large batches (256), dropout (0.5) and Adam

optimizer [13] with a small learning rate ( $1e-4$ ). For each experimental run, we generated a new training and validation set by combining the original sets, shuffling and sampling them in proportions of 19 : 1, thereby resulting in a validation set of size 5%.

Model	Resampled Valid. Set			Original Train. Set			Original Valid. Set		
	Prec.	Recall	F-1	Prec.	Recall	F-1	Prec.	Recall	F-1
IF-1C	0.630	0.435	0.481	0.683	0.497	0.545	0.690	0.499	0.548
SFN	0.759	0.758	0.758	0.753	0.692	0.707	0.762	0.704	0.717

**Table 1.** Summary of experimental results. All columns contain average scores achieved by 5 separately trained models on resampled training and validation sets. We also present scores achieved by the models on original sets (in the evaluation mode).

In Tab. 1 we present a comparison of average scores achieved by our baseline models using single classifier (IF-1C) and the Supporting Facts Networks (SFN). Our results clearly indicate the advantage of using *'supporting facts'* over the baseline model with a single classifier. The SFN model by our team achieved a best score of (0.558 Accuracy, 0.582 BLEU score) on the test set as indicated by the CrowdAI leaderboard. One of the reasons for such a significant drop in performance is due to the presence of new answers classes in the test set that were not present both in the original training and validation sets.

## 5 Summary

In this work, we introduced a new model called **Supporting Facts Network** (SFN), that leverages knowledge learned from combinations of upstream tasks in order to benefit additional downstream tasks. The model incorporates domain knowledge that we gathered from a thorough analysis of the dataset, resulting in specialized input fusion methods and five separate, category-specific classifiers. It comprises of two pre-trained shared modules followed by a reasoning module jointly trained with five classifiers using the multi-task learning approach. Our models were found to train faster and to deal much better with label distribution shifts under a small imbalanced data regime.

Among the five categories of samples present in the VQA-Med dataset, **C4** and **Binary** turned out to be extremely difficult to learn, for several reasons. First, there were 1483 unique answer classes assigned to 3082 training samples related to **C4**. Second, both **C4** and **Binary** required more complex reasoning and, besides, might be impossible to conclude by looking only at the question and content of the image. However, our observation that some of the information from simpler categories might be useful during reasoning on more complex ones, we refined the model by adding supporting networks. Given, modality, imaging

plane and organ typically help narrow down the scope of disease conditions and/or answer whether or not an abnormality is present. Our empirical studies prove that this approach performs significantly better, leading to an 18 point improvement in F-1 score over the baseline model on the original validation set.

## References

1. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., Parikh, D.: VQA: Visual Question Answering. In: Proceedings of the IEEE international conference on computer vision. pp. 2425–2433 (2015)
2. Ardila, D., Kiraly, A.P., Bharadwaj, S., Choi, B., Reicher, J.J., Peng, L., Tse, D., Etemadi, M., Ye, W., Corrado, G., Naidich, D.P., Shetty, S.: End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nature Medicine* (2019)
3. Ben Abacha, A., Hasan, S.A., Datla, V.V., Liu, J., Demner-Fushman, D., Müller, H.: VQA-Med: Overview of the medical visual question answering task at imageclef 2019. In: CLEF 2019 Working Notes. CEUR Workshop Proceedings, CEUR-WS.org < <http://ceur-ws.org/> >, Lugano, Switzerland (September 09-12 2019)
4. Caruana, R.: Multitask learning. *Machine learning* **28**(1), 41–75 (1997)
5. Desta, M.T., Chen, L., Kornuta, T.: Object-based reasoning in VQA. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1814–1823. IEEE (2018)
6. Fukui, A., Park, D.H., Yang, D., Rohrbach, A., Darrell, T., Rohrbach, M.: Multi-modal compact bilinear pooling for visual question answering and visual grounding. In: EMNLP (2016)
7. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* **9**(8), 1735–1780 (1997)
8. Hudson, D.A., Manning, C.D.: Compositional attention networks for machine reasoning. In: CVPR (2018)
9. Hudson, D.A., Manning, C.D.: Gqa: a new dataset for compositional question answering over real-world images. arXiv preprint arXiv:1902.09506 (2019)
10. Ionescu, B., Müller, H., Péteri, R., Cid, Y.D., Liauchuk, V., Kovalev, V., Klimuk, D., Tarasau, A., Ben Abacha, A., Hasan, S.A., Datla, V., Liu, J., Demner-Fushman, D., Dang-Nguyen, D.T., Piras, L., Riegler, M., Tran, M.T., Lux, M., Gurrin, C., Pelka, O., Friedrich, C.M., de Herrera, A.G.S., Garcia, N., Kavallieratou, E., del Blanco, C.R., Rodríguez, C.C., Vasilopoulos, N., Karampidis, K., Chamberlain, J., Clark, A., Campello, A.: ImageCLEF 2019: Multimedia retrieval in medicine, lifelogging, security and nature. In: Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the 10th International Conference of the CLEF Association (CLEF 2019), LNCS Lecture Notes in Computer Science, Springer, Lugano, Switzerland (September 9-12 2019)
11. Kazemi, V., Elqursh, A.: Show, ask, attend, and answer: A strong baseline for visual question answering. arXiv preprint arXiv:1704.03162 (2017)
12. Kim, J.H., On, K.W., Lim, W., Kim, J., Ha, J.W., Zhang, B.T.: Hadamard product for low-rank bilinear pooling. In: ICLR (2017)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
14. Kornuta, T.: PyTorchPipe. <https://github.com/ibm/pytorchpipe> (2019)
15. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *nature* **521**(7553), 436 (2015)

16. Malinowski, M., Doersch, C.: The Visual QA devil in the details: The impact of early fusion and batch norm on CLEVR. In: ECCV'18 Workshop on Shortcomings in Vision and Language (2018)
17. Malinowski, M., Fritz, M.: A multi-world approach to question answering about real-world scenes based on uncertain input. In: Advances in neural information processing systems. pp. 1682–1690 (2014)
18. Marois, V., Jayram, T., Albouy, V., Kornuta, T., Bouhadjar, Y., Ozcan, A.S.: On transfer learning using a MAC model variant. In: NeurIPS'18 Visually-Grounded Interaction and Language (ViGIL) Workshop (2018)
19. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
20. Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). pp. 1532–1543 (2014)
21. Santoro, A., Raposo, D., Barrett, D.G., Malinowski, M., Pascanu, R., Battaglia, P., Lillicrap, T.: A simple neural network module for relational reasoning. In: Advances in Neural Information Processing Systems. pp. 4967–4976 (2017)
22. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
23. Singh, A., Natarajan, V., Shah, M., Jiang, Y., Chen, X., Batra, D., Parikh, D., Rohrbach, M.: Towards vqa models that can read. arXiv preprint arXiv:1904.08920 (2019)
24. Syeda-Mahmood, T., Walach, E., Beymer, D., Gilboa-Solomon, F., Moradi, M., Kisilev, P., Kakrania, D., Compas, C., Wang, H., Negahdar, R., et al.: Medical sieve: a cognitive assistant for radiologists and cardiologists. In: Medical Imaging 2016: Computer-Aided Diagnosis. vol. 9785, p. 97850A. International Society for Optics and Photonics (2016)