

# SEGMENTATION-GUIDED ATTENTION FOR VISUAL QUESTION ANSWERING FROM REMOTE SENSING IMAGES

Lucrezia Tosato<sup>\*a,b</sup>, Hichem Boussaid<sup>\*,a</sup>, Flora Weissgerber<sup>b</sup>,  
Camille Kurtz<sup>a</sup>, Laurent Wendling<sup>a</sup>, Sylvain Lobry<sup>a</sup>

<sup>a</sup>LIPADE, Université Paris Cité, 75006 Paris, France

<sup>b</sup>DTIS, ONERA, Université Paris Saclay, FR-91123 Palaiseau, France

## ABSTRACT

Visual Question Answering for Remote Sensing (RSVQA) is a task that aims at answering natural language questions about the content of a remote sensing image. The visual features extraction is therefore an essential step in a VQA pipeline. By incorporating attention mechanisms into this process, models gain the ability to focus selectively on salient regions of the image, prioritizing the most relevant visual information for a given question. In this work, we propose to embed an attention mechanism guided by segmentation into a RSVQA pipeline. We argue that segmentation plays a crucial role in guiding attention by providing a contextual understanding of the visual information, underlying specific objects or areas of interest. To evaluate this methodology, we provide a new VQA dataset that exploits very high-resolution RGB orthophotos annotated with 16 segmentation classes and question/answer pairs. Our study shows promising results of our new methodology, gaining almost 10% of overall accuracy compared to a classical method on the proposed dataset.

**Index Terms**— Visual Question Answering, Attention, Segmentation, Natural Language Processing

## 1. INTRODUCTION

Visual Question Answering (VQA) is designed to deliver natural language answers to open-ended queries about the content of an image [1]. Remote Sensing Visual Question Answering (RSVQA) is the application of VQA to remote sensing images, introduced in [2]. Recent advancements in computer vision and natural language processing have elevated the performance of standards on both VQA and RSVQA benchmarks. Notably, Large Language Models have become capable of executing knowledge-based VQA, utilizing external knowledge and commonsense reasoning [3]. Large language models have been used in RSVQA and

showed great potential to encode the question and to fuse the question with the classes present in the satellite images [4]. Additionally, transformers, such as VisualBERT, have been effectively used to fuse the image and the text [5]. Attention mechanisms have also been used in RSVQA to enhance the features by considering the alignments between spatial positions and words [6].

In the field of computer vision, [7] and [8] demonstrate that incorporating attention mechanisms enhances outcomes compared to relying solely on the image.

In natural images, semantic segmentation has been used to guide visual question answering towards the object of interest [9], as well as to guide attention [10]. To the best of our knowledge, these techniques have not yet been used in the field of remote sensing visual question answering.

Our contribution in this work is to propose to embed a segmentation-guided attention mechanism into a RSVQA pipeline. To compute attention weights we consider a multi-channel semantic segmentation, which differs from traditional one-channel semantic segmentation maps by allowing overlapping of objects of different classes. We showcase the interest of this methodological contribution on a new RSVQA dataset centered on the Ile-de-France region of France. It contains very high resolution airborne images, segmentation annotations as well as the automatically constructed question/answer pairs.

The remainder of this article is organized as follows. The proposed dataset is first described in section 2. Section 3 then introduced our novel RSVQA pipeline, including segmentation-guided attention. Finally, the results are presented and discussed in section 4.

## 2. DATASET

### 2.1. Very high resolution orthophotos (BD ORTHO) with segmentation annotation

BD ORTHO is a database of aerial optical images at a resolution of 20cm obtained from the French National Geographic Institute (IGN). The Very High Resolution (VHR) RGB-patches are derived from the subdivision of the BD ORTHO

\* Lucrezia Tosato and Hichem Boussaid contributed equally.

This work is supported by Agence Nationale de la Recherche (ANR) under the ANR-21-CE23-0011 project.

The experiments conducted in this study were performed using HPC/AI resources provided by GENCI-IDRIS (Grant 2023-AD011012735R2).

tiles into patches measuring  $1000 \times 1000$  pixels (equivalent to  $200m \times 200m$ ).

The IGN also provides a vectorial description of the French territory in the BD TOPO database. From the latter, we extract 16 classes for our multi-class segmentation annotations: Buildings (Building, Cemetery, Sports Field, Water Tank, Pylon, Surface Construction), Land Use (Foreshore Zone, Vegetation Zone), Water Area, Transport (Airfield, Transportation Construction, Road, Railway), Regulated Areas (Public Forest, National Park), and Services and Activities (that encompasses Museums, Monuments, Schools, etc).

We select four adjacent departments from the Ile-de-France region: Paris, Hauts-de-Seine, Val-de-Marne and Seine-Saint-Denis, resulting in 16'274 patches of size  $1000 \times 1000$  pixels.


## 2.2. Question/answer pairs

Following a procedure similar to [11], we propose an automated approach to generate sets of questions and answers linked to individual VHR patches. This method fully leverages the BD TOPO database, encompassing both general geographical features, such as buildings and water areas, as well as more specific entities like museums and lakes.

For a given VHR patch  $p_{VHR}$ , we retrieve from the BD TOPO the collection of geo-located objects that are present in the geographical extent of  $p_{VHR}$ . The objects are characterized by one element present in BD TOPO that we call a class, e.g. "road". We define nine question types, divided in four categories:

1. **Class questions** – that consider only one class of objects at a time. The questions are divided into **presence** (a), **count** (b) and **density** (c);
2. **Objects questions** – that consider the **absolute location** in the image (a) or the **area** (b) of a specific object;
3. **Two-class questions** – that **compare** the number of objects of two different classes;
4. **Object relation questions** – that consider the **relative location** (a) and the **distance** (b) of two specific objects from two different classes, or the absolute location of the **nearest** object from one class to a pre-selected object from another class or a pixel position (c).

One of the challenges of constructing a VQA dataset stochastically is to balance the question type and the answer type. This is a requirement to reduce language biases [12]. To balance the question type, we first randomly generate 10 questions per question type for each VHR patch  $p_{VHR}$ . Then, to balance the answer type, we define a maximum-number-of-questions per answer type  $N_{Q,A}$  for each question type. Only the questions with an answer type less present than  $N_{Q,A}$  are kept. To get a sufficient number of questions, the VHR



**Fig. 1.** Distribution of answers by question type. We omit numerical answers labeling and we show them ordered. The maximum numerical values are 280 (counting questions), 40000m2 (area questions), 273m (distance questions).

patches are browsed twice. Applying this procedure results in a total of 146'848 question/answer pairs. Their distribution is graphically represented in Figure 1.

## 2.3. Dataset splitting

The VHR patches are randomly attributed to the training, validation and test sets with a proportion of 60%, 20%, and 20% respectively. We use the same dataset splitting for both the segmentation auxiliary task and the complete pipeline.

## 3. METHOD

We introduce a novel approach for addressing Remote Sensing Visual Question Answering (RSVQA) by leveraging segmentation to direct the attention mechanism. The architecture of our pipeline is illustrated in Figure 2.

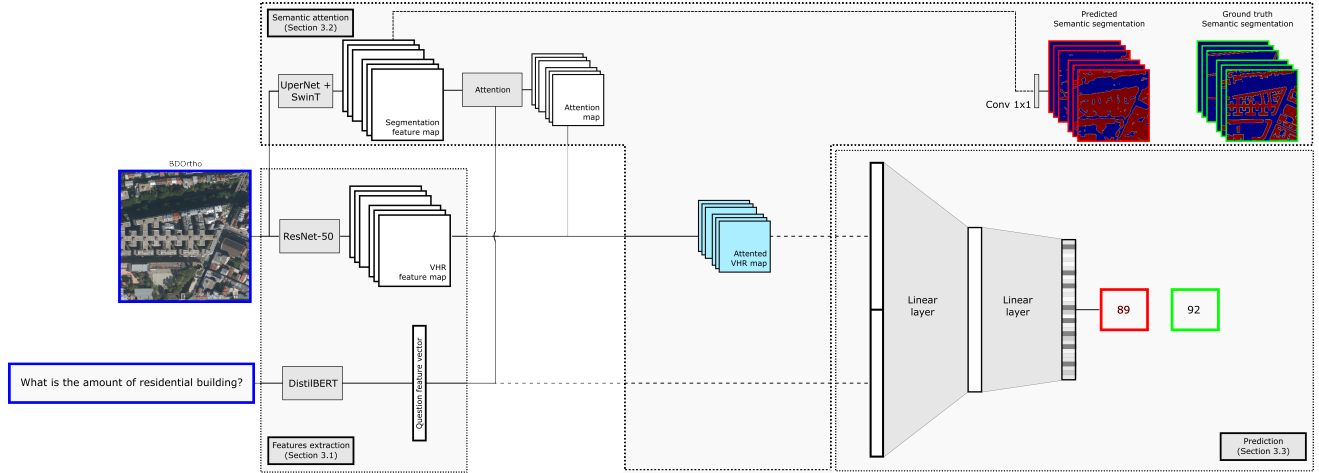
### 3.1. Features extraction

To extract visual features, noted  $f_{VHR}$ , from the VHR patches, we use a frozen ResNet-50 model [13] pre-trained on ImageNet from which we remove the last fully-connected layer.

The textual features  $f_q$  from the question  $q$  are extracted with a DistilBERT encoder [14], trained on the BookCorpus dataset [15] and frozen in the rest of this study.

### 3.2. Segmentation-guided attention

In our approach, we incorporate an auxiliary per-class segmentation task to guide the computation of the attention



**Fig. 2.** Graphical outline of the proposed architecture. The inputs (very high resolution remote sensing images and language questions) are shown in blue frames, the outputs in red frames (answer and segmentation) and the ground truths (answer and segmentation) in green frames.

weight. We assume that the auxiliary task allows a better initialization of the attention module by explicitly introducing a first layer of semantics in the features space. This allows us to facilitate the learning process of the attention module.

We first train the segmentation module by fine-tuning a UPerNet model with a Swin Transformer backbone introduced in [16]; pre-trained on ADE20K [17] on the 16 segmentation classes. To do so, we add a convolutional layer that maps the feature maps  $f_{seg}$  to the 16 channels output. This output is finally interpolated to the spatial size of  $p_{VHR}$  to obtain the segmentation result. The weights of the segmentation module are then kept frozen.

We obtain the attention weights by applying a linear layer (with dropout of  $d = 0.5$ ) on  $f_q$  to map it to a 250-dimensional space. Similarly, we use a  $1 \times 1$  convolution (with a dropout of  $d$ ) to obtain a 250 channels representation of  $f_{seg}$ . We concatenate both representations from the segmentation and the text, and apply a ReLU activation function followed by a convolution to obtain two attention glimpses  $a$ . Finally, we apply  $a$  to the visual features  $f_{VHR}$ .

### 3.3. Prediction

We concatenate the attended version of  $f_{VHR}$  with  $f_q$  to obtain a global representation of the inputs. This representation is mapped to a 1000-dimensional output (corresponding to the 1000 most frequent answers from the training set) using a 2-layer perceptron. As for the attention, we use the ReLU activation function and a dropout of  $d$ . We train the model using a cross-entropy loss, optimized with Adam, with a learning rate of  $10^{-6}$  and a batch size of 4 samples.

### 3.4. Evaluation

We evaluate the segmentation results with the overall precision / recall and the overall F1-score.

Three metrics are used to evaluate the VQA results: the per-question type accuracy, the overall accuracy (OA) and the average accuracy (AA). The per-question type accuracy is defined as the ratio of correct answer with the total number of questions for one of the nine question types. The OA is the ratio of correct answers with the total number of questions in the dataset. Finally, the AA is the average of the per-question type accuracies.

## 4. RESULTS AND DISCUSSION

### 4.1. Segmentation auxiliary task

Our model is trained using a NVIDIA GeForce RTX 4090 24G GPU. Our model achieves an overall precision and recall of 53.6% and 73.2% respectively and an overall F1-score of 60.1%. These results are obtained using a precision-recall curve with 20 different thresholds, to identify on the validation set the optimal threshold value for each class.

### 4.2. VQA task

Our model is trained for 100 hours, using one NVIDIA V100-16G GPU. The overall results for the VQA task are presented in Table 1. We observe an AA of 43.24% and an OA of 45.44% using our segmentation-guided attention.

To study the impact of the segmentation-guided attention we design two ablation studies. The first one uses a vanilla attention mechanism (i.e. not guided by the segmentation) and the second one does not use an attention mechanism.

Model	Model param.		Per-question accuracy									Dataset metrics	
	Att.	Seg.	1.(a)	1.(b)	1.(c)	2.(a)	2.(b)	3.	4.(a)	4.(b)	4.(c)	AA	OA
Proposed	✓	✓	<b>89.36</b>	<b>26.36</b>	<b>23.42</b>	<b>14.61</b>	<b>56.98</b>	<b>93.08</b>	<b>13.76</b>	<b>74.19</b>	<b>17.21</b>	<b>43.24</b>	<b>45.44</b>
A1	✓		85.79	22.98	20.99	13.82	44.69	82.94	11.91	74.19	15.60	39.40	41.43
A2			80.94	10.50	20.38	12.96	40.94	74.00	13.41	59.80	13.44	34.91	36.26

**Table 1.** Results of our proposed model and ablation studies. All of the results are accuracy percentages.

These ablation studies demonstrate that using segmentation-guided attention greatly improves the performances, with a gain of 10% in overall accuracy and 5% on the vanilla attention mechanism.

A visual example of our model is shown in Figure 3. In table 1 an improvement can be observed in performance using segmentation-guided attention across each of the question types. In almost all cases, the improvement is gradual, transitioning from A2 to A1 in the proposed method. This demonstrates that the attention mechanism enables the neural network to selectively focus on specific input parts, thereby assisting the RSVQA task. However, the auxiliary task of segmentation increases the level of supervision in the process, leading to better results.


We observe lower prediction accuracy for specific questions (2a, 4a, 4c) that involve determining objects' relative and absolute positions. This challenging task likely contributes to the decreased performance in these cases. It is indeed difficult to simultaneously recognize remote sensing objects and their spatial relationship from end-to-end only relying on present deep learning networks [18]. To increase the accuracy of such classes, having segmentation maps available, one could use histograms of forces to model the directional spatial relations between geo-localised objects as proposed in [19].

We can observe that in class density questions (1c), we have a slightly lower accuracy, which is an interesting result when compared to the performances of the area questions (2b). These two tasks, although similar, present a substantial difference: the area prediction is indeed made by numerical values greater than zero and integer, while the density predictions are all values from 0 to 1. We believe that the difference in the scale of results imposes a higher sensitivity and precision in the density predictions. Another reason might be that the number of density questions is smaller than the others.

## 5. CONCLUSIONS

In this study, we apply a segmentation-guided attention model in the context of RSVQA on a new dataset built from both BD ORTHO high-resolution images and BD TOPO for the segmentation annotation, as well as the questions/answers pairs.

From the preliminary results, we observe that segmentation succeeds in directing attention more effectively than attention alone. We believe that by using 16-channel segmentation, attention identifies the channels related to a specific word

	Question	Ground Truth	Prediction
	Is there an annexe building?	Yes	Yes
	Are there less residential building than field of hop plants?	No	No
	What is the area of the road intersection?	25.00m2	25.00m2
	Where is the closest annexe building to the transportation construction?	Left of the image	Left of the image

**Fig. 3.** Example of an image in department Hauts-de-Seine, 92 with questions, ground truths and predictions.

in the question and thus it is easier to locate the correct object in the final image. Further experiments with a more complete dataset are necessary to verify that the results can generalize to different geographical areas and more diverse questions.

## 6. REFERENCES

- [1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, "VQA: Visual Question Answering," in *IEEE/CVF ICCV*, pp. 2425–2433, 2015.
- [2] S. Lobry, D. Marcos, J. Murray, and D. Tuia, "RSVQA: Visual Question Answering for Remote Sensing Data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 12, pp. 8555–8566, 2020.
- [3] Y. Hu, H. Hua, Z. Yang, W. Shi, N. A. Smith, and J. Luo, "PROMPTCAP: Prompt-guided image captioning for VQA with GPT-3," in *IEEE/CVF ICCV*, pp. 2963–2975, 2023.
- [4] C. Chappuis, V. Zermatten, S. Lobry, B. Le Saux, and D. Tuia, "Prompt-RSVQA: Prompting visual context to a language model for remote sensing visual question answering," in *IEEE/CVF CVPR*, pp. 1372–1381, 2022.
- [5] T. Siebert, K. N. Clasen, M. Ravanbakhsh, and B. Demir, "Multi-modal fusion transformer for visual question answering in remote sensing," in *SPIE Remote Sensing*, pp. 162–170, 2022.
- [6] X. Zheng, B. Wang, X. Du, and X. Lu, "Mutual attention inception network for remote sensing visual question answering," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2021.



- [7] Q. Zhang, Z. Lei, Z. Zhang, and S. Z. Li, "Context-aware attention network for image-text retrieval," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3536–3545, 2020.
- [8] X. Wei, T. Zhang, Y. Li, Y. Zhang, and F. Wu, "Multi-modality cross attention network for image and sentence matching," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10941–10950, 2020.
- [9] J. Wu and R. J. Mooney, "Faithful multimodal explanation for visual question answering," *arXiv preprint arXiv:1809.02805*, 2018.
- [10] C. Gan, Y. Li, H. Li, C. Sun, and B. Gong, "VQS: Linking segmentations to questions and answers for supervised attention in VQA and question-focused semantic segmentation," in *IEEE/CVF ICCV*, 2017.
- [11] S. Lobry, J. Murray, D. Marcos, and D. Tuia, "Visual question answering from remote sensing images," in *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pp. 4951–4954, IEEE, 2019.
- [12] C. Chappuis, E. Walt, V. Mendez, S. Lobry, B. L. Saux, and D. Tuia, "The curse of language biases in remote sensing VQA: the role of spatial attributes, language diversity, and the need for clear evaluation," *arXiv preprint arXiv:2311.16782*, 2023.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE/CVF CVPR*, pp. 770–778, 2016.
- [14] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019.
- [15] Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler, "Aligning books and movies: Towards story-like visual explanations by watching movies and reading books," in *IEEE/CVF ICCV*, 2015.
- [16] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin Transformer: Hierarchical vision transformer using shifted windows," in *IEEE/CVF ICCV*, 2021.
- [17] B. Zhou, H. Zhao, X. Puig, T. Xiao, S. Fidler, A. Barriuso, and A. Torralba, "Semantic understanding of scenes through the ADE20K dataset," *International Journal of Computer Vision*, vol. 127, pp. 302–321, 2019.
- [18] W. Cui, F. Wang, X. He, D. Zhang, X. Xu, M. Yao, Z. Wang, and J. Huang, "Multi-scale semantic segmentation and spatial relationship recognition of remote sensing images based on an attention model," *Remote Sensing*, vol. 11, no. 9, p. 1044, 2019.
- [19] M. Faure, S. Lobry, C. Kurtz, and L. Wendling, "Embedding spatial relations in visual question answering for remote sensing," in *2022 26th International Conference on Pattern Recognition (ICPR)*, pp. 310–316, IEEE, 2022.