



HAL
open science

Weakly Supervised Marine Animal Detection from Remote Sensing Images Using Vector-Quantized Variational Autoencoder

Minh-Tan Pham, Hugo Gangloff, Sébastien Lefèvre

► **To cite this version:**

Minh-Tan Pham, Hugo Gangloff, Sébastien Lefèvre. Weakly Supervised Marine Animal Detection from Remote Sensing Images Using Vector-Quantized Variational Autoencoder. IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium, Jul 2023, Pasadena, United States. pp.5559-5562, 10.1109/IGARSS52108.2023.10282672 . hal-04253903

HAL Id: hal-04253903

<https://hal.science/hal-04253903v1>

Submitted on 23 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WEAKLY SUPERVISED MARINE ANIMAL DETECTION FROM REMOTE SENSING IMAGES USING VECTOR-QUANTIZED VARIATIONAL AUTOENCODER

Minh-Tan Pham^{*}, Hugo Gangloff^{*,†}, Sébastien Lefèvre^{*}

^{*} Université Bretagne Sud, UMR IRISA 6074, 56000 Vannes, France

[†] INRAE, AgroParisTech, UMR MIA Paris-Saclay, 91120 Palaiseau, France

minh-tan.pham@irisa.fr, hugo.gangloff@inrae.fr, sebastien.lefevre@irisa.fr

ABSTRACT

This paper studies a reconstruction-based approach for weakly-supervised animal detection from aerial images in marine environments. Such an approach leverages an anomaly detection framework that computes metrics directly on the input space, enhancing interpretability and anomaly localization compared to feature embedding methods. Building upon the success of Vector-Quantized Variational Autoencoders in anomaly detection on computer vision datasets, we adapt them to the marine animal detection domain and address the challenge of handling noisy data. To evaluate our approach, we compare it with existing methods in the context of marine animal detection from aerial image data. Experiments conducted on two dedicated datasets demonstrate the superior performance of the proposed method over recent studies in the literature. Our framework offers improved interpretability and localization of anomalies, providing valuable insights for monitoring marine ecosystems and mitigating the impact of human activities on marine animals.

Index Terms— Anomaly detection, weakly-supervised learning, aerial remote sensing, marine animal monitoring

1. INTRODUCTION

With the increasing exploitation of marine resources nowadays, monitoring human activities in the sea environments becomes crucial since they could have a significant impact on marine animals [1]. To prevent negative influences from these activities to the ecosystem, aerial surveys are exploited to capture and detect the presence of marine animals for monitoring their population and behavior. However, the visual analysis of such aerial images and videos is very time-consuming and often requires expert knowledge. The emerging development of deep learning techniques such as Convolutional Neural Networks (CNNs) enables us to automatically perform this task. Nevertheless, most of the effective methods using CNNs for

animal detection from aerial images are conducted using supervised learning [2, 3]. To reduce the high cost of data annotation, recent studies have investigated weakly-supervised approaches such as [4, 5]. In these studies, only image-level annotations (which are faster and cheaper to obtain) are used for training instead of object-level labels (bounding boxes) required by fully supervised frameworks. In [5], the authors formulate the weakly-supervised animal detection as an anomaly detection (AD) task where deep networks are first trained only on empty images (i.e., normal data), while the prediction detects animals as anomalies. To do so, they exploit the CNN feature embedding to model the statistical distribution of normal data and compute the anomaly score using distribution-based distance metrics. Compared to previous studies of AD on natural images [6, 7], the adaptations and improvements proposed in [5] make this approach more suitable for remote sensing data, as proved by its higher detection performance.

In this work, we adopt another anomaly detection framework using a reconstruction-based approach to perform weakly-supervised animal detection from aerial images. This framework is appealing given it relies on metrics computed on the input space of images. As such, it could provide easier interpretability and localization of the anomalies than feature embedding-based approach [8]. Indeed, we build upon our recent work on VQ-VAEs (Vector-Quantized Variational Autoencoders) [9] that successfully performed AD on standard computer vision datasets. Here, we investigate its behavior and perform a comparison with respect to literature works in the context of marine animal detection from remote sensing data. We also propose some adaptations of the VQ-VAE framework to remain efficient on noisy data. Experimental results on two dedicated datasets show better performance than feature embedding methods proposed recently in [5–7].

In the rest of this paper, Sec. 2 overviews the the VQ-VAE framework proposed in [9] for AD and provides our improvements to deal with noise in images using the hysteresis thresholding technique [10]. Sec. 3 describes our experiments conducted on two aerial image datasets used in [5] in the context of weakly supervised marine animal detection. We draw some conclusions and potential future works in Sec. 4.

This work was done in the context of the SEMMACAPE project, which benefits from an ADEME (*Agence de la transition écologique*) grant under the “Sustainable Energies” call for research projects (2018–2019).

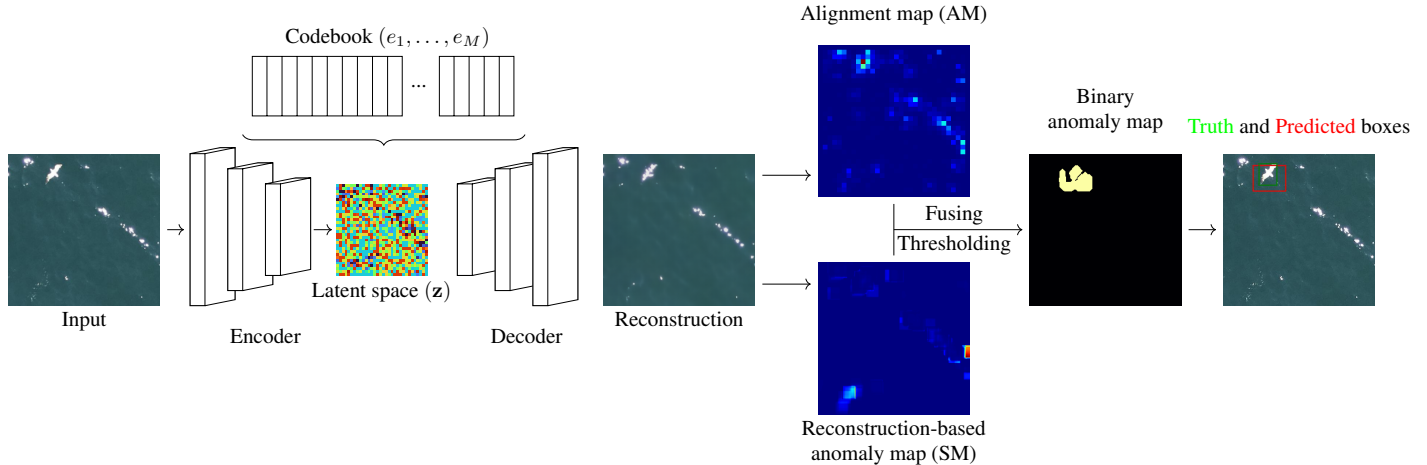


Fig. 1: Overview of the proposed AD framework using VQ-VAE. To deal with noisy data in the context of marine animal detection, both AM and SM are thresholded before being fused to produce the final anomaly map (see Algorithm 1).

2. METHODOLOGY

2.1. VQ-VAEs for anomaly detection

Reconstruction-based AD methods compute anomaly score based on the distance between the input image and its reconstructed version given by a deep network. VAEs are the most popular models that have been adopted within this context. First introduced in [11], VQ-VAEs with discrete latent variables have been recently shown to provide more accurate reconstructions than standard VAEs. This makes VQ-VAEs promising models for reconstruction-based AD, since they could result in less noisy residual images. This higher performance has been proved in [9] on different vision image benchmarks. Fig. 1 illustrates the workflow of VQ-VAEs which is quite similar to classical VAEs, apart from their discrete latent space linked to a particular codebook vector. To compute the final anomaly map, [9] proposed to fuse two components including the reconstruction-based anomaly map (denoted by SM) computed using the Structure Similarity Index Measure (SSIM), and the alignment map (AM) computed from the latent space of the VQ-VAE. AM is then upsampled to the image scale based on morphological dilation operation and fused with the SM to produce the final anomaly map. For more details about the computations and fusion of these two metrics, we refer readers to [9]. To this end, the principle of measuring anomaly score based on both reconstruction and discrete latent space of the model makes the VQ-VAE different and more effective than standard VAE-based AD models.

2.2. Adaptations to deal with noisy data

Similar to most applications of deep frameworks developed on natural images into the remote sensing field, the use of VQ-VAEs to detect marine animals from aerial images is not

Algorithm 1 Using the VQ-VAE for AD in noisy data.

Input:

- SM : the SSIM anomaly map
- AM : the alignment anomaly map
- $\lambda_{SM}, \lambda_{AM}$: thresholds for SM and AM segmentation

Output: $Amap$: the final binary anomaly map

1. Segment SM at threshold $\lambda_{SM} \rightarrow \overline{SM}$
 2. Segment AM at threshold $\lambda_{AM} \rightarrow \overline{AM}$
 3. Get the connected components of $\overline{SM} \rightarrow \widetilde{SM}$
 4. Keep the connected components of \widetilde{SM} which have non null intersection with $\overline{AM} \rightarrow Amap$
-

direct and requires some adaptations to consider the data particularities. Indeed, aerial images acquired over the sea surface for marine monitoring often face with some issues:

- The sun glare which saturates many images makes it difficult to discriminate from the birds which we want to detect, thus tends to induce many false positives.
- The appearance of underwater animals (dolphin, porpoise, etc.) is very varied, due to their positions from the sea surface. Moreover, their movements also induce foam that makes the image very noisy.

To address these complex scenarios, we propose a dedicated approach using thresholding by hysteresis technique [10] (i.e. double thresholding) that uses the segmented AM as markers to keep connected components of interest from the segmented SM . Such an approach exploits the AM as a way to discard false positive results from the SM . The proposed method is summarized in Algorithm 1.

3. EXPERIMENTS

3.1. Datasets and setup

Our experiments are conducted on two dedicated marine animal image datasets which are called *Semmacape* and *Kelonia*, recently introduced and studied in [5]. We adopt the similar data splits to perform AD for weakly-supervised animal detection: only normal data (images without animals) are used for training. The proposed approach is then compared against three recent methods including PaDiM [6], OrthoAD [7] and PaDiM+NF [5]. To conduct experiment on Semmacape data which are highly noisy, our proposed adaptation in Sec. 2.2 is adopted. For model and training setup, we basically follow the setting of the standard VQ-VAE architecture in [9, 11]. For the three reference methods, we also adopt the similar setting as done in [5].

3.2. Results and analysis

Table 1 shows better behavior of our VQ-VAE model on the two marine animal image datasets compared to three reference methods [5–7]. From the table, a high gain of F1-score is achieved for both datasets (about 12% better than the second best method). It should be noted that since VQ-VAEs could provide highly accurate reconstructions, they reduce the noise in the residual images. Therefore, the number of false detections is reduced, which yields high precision but may affect the recall. We observe this behavior from the first table on the Semmacape data.

Method	F1-score	Recall	Precision
PaDiM [6]	0.383	0.434	0.343
OrthoAD [7]	0.458	0.373	0.594
PaDiM+NF [5]	0.530	0.757	0.408
Ours	0.636	0.497	0.884

(a) Results on the Semmacape dataset.

Method	F1-score	Recall	Precision
PaDiM [6]	0.504	0.443	0.586
OrthoAD [7]	0.571	0.514	0.643
PaDiM+NF [5]	0.568	0.559	0.578
Ours	0.726	0.754	0.701

(b) Results on the Kelonia dataset.

Table 1: Comparative results on two marine animal datasets.

We now provide some qualitative assessment on the use of VQ-VAE on the noisy Semmacape images. Fig. 2 illustrates some favorable cases of good results yielded by the proposed method, while Fig. 3 shows some failed cases when dealing with different complex scenarios. First, it is interesting to notice from the figures that the model seems to reconstruct the anomalies (i.e., the animals) into sun glare or foam. This is an expected behaviour since the VQ-VAE has been trained

without being presented any animal: it is thus unable to correctly reconstruct the latter. From Fig. 2, apart from the simplest cases where the animal is isolated on an almost uniform background (as in the first row), the proposed approach also works well on usual cases where foam and sun glare appear and disturb the images. In Fig. 3, the first image is devoid of sun glare or foam but it depicts the case where dolphins are present at different distances from the sea surface. Our approach only detects 2 out of 5 dolphins. Then, the second row illustrates the most complex scenario where images are corrupted by a lot of sun glare. In this case, the *SM* map exhibits many false positives while the *AM* is unable to correctly localize the anomalies. Thus, we end up with no detection at all after the double thresholding step.

4. CONCLUSIONS

The context of AD enables us to take advantage of the huge amount of empty data which is devoid of interest in many traditional approaches. We have reported promising results in weakly-supervised detection of marine animals from aerial images by adapting the VQ-VAE model to the particularly noisy data. The VQ-VAE model learns a representation of the normality (empty images with no animal) during the training step. Then at prediction time, anomalies are spotted thanks to dedicated metrics and a double thresholding approach.

One of the next steps in this unsupervised context using VQ-VAE architectures is to perform the counting and classification of the anomalies directly from their discrete latent space with codebooks, without their explicit detection. This could lead to methodological and practical novelties that could be of crucial importance for the community. Future works may also focus on the use of self-supervised approach [12, 13] to foster the representative capacity of VQ-VAE latent space.

5. REFERENCES

- [1] L. Bergström, L. Kautsky, T. Malm, R. Rosenberg, M. Wahlberg, N. Å. Capetillo, and D. Wilhelmsson, “Effects of offshore wind farms on marine wildlife—a generalized impact assessment,” *Environmental Research Letters*, vol. 9, no. 3, p. 034012, 2014.
- [2] F. Maire, L. M. Alvarez, and A. Hodgson, “Automating marine mammal detection in aerial images captured during wildlife surveys: a deep learning approach,” in *Australasian Joint Conference on Artificial Intelligence*, pp. 379–385, Springer, 2015.
- [3] L. B. Boudaoud, F. Maussang, R. Garello, and A. Chevallier, “Marine bird detection based on deep learning using high-resolution aerial images,” in *IEEE OCEANS*, 2019.




Fig. 2: Some favorable cases processed by the VQ-VAE model from the Semmacape dataset.




Fig. 3: Some complex cases processed by the VQ-VAE model from the Semmacape dataset.

- [4] C. Fasana, S. Pasini, F. Milani, and P. Fraternali, “Weakly supervised object detection for remote sensing images: A survey,” *Remote Sensing*, vol. 14, no. 21, p. 5362, 2022.
- [5] P. Berg, D. Santana Maia, M.-T. Pham, and S. Lefèvre, “Weakly supervised detection of marine animals in high resolution aerial images,” *Remote Sensing*, vol. 14, no. 2, p. 339, 2022.
- [6] T. Defard, A. Setkov, A. Loesch, and R. Audigier, “Padim: a patch distribution modeling framework for anomaly detection and localization,” in *ICPR*, pp. 475–489, Springer, 2021.
- [7] J.-H. Kim, D.-H. Kim, S. Yi, and T. Lee, “Semi-orthogonal embedding for efficient unsupervised anomaly segmentation,” *arXiv preprint arXiv:2105.14737*, 2021.
- [8] L. Ruff, R. A. Vandermeulen, N. Görnitz, A. Binder, E. Müller, K.-R. Müller, and M. Kloft, “Deep semi-supervised anomaly detection,” *arXiv preprint arXiv:1906.02694*, 2019.
- [9] H. Gangloff, M.-T. Pham, L. Courtrai, and S. Lefèvre, “Leveraging vector-quantized variational autoencoder inner metrics for anomaly detection,” in *ICPR*, 2022.
- [10] J. Angulo, T. Nguyen-Khoa, Z. A. Massy, T. Drüeke, and J. Serra, “Morphological quantification of aortic calcification from low magnification images,” *Image Analysis & Stereology*, vol. 22, no. 2, pp. 81–89, 2003.
- [11] A. Van Den Oord, O. Vinyals, *et al.*, “Neural discrete representation learning,” in *NeurIPS*, vol. 30, 2017.
- [12] P. Berg, M.-T. Pham, and N. Courty, “Self-supervised learning for scene classification in remote sensing: Current state of the art and perspectives,” *Remote Sensing*, vol. 14, no. 16, p. 3995, 2022.
- [13] A. Bauer, “Self-supervised training with autoencoders for visual anomaly detection,” *arXiv preprint arXiv:2206.11723*, 2022.