



HAL
open science

Variational Autoencoder with Gaussian Random Field prior: Application to unsupervised animal detection in aerial images

Hugo Gangloff, Minh-Tan Pham, Luc Courtrai, Sébastien Lefèvre

► To cite this version:

Hugo Gangloff, Minh-Tan Pham, Luc Courtrai, Sébastien Lefèvre. Variational Autoencoder with Gaussian Random Field prior: Application to unsupervised animal detection in aerial images. ISPRS Journal of Photogrammetry and Remote Sensing, 2024, 218, pp.600-609. 10.1016/j.isprsjprs.2024.09.028 . hal-04719231

HAL Id: hal-04719231

<https://hal.science/hal-04719231v1>

Submitted on 3 Oct 2024

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.

Variational Autoencoder with Gaussian Random Field prior: application to unsupervised animal detection in aerial images

Hugo Gangloff¹, Minh-Tan Pham², Luc Courtrai², Sébastien Lefèvre²

¹*Université Paris-Saclay, INRAE, AgroParisTech, UMR MIA Paris-Saclay, 91120, Palaiseau, France*

²*IRISA, Université Bretagne Sud, UMR 6074, 56000 Vannes, France*

Abstract

In real world datasets of aerial images, the objects of interest are often missing, hard to annotate and of varying aspects. The framework of unsupervised Anomaly Detection (AD) is highly relevant in this context, and Variational Autoencoders (VAEs), a family of popular probabilistic models, are often used. We develop on the literature of VAEs for AD in order to take advantage of the particular textures that appear in natural aerial images. More precisely we propose a new VAE model with a Gaussian Random Field (GRF) prior (VAE-GRF), which generalizes the classical VAE model, and we provide the necessary procedures and hypotheses required for the model to be tractable. We show that, under some assumptions, the VAE-GRF largely outperforms the traditional VAE and some other probabilistic models developed for AD. Our results suggest that the VAE-GRF could be used as a relevant VAE baseline in place of the traditional VAE with very limited additional computational cost. We provide competitive results on the MVTec reference dataset for visual inspection, and two other datasets dedicated to the task of unsupervised animal detection

*Corresponding author.

Email addresses: hugo.gangloff@inrae.fr (Hugo Gangloff¹), minh-tan.pham@univ-ubs.fr (Minh-Tan Pham²), luc.courtrai@univ-ubs.fr (Luc Courtrai²), sebastien.lefevre@univ-ubs.fr (Sébastien Lefèvre²)

in aerial images.

Keywords: Variational autoencoders, anomaly detection, Gaussian random fields, aerial images

1. Introduction

This article introduces a new deep probabilistic model followed by its application to real world data. First, a new model of Variational Autoencoders (VAEs) with a Gaussian Random Field (GRF) prior is presented. This offers a relevant way to model images with strong spatial correlations. Second, the VAE-GRF is used in the context of Anomaly Detection (AD). More precisely, we address the real world application of unsupervised detection of animals in aerial images. Both of these topics are now presented.

1.1. Variational Autoencoders and some extensions

Our work focuses on VAEs. They are generative probabilistic models, widely used in the context of anomaly detection. They are popular for several reasons. First, they are derived from a sound probabilistic background and robust training procedures have been developed for such models. VAEs are also widely used in unsupervised settings such as ours. Finally, since a VAE is a generative model, samples can be generated from the model for any kind of purpose once the model has been trained.

In this article, we introduce the VAE-GRF model for images which makes use of a prior in the form of a stationary bi-dimensional GRF on the torus. Indeed, we construct a convolutional latent space in which we assume that the model learns a compressed representation of the input images. Therefore, such an approach offers a refinement of the prior distribution over the latent random variables as compared to the independent and identically distributed standardized Gaussian prior from the traditional VAE context. We study the

advantages of the VAE-GRF in terms of modeling. We also demonstrate how the stationary and torus assumptions are used to develop a model with efficient computations despite the full covariance structure of the prior.

Over the last few years, several works have considered more complex priors for VAEs to be able to propose a more relevant modeling of the data. A recent review is available in [11]. Following this review, our work fall in the category called *distributional VAEs*, in which we consider *plugging* another (prior) distribution in the VAE model, without modifying the model architecture. Three main works falls into this category. Hyperspherical VAEs [7] along with VAEs with Gaussian mixture prior [10] are perhaps the most popular works which refine the prior law but they do not address the question of latent variable correlations. The closest works to ours are perhaps VAEs with Gaussian Process priors, as explored in [6], [30] or in [19]. Those works propose a Gaussian Process to model inter-sample correlations, which, by design induces some correlations on the latent random variables too. However the latent spaces they use are one dimensional or factorized as Markov chains which differs from the convolutional latent space we propose. Indeed, they aim at encoding inter-sample correlations between hidden random variables while we aim at encoding the spatial correlations in the latent space, for each sample. Therefore, following the previously mentioned review and up to the best of our knowledge, our article is the first study that addresses the question of spatial correlation within latent variables of VAEs for image processing. We also propose an efficient way to implement such a model.

Therefore, as opposed to these approaches, we use a bi-dimensional GRF, which is more costly and requires different hypotheses to maintain a tractable model (stationarity and torus). In addition, since AD is not the focus of these approaches and since it remains unclear how Gaussian Process VAEs could be

adapted for real-world image processing, we will discard them from comparisons.

1.2. Anomaly Detection in aerial images

In this article, we address the problem of unsupervised detection of animals in aerial images. Detecting, tracking and counting animals are real-world applications that are more and more studied in the literature, especially thanks to the unprecedented availability of Unmanned Aerial Vehicles (UAV) [1] [5] [26]. However, to the best of our knowledge, very few works consider the unsupervised context, which does not require the tedious and costly annotation step [3]. As the number of data keeps increasing, it becomes a necessity to develop reliable unsupervised approaches which can solve these tasks.

In the unsupervised context, we can not use any annotation to learn a representation of the animals. Moreover, in the context of animal detection with UAV, most of the captured images are empty, *i.e.*, they do not contain any animal. For these reasons, in this article, we resort to the principles of unsupervised AD in order to detect animals. AD is a vast research field and several reviews have already been published on the topic [34] [27]. An introduction to AD is proposed in Sec. 3.1. In the context of this article, the AD approaches we will compare the VAE-GRF to are presented in Sec. 3.3.

1.3. Outline of the article

The outline of the article is the following. We first present the new VAE-GRF model along with some theoretical backgrounds motivating the work. We then briefly review unsupervised AD and assess the new VAE-GRF on the AD task on a standard public dataset (dedicated to visual inspection). Finally, we use the AD context to perform unsupervised animal detection in aerial images from two datasets. In all the experiments, the new model is compared with other state-of-the-art approaches.

Remark: The code of the VAE-GRF model to reproduce the experiments presented in this article is available at https://github.com/HGangloff/vae_grf.

2. Variational Autoencoders with Gaussian Random Field prior

2.1. Gaussian Random Fields

2.1.1. Definition

The probability density function of a GRF [33] with mean $\boldsymbol{\mu}$ and covariance matrix Σ (forming a set of parameters $\{\boldsymbol{\mu}, \Sigma\}$) is given by

$$p(\mathbf{x}) = (2\pi)^{-\frac{n}{2}} |\Sigma|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})\right), \quad (1)$$

where $|\cdot|$ refers to the determinant of a matrix and n is the dimension of \mathbf{x} . The GRF will be associated to the regular graph defined by an image: a node is associated to each pixel and nodes linked with each other in the graph are adjacent pixels (we consider the 8-nearest neighbors for each pixel). Then $\mathbf{x} = \{x\}_{s \in \mathcal{S}}$ where \mathcal{S} is the set of nodes of the graph.

In the rest of the paper, we will consider GRFs with stationary mean and stationary covariance matrix, *i.e.*, the mean is the same for each x_s and the covariance between two x_s only depends on their distance (distance that will be defined later).

2.1.2. Spectral properties

In this article, along with the stationarity assumption, we will formulate the *torus assumption* on our images. In the torus assumption, the image borders are supposed wrapped like on a torus. One can observe that a stationary GRF defined on an image with torus assumption will have a covariance matrix which is block-circulant with circulant blocks [33]. Therefore, such a covariance matrix does not need to be fully stored, it is entirely defined by a smaller matrix called the *base* matrix. Indeed, for a matrix of dimension $l_x \times l_y$, its covariance matrix

C has dimension $l_x l_y \times l_x l_y$. However, if both the stationarity and torus assumptions are made, the covariance matrix is block-circulant with circulant blocks and is entirely defined by its *base* matrix, $\text{base}(C)$, with dimension $l_x \times l_y$. All details about circulant and block-circulant matrices can be found in [33]. Matrix operations for block-circulant matrices with circulant blocks are efficiently computed with the Fourier transforms and are called the *spectral properties*¹. The formulas we use in the article are given in App. Appendix A.

2.2. VAE-GRF: model definition

2.2.1. Generative model architecture

Recall that VAEs are generative probabilistic models which aim at modeling a distribution over the observations $p_{\theta}(\mathbf{x})$ with the help of latent random variables \mathbf{z} , such that, $p_{\theta}(\mathbf{x}) = \int p_{\theta}(\mathbf{x}, \mathbf{z}) d\mathbf{z} = \int p_{\theta}(\mathbf{x}|\mathbf{z}) p_{\theta}(\mathbf{z}) d\mathbf{z}$. However, in VAEs, computing the posterior $p_{\theta}(\mathbf{z}|\mathbf{x}) = p_{\theta}(\mathbf{x}, \mathbf{z})/p_{\theta}(\mathbf{x}) = p_{\theta}(\mathbf{x}, \mathbf{z})/\int p_{\theta}(\mathbf{x}|\mathbf{z}) p_{\theta}(\mathbf{z}) d\mathbf{z}$ is intractable (the model likelihood, at the denominator, is, in general, an intractable integral). In VAEs, we thus introduce a variational distribution $q_{\varphi}(\mathbf{z}|\mathbf{x})$ that aims at approximating $p_{\theta}(\mathbf{z}|\mathbf{x})$ and whose parameters φ are learnt during the optimization process.

The VAE-GRF model we propose is first composed of a stochastic encoding network, with input \mathbf{x} , which maps to a convolutional latent space associated to the realizations of a random variable \mathbf{z} , following the ideas from [35] [13]. Thus, \mathbf{z} has dimension $N = n_x \times n_y \times n_z$ (width×height×depth). The outputs of the encoder, $L = \text{diag}(\sigma_1^2, \dots, \sigma_N^2)$ and $\mathbf{m} \in \mathbb{R}^N$, parametrize a variational posterior distribution, $q_{\varphi}(\mathbf{z}|\mathbf{x})$, chosen as independent Gaussian random variables; we then have $q_{\varphi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \mathbf{m}, L)$. For a reason that will be clarified later, we factorize on the depth dimension such that $q_{\varphi}(\mathbf{z}|\mathbf{x}) = \prod_{k=1}^{n_z} \mathcal{N}(\mathbf{z}_k; \mathbf{m}_k, L_k)$, and then $L_k =$

¹For example a one-dimensional circular convolution of a M -length signal has computational complexity $\mathcal{O}(M^2)$. This reduces to $\mathcal{O}(M \log M)$ thanks to the Fast Fourier Transform algorithm [12].

$\text{diag}((\sigma_k^2)_1, \dots, (\sigma_k^2)_{n_x n_y})$. The model is then composed of a stochastic decoder network whose output $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_{l_x l_y})$ parametrizes a product of independent Continuous Bernoulli random variables [25], $p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) = \prod_{k=1}^{l_x l_y} \mathcal{CB}(x_k, \lambda_k)$. A realization of this stochastic decoder corresponds to a reconstruction, denoted $\hat{\mathbf{x}}$, of the input image \mathbf{x} by the model.

2.2.2. GRF prior

Let us first be more specific about the structure of the latent space with GRF prior with parameters $\bar{\boldsymbol{\theta}}$. We consider the zero-mean stationary and toroidal GRF prior on the $n_x \times n_y$ dimension² and we consider the components of \mathbf{z} to remain independent on the depth dimension n_z . Hence $p_{\bar{\boldsymbol{\theta}}}(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \Sigma) = \prod_{k=1}^{n_z} \mathcal{N}(\mathbf{z}_k; \mathbf{0}, \Sigma_k)$, where Σ_k are Symmetric Positive Definite (SPD) matrices, and the computations can still be done in a parallel manner on this dimension because the computations fall back to n_z parallel computations involving bi-dimensional GRFs which shared parameters. We go one step further by sharing the parameters between these GRFs, thus, $p_{\bar{\boldsymbol{\theta}}}(\mathbf{z}) = \prod_{k=1}^{n_z} \mathcal{N}(\mathbf{z}_k; \mathbf{0}, \Sigma)$. Note that such a parameter sharing is also proposed in the one dimensional case of the factorized Gaussian Process VAEs of [19]. Note also that, since the diagonal covariance matrix used in the standardized Gaussian prior in the classical VAE model is comprised in the set of the SPD matrices yielding a GRF, the VAE-GRF model is a strict generalization of the classical VAE model.

We emphasize the fact that, even when it is not mentioned, we never fully compute the full covariance matrix Σ but only its base matrix. Then, note that the estimation of the covariance matrix Σ needs to yield a SPD matrix. In our work, to ensure this requirement, the covariance matrix Σ is assumed to be generated by Matern correlation functions [23] which form a class of correlation

²In all the following we consider, without loss of generality, that $l_x = l_y$ and $n_x = n_y$: square input images and square latent space images.

functions defined in \mathbb{R}^2 by the general equation: $\forall a \in \mathbb{R}^2, \forall b \in \mathbb{R}^2$,

$$\rho_\nu(a, b; r) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|a - b\|_t}{r} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\|a - b\|_t}{r} \right), \quad (2)$$

where Γ is the gamma function, K_ν is the modified Bessel function of the second kind, ν is a non-negative parameter, r is the non-negative range parameter and $\|\cdot\|_t$ is the Euclidean distance on an image (of dimensions $l_x \times l_y$) with torus assumption. For $a = (a_1, a_2) \in \mathbb{R}^2, b = (b_1, b_2) \in \mathbb{R}^2$, the Euclidean distance on the torus is defined by

$$\|a - b\|_t = \left(\min(|a_1 - b_1|, l_x - |a_1 - b_1|)^2 + \min(|a_2 - b_2|, l_y - |a_2 - b_2|)^2 \right)^{\frac{1}{2}}. \quad (3)$$

From Eq. (2) several standard correlation functions can be derived:

- the exponential or Matern- $\frac{1}{2}$ (Eq. (2) with $\nu = \frac{1}{2}$) correlation function on the torus which reads

$$\rho^{1/2}(a, b; r) = \exp \left(- \frac{\|a - b\|_t}{r} \right). \quad (4)$$

- the Matern- $\frac{3}{2}$ (Eq. (2) with $\nu = \frac{3}{2}$) correlation function on the torus which reads

$$\rho^{3/2}(a, b; r) = \left(\frac{\|a - b\|_t}{r} + 1 \right) \exp \left(- \frac{\|a - b\|_t}{r} \right). \quad (5)$$

- the Gaussian or Matern- ∞ (Eq. (2) with $\nu \rightarrow \infty$) correlation function on the torus which reads

$$\rho^\infty(a, b; r) = \exp \left(- \left(\frac{\|a - b\|_t}{r} \right)^2 \right). \quad (6)$$

These instances of the Matern correlation function are commonly used in spatial statistics [14] [23]. They differ in the way they induce a decrease with the distance in the correlation between the random variables. Fig. 1 illustrates the different correlation functions that will be experimentally compared in Ap-

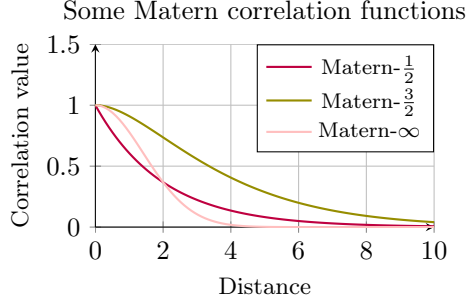


Figure 1: Illustration of some specific instances of Matern correlation function. We have $r = 2$ for all curves.

pendix Appendix B.1.

Note that to form the covariance matrix, it remains to multiply the correlation function with the GRF variance σ^2 , leading to the following covariance between two random variables $(z_k)_s$ and $(z_k)_{s'}$, $\forall k \in \{1, \dots, n_z\}$:

$$\text{Cov}((z_k)_s, (z_k)_{s'}) = (\Sigma)_{s,s'} = \sigma^2 \rho(s, s'; r). \quad (7)$$

We also note that the variance parameter σ^2 and the range r are constants for the whole GRF to respect a stationary covariance structure (and for all the n_z GRFs). Since we assume that the mean parameter $\boldsymbol{\mu} = \mathbf{0}$, the set of parameters defining the GRF prior is $\bar{\boldsymbol{\theta}} = \{r, \sigma^2\}$. In our approach, the prior parameters will be learnt as additional parameters for the network by log-likelihood maximization of the observed image \mathbf{x} over which we also assume the same GRF structure as prior. This latter GRF has parameter $\{\tilde{r}, \tilde{\sigma}^2\}$. Using the estimation of $\bar{\boldsymbol{\theta}}$, we get the GRF over z parameter with the assumed relations:

$$\begin{cases} r_k = \lambda \tilde{r}, \\ \sigma^2 = \tilde{\sigma}^2, \end{cases} \quad (8)$$

where λ is set to the ratio of the latent image size over the input image size: $\lambda = \lfloor \frac{n_x}{l_x} \rfloor = \lfloor \frac{n_y}{l_y} \rfloor$. The details about these computations are given in Sec. 2.3.2.

Remark: Importantly, in the VAE-GRF, the encoder, the latent space and the decoder of the model have exactly the same number of parameters as the classical VAE and, overall, the VAE-GRF does not make use of any additional module. We only need to store two more scalar parameters in the VAE-GRF model: the range and the variance of the prior. Indeed, we propose a refinement in the prior modeling, yielding improved results, thus with a very limited additional computational cost.

2.3. Training the model

Classically, VAEs are trained by maximizing a lower-bound on the log-likelihood, called the Evidential Lower Bound (ELBO), denoted $\mathcal{E}_{\boldsymbol{\theta}, \boldsymbol{\varphi}}(\mathbf{x})$, which reads

$$\log p_{\boldsymbol{\theta}}(\mathbf{x}) \geq \underbrace{\mathbb{E}_{q_{\boldsymbol{\varphi}}(\mathbf{z}|\mathbf{x})}[\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})] - \text{KL}(q_{\boldsymbol{\varphi}}(\mathbf{z}|\mathbf{x})||p_{\bar{\boldsymbol{\theta}}}(\mathbf{z}))}_{\mathcal{E}_{\boldsymbol{\theta}, \boldsymbol{\varphi}}(\mathbf{x})}, \quad (9)$$

where KL denotes the Kullback-Leibler divergence between two probability distributions.

For more flexibility in the training, we use the β -ELBO [18] (with an additional stop gradient operator) which reads

$$\mathcal{E}_{\boldsymbol{\theta}, \boldsymbol{\varphi}, \beta}(\mathbf{x}) = \mathbb{E}_{q_{\boldsymbol{\varphi}}(\mathbf{z}|\mathbf{x})}[\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})] - \beta \text{KL}(q_{\boldsymbol{\varphi}}(\mathbf{z}|\mathbf{x})||p_{\bar{\boldsymbol{\theta}}}(\mathbf{z})). \quad (10)$$

The reason for introducing the β scalar is that we can modulate how strong we want the posterior to fit to the prior, in other words, how strong the GRF assumption is.

Then, the VAE-GRF is trained by alternatively optimizing the prior parameters $\bar{\boldsymbol{\theta}}$ and the VAE-GRF parameters $(\boldsymbol{\theta}, \boldsymbol{\varphi})$ by maximizing, respectively, the

log-likelihood and the ELBO:

$$\begin{cases} \max_{\bar{\theta}} \log p_{\bar{\theta}}(\mathbf{x}), \\ \max_{\theta, \varphi} \mathcal{E}_{\theta, \varphi, \beta}(\mathbf{x}). \end{cases} \quad (11)$$

From the practical viewpoint we use the Adam optimizer [20] to update the VAE-GRF parameters. On the other hand, we use the Rprop optimizer [32] to update the two prior parameters, thus benefiting from an adaptative learning rate.

2.3.1. The ELBO term

The ELBO is composed of a first term similar to a cross-entropy which favors reconstructions similar to the input, it is called the *reconstruction term* and its computation is the same as classically done in VAEs. The Kullback-Leibler divergence term³ can be interpreted as a regularization term which pushes the posterior to match the prior during the training. We have by definition

$$\text{KL}(q_{\varphi}(\mathbf{z}|\mathbf{x})||p_{\bar{\theta}}(\mathbf{z})) = \mathbb{E}_{q_{\varphi}(\mathbf{z}|\mathbf{x})} \left[\log \frac{q_{\varphi}(\mathbf{z}|\mathbf{x})}{p_{\bar{\theta}}(\mathbf{z})} \right], \quad (12)$$

and we can show that

$$\begin{aligned} \text{KL}(q_{\varphi}(\mathbf{z}|\mathbf{x})||p_{\bar{\theta}}(\mathbf{z})) &= \text{KL}(\mathcal{N}(\mathbf{z}; \mathbf{m}, L)||\mathcal{N}(\mathbf{z}; \mathbf{0}, \Sigma)), \\ &= \frac{1}{2} \sum_{k=1}^{n_k} \left(-\log|L_k| - n_x n_y + \log|\Sigma| \right. \\ &\quad \left. + \text{tr}(\Sigma^{-1} \mathbf{m}_k \mathbf{m}_k^T) + \text{tr}(\Sigma^{-1} L_k) \right), \end{aligned} \quad (13)$$

where tr refers to the trace operator. In Sec. 2.4, we explain how the computations of Eq. (13) can remain cheap in our VAE-GRF model, thanks to the stationary GRF prior on torus.

³In this section, we ignore the stop gradient operator whose role has been clarified in the previous section.

Remark: In the common case of a VAE with a standardized Gaussian as prior [21], Eq. (13) falls back to

$$\begin{aligned} \text{KL}(q_{\boldsymbol{\varphi}}(\mathbf{z}|\mathbf{x})||p_{\bar{\boldsymbol{\theta}}}(\mathbf{z})) &= \text{KL}(\mathcal{N}(\mathbf{z}; \mathbf{m}, L)||\mathcal{N}(\mathbf{z}; \mathbf{0}, I)), \\ &= \frac{1}{2} \sum_{i=1}^N (-\log \sigma_i^2 - 1 + m_i^2 + \sigma_i^2). \end{aligned} \quad (14)$$

2.3.2. The log-likelihood term

The second term of the loss function is the log-likelihood of the observed image upon which we assume a stationary GRF prior on the torus. With $\bar{\boldsymbol{\theta}}$ defined as before, we have

$$\log p_{\bar{\boldsymbol{\theta}}}(\mathbf{x}) = -\frac{l_x l_y}{2} \log 2\pi - \frac{1}{2} \log |\bar{\Sigma}| - \frac{1}{2} \mathbf{x}^T \bar{\Sigma}^{-1} \mathbf{x}, \quad (15)$$

where $\bar{\Sigma}$ is the covariance matrix generated from the selected correlation function and the parameter $\bar{\boldsymbol{\theta}}$. Again, the spectral properties are usable, and the details for the computation of Eq. (15) are given in Sec. 2.4.

2.4. Computational efficiency

Recall that the stationarity of the GRF and the torus assumption make the properties detailed in Sec. 2.1.2 usable. The latter are a critical element in our approach. We now detail how the terms in the objectives (Eq. (11)) can be efficiently computed both in terms of time and memory complexity. We review each of the terms that are found in Eqs. (13) and (15). A naive computation of these terms would be impossible because they would lead to excessive memory allocations for standard GPUs. Thus, in the VAE-GRF:

- All determinants can be computed directly with Eq. (A.3).
- All matrix inversions can be computed with Eq. (A.1).
- The term $\text{tr}(\Sigma^{-1} \mathbf{m} \mathbf{m}^T)$ can be computed by first computing $\Sigma^{-1} \mathbf{m}$. We

then multiply this result element-by-element with \mathbf{m}^T and sum the resulting vector.

- The term $\text{tr}(\Sigma^{-1}L)$ is equivalent to multiplying each element of the diagonal matrix $L = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$ by the element at position $(0,0)$ of matrix Σ^{-1} and then summing the resulting matrix.

Note that all the operations involve the base matrices of the full covariance matrices and thus, no matrix with greater dimension than $n_x \times n_y$ is stored or allocated. Moreover, an increase in time efficiency is also due to the use of Fourier transforms in the spectral properties. This way, all these operations can be implemented on GPU, differentiated and computed in a batched manner thanks to the Pytorch [29] library.

Example: Let us consider, for example, the naive computation of $\Sigma^{-1}\mathbf{m}_k$ (with a non stationary Σ^{-1}). In such case Σ^{-1} has dimension 1024×1024 and we would need to allocate batches of size $1024 \times 1024 \times 256$. Consider a batch size of 16 and 32-bit float precision, the latent space alone would require $16 \times 1024 \times 1024 \times 256 \times 32\text{bits} \approx 16\text{GBytes}$ without our specific hypothesis. This represents allocations for the computation of the KL term only and thus illustrates that the procedure we develop is crucial for the VAE-GRF to be implemented on standard GPUs.

3. Unsupervised anomaly detection

3.1. Principle

Anomaly Detection (AD) refers to the task of detecting observations that deviate from some underlying concept of normality [16]. It is an popular research topic with a vast literature [31] but recently, it has been revolutionized by deep learning approaches which have yielded new state-of-the-art results thanks to the

unprecedented possibilities of capturing and modeling the normality [34]. The most popular family of approaches to AD is composed of the reconstruction-based methods, in which distances (ℓ -2 distance, SSIM distance [37], etc.) are computed between the inputs and the reconstructions in order to locate the anomalies. The summary works of [2] and [38] illustrate these approaches upon which we base our work.

The principle of unsupervised AD is the following. A representation of the normality is learnt thanks to a deep model which is trained on normal samples devoid of anomalies. Then, at testing time, metrics can be used to detect a change in behaviour of the model when the latter is presented an anomalous sample. VAEs and Vector Quantized (VQ)-VAEs are popular models for this unsupervised task. The former models have appeared first [2] [38] [22] [9] while the latter models have been explored very recently [36] [13].

Remark: Strictly speaking, in the setting of this article, the approach should be called *weakly supervised* since some supervision is needed to make sure the training is done only on images without anomalies. However, in this context, *unsupervised* and *weakly supervised* are found interchangeably in the literature.

3.2. VAE-GRF for Anomaly Detection

We now precisely describe our VAE-GRF model for AD. The model is classically trained on a dataset of normal samples to learn a representation of the normality. Then, we propose to detect the anomalies at testing time from an original metric defined on the latent space coupled with a reconstruction based metric.

We show that the VAE-GRF is a robust approach, competitive with comparable approaches from the state-of-the-art provided that some assumptions

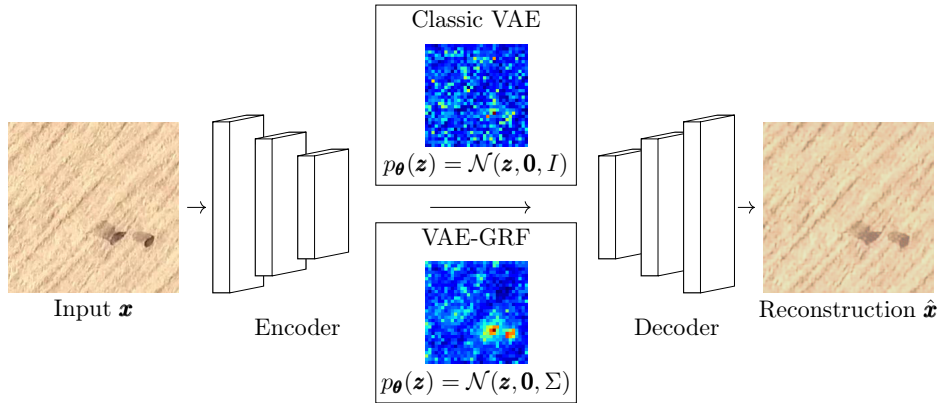


Figure 2: The classic VAE and VAE-GRF architectures. We illustrate the pixel-wise MAD metric (Eq. 16) computed from the latent space in both cases. The interpretability and improved results of the metric can be seen in the case of the VAE-GRF as opposed to VAE with classical standardized Gaussian prior.

hold on the datasets. Indeed, we will show that the new metric defined on the latent space requires, as expected, that the stationary and torus assumptions made for the GRF hold. This is the case for some texture images as we illustrate in Sec. 3.5, and also for some aerial images in remote sensing as we present in Sec. 4.

Let us now introduce our new metric for AD that will be used in all the following experiments. Let MAD (from Mean Absolute Deviation) be the $n_x \times n_y$ anomaly map that we compute from the latent space. Each of the pixel of the anomaly map is computed as the mean absolute deviation from the mean of the same location of the output of the convolutional encoder \mathbf{z} , *i.e.*, $\forall x \in \{1, \dots, n_x\}, \forall y \in \{1, \dots, n_y\}$,

$$MAD_{x,y} = \frac{1}{n_z} \sum_{k=1}^{n_z} \left| z_{x,y,k} - \frac{1}{n_z} \sum_{k=1}^{n_z} z_{x,y,k} \right|. \quad (16)$$

Then, the MAD anomaly map is upsampled to the dimension of the original image and AD can be performed.

Let us also define a reconstruction-based anomaly map, called SM , which

uses the Structural Similarity Index Measure (SSIM) [37]. For each pixel i , we have:

$$SM(x_i) = \text{SSIM}(\mathbf{p}_i, \mathbf{q}_i) = \frac{(2\mu_{\mathbf{p}}\mu_{\mathbf{q}} + c_1)(2\sigma_{\mathbf{p}\mathbf{q}} + c_2)}{(\mu_{\mathbf{p}}^2 + \mu_{\mathbf{q}}^2 + c_1)(\sigma_{\mathbf{p}}^2 + \sigma_{\mathbf{q}}^2 + c_2)}, \quad (17)$$

where \mathbf{p}_i (resp. \mathbf{q}_i) is a patch around pixel i of \mathbf{x} (resp. $\hat{\mathbf{x}}$). μ_p , σ_p and σ_{pq} represent, respectively, the mean, the standard deviation and the covariance of the patches. The scalars are set to $c_1 = 0.01$ and $c_2 = 0.03$ [37].

From the two previous anomaly maps we form a new original anomaly map using an element-wise multiplication, which is denoted $MAD \odot SM$. The two anomaly maps MAD and $MAD \odot SM$ are at the core of our contribution for improved AD, along with the GRF prior. These anomaly maps will then be compared with state-of-the-art models presented in the next section.

For completeness, Fig. 2 graphically illustrates the model and MAD anomaly map. The encoder-decoder structure is, however, identical to that of a traditional VAE. This is remarkable since the improvements we propose are based only on an original refinement of the probabilistic model.

3.3. Related work

To date, most of the best performing methods are based on feature extraction [8] [24]. The latter always perform slightly better than VAE-based approaches. However, generative models offer other possibilities to work with aspects related to the probabilistic framework, sample generation or latent space interpretability [36] [19].

Such aspects are also crucial to our study, thus we will compare the VAE-GRF with generative models from the literature: comparable generative models are based on a baseline VAE and also rely on a modelization refinement. Moreover, we favor related works with no additional modules attached (*e.g.* discriminator modules [35]). In such setting, we are able to fully evaluate the gain offered by the GRF prior and the new AD metrics.

We now list the most notable comparable approaches:

- We compare our MAD and $MAD \odot SM$ metrics to the already existing SM metric introduced in [4].
- An interesting idea quite similar to ours is the Visually Explained Variational Autoencoder (VEVAE) [22]. However the code is not available and we were not able to replicate their study. Hence we compare the VAE-GRF to the VE-VAE on the MVTEC experiment only in Sec. 3.5.
- AD metrics similar to ours have also been introduced in [38]. The magnitude of the gradient of the loss $|\mathcal{E}_{\theta, \varphi}(\mathbf{x})|$ is used as a metric to localize anomalies and it can be multiplied by another anomaly map. Hence we also straightforwardly test the metric $|\mathcal{E}_{\theta, \varphi}(\mathbf{x})| \odot SM$.
- We also test the iterative procedure proposed by [9]. It consists in a refinement of the VAE reconstruction. We limit this approach to 15 iterations of projection because it is very slow and empirically performed best. The refinement is performed before computing the SM anomaly map: we refer to this approach as SM grad.

Thus, except for VEVAE, all the comparable approaches are from our reimplementations because the code were unavailable. All these reimplementations have the same classic VAE baseline described in Sec. 3.4 (except for the VE-VAE).

3.4. Network architecture

All the inputs are resized to the size 256×256 ($l_x = l_y = 256$). Then:

- The encoder consists in the first three layers of the ResNet18 neural network [17], followed by a 1×1 convolutional layer setting the depth dimension to $2 * 256 = 512$ (kernel size 1, stride 1 and padding 0).

- The latent space image has width and height 32 ($n_x = n_y = 32$) and depth $n_z = 256$.
- The decoder is first composed of a deconvolutional layer with input dimension 256 and output 128 (kernel size 1, stride 1, and padding 0). It is then stacked with three deconvolutional layers (kernel size 4, stride 2 and padding 1), each followed by a ReLU activation and Batch Normalization.

This architecture originates from ideas found in [22] or [13].

3.5. Effects of the GRF prior: experiments on MVTEC textures

3.5.1. Presentation

In this section we experimentally investigate the role of the prior in the VAE-GRF model. This experiment is based on the texture images from the MVTEC dataset [4] which is the standard dataset for AD. The dataset provides normal and abnormal (defective) RGB images of industrial goods from 15 different categories. The results reported below compare the VAE-GRF with Matern- $\frac{3}{2}$ correlation function to other models. In Appendix B.1, we compare several choices of correlation functions.

3.5.2. Results

Tab. 1 and 2 summarize the results in terms of pixel-wise area under the receiver operating characteristic curve (ROCAUC), computed with the final anomaly map and the ground truth. This is the classical score used for AD. Compared to the traditional VAE model (Tab. 1), on the images which respect the most the stationary GRF assumption, *i.e.* texture images, the VAE-GRF performs better, especially when it comes to the new metrics we introduced (see the *MAD* or *MAD* \odot *SM* metric on *Leather*, *Tile* and *Wood*). On images where the hypothesis of a stationary GRF is the most erroneous (see *Carpet* or *Hazelnut*), the *MAD* metric does not increase. We then can see the interest of the refinement in the VAE modeling and the potential gain in the *MAD*

Category		VAE			VAE-GRF		
		SM	MAD	$MAD \odot SM$	SM	MAD	$MAD \odot SM$
Stationary textures	Leather ($\beta = 100$)	<u>0.97</u>	0.92	0.94	0.95	<u>0.95</u>	0.98
	Tile ($\beta = 0.01$)	<u>0.93</u>	0.69	<u>0.93</u>	<u>0.93</u>	<u>0.87</u>	0.95
	Wood ($\beta = 100$)	0.70	0.75	0.76	0.78	0.86	<u>0.80</u>
Non stationary textures	Carpet ($\beta = 0.01$)	<u>0.91</u>	0.60	0.92	0.90	<u>0.61</u>	<u>0.91</u>
	Grid ($\beta = 0.01$)	<u>0.96</u>	0.66	<u>0.96</u>	0.95	<u>0.76</u>	0.97
	Hazelnut ($\beta = 1$)	0.94	0.94	0.96	0.98	<u>0.91</u>	0.98

Table 1: ROCAUC scores for pixel-wise AD on the texture images from the MVTEC dataset. Best scores appear in bold and second best are underlined. Scores from the main contribution of this article are in purple.

metric, as well as its limitations when the stationary GRF prior assumption is violated. That is why we only focus our analysis on MVTEC texture images. Similar conclusions can be made when we compare to the other approaches in the literature (Tab. 2). The VAE-GRF is always on par or better than the state-of-the-art results. The best gains coming from the texture images.

Note also that, globally, both our VAE and GRF-VAE are competitive against state-of-the-art approaches which also suggests that the proposed architecture and training procedures are highly relevant and optimized (see Sec. 3.4). Finally, Fig. 3 illustrates the experiment on some images of the MVTEC database.

Remark: The choice of the correlation function as well as the choice of the hyperparameter β is empirically made in our study through gridsearch procedures that we present in Appendix B.1 and Appendix B.2.

4. Unsupervised animal detection in aerial images with VAEs-GRF

In this section we present how the VAE-GRF model can be used to solve the real world problem of detection in aerial images, where the landscape can be assumed to be well modeled by a stationary texture. We focus on the task of unsupervised animal detection on two different datasets. The context of

	Category	VAE-GRF	Liu et al. [22]	Zimmerer et al. [38]		Dehaene et al. [9]
		$MAD \odot SM$	VEVAE	$ \nabla_{\mathbf{x}} \mathcal{E}_{\theta, \varphi}(\mathbf{x}) $	$ \nabla_{\mathbf{x}} \mathcal{E}_{\theta, \varphi}(\mathbf{x}) \odot SM$	SM grad
Stationary textures	Leather	0.98	<u>0.95</u>	0.55	0.94	<u>0.95</u>
	Tile	0.95	<u>0.80</u>	0.62	0.78	0.78
	Wood	0.80	<u>0.77</u>	0.54	0.76	0.74
Non stationary textures	Carpet	0.91	0.78	0.61	<u>0.89</u>	<u>0.89</u>
	Grid	0.97	0.73	0.54	<u>0.92</u>	0.90
	Hazelnut	0.98	0.98	0.80	0.98	<u>0.92</u>

Table 2: ROCAUC scores for pixel-wise AD on the texture images from the MVTEc dataset. Comparisons with other methods. Best scores appear in bold and second best are underlined. Scores from the main contribution of this article are in purple.

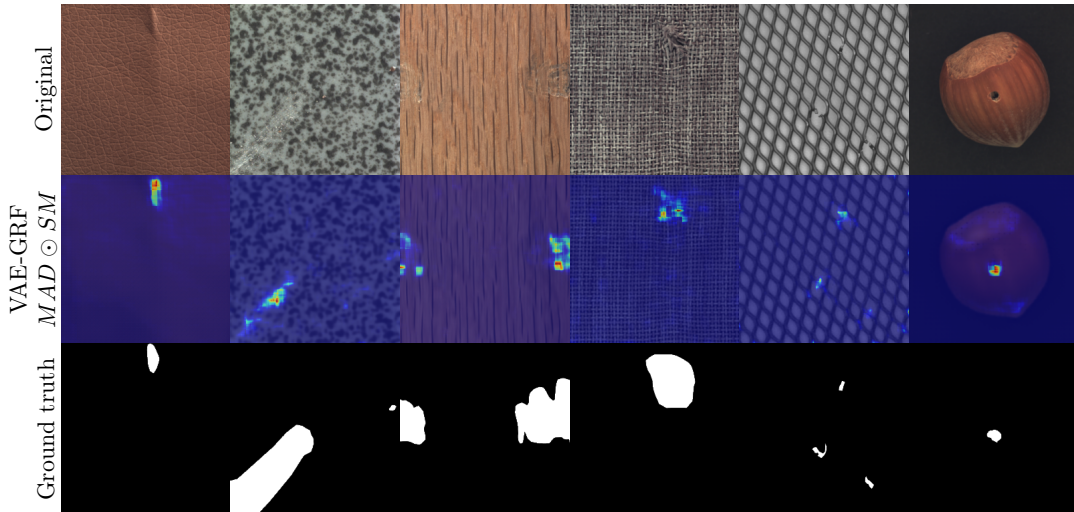


Figure 3: Selected illustrations for the MVTEc experiment. Anomaly maps and reconstructions from our proposed approach. *Top row*: original image. *Middle row*: the anomaly maps overlaying the VAE-GRF reconstruction. *Bottom row*: The segmented anomalies. From left to right: the *Leather*, *Tile*, *Wood*, *Carpet*, *Grid* and *Hazelnut* categories.

AD described previously is used here. Indeed, many images are empty in our datasets and we can consider animals as anomalies.

In the following experiments, we show that the VAE-GRF model and the associated metrics are relevant and provide competitive results against state-of-the-art approaches.

4.1. Livestock dataset

4.1.1. Presentation

The open Livestock dataset [15] regroups aerial images of livestock over grassland in which animals have been annotated. We again wish to perform an unsupervised detection of the animals using an AD approach. It is possible since we have a total of 3430 empty images to perform training and 890 images containing at least one animal.

4.1.2. Experiments & Results

The network described in Sec. 3.5 is used here, as well as the pixel-wise AD metrics from Sec. 3.2.

Tab. 3 gives the score in terms of pixel-wise ROCAUC, for all the models (we have discarded some of the worst performing approaches, ranking according to the MVTEC results). We can see that the VAE-GRF performs best with the $MAD \odot SM$ metric. Again, the improvements brought by the GRF prior could be expected since the images seem to really fit the stationary GRF hypothesis. Fig. 4 provides a graphical illustration of the experiment over some images of the dataset.

Remark: The MAD metric is all the more interesting as it is a metric which is only computed from the latent space, *i.e.*, at testing time, we do not need to perform any computations with the decoder. This then represents an interesting computational gain (one can calculate that this saves more than 40,000 two dimensional convolutions when using the decoder described in Sec. 3.4).

	VAE $MAD \odot SM$	VAE-GRF $MAD \odot SM$	Zimmerer et al. [38] $ \nabla_{\mathbf{x}} \mathcal{E}_{\theta, \varphi}(\mathbf{x}) $ $ \nabla_{\mathbf{x}} \mathcal{E}_{\theta, \varphi}(\mathbf{x}) \odot SM$		Dehaene et al. [9] SM grad
ROCAUC	<u>0.77</u>	0.80	<u>0.77</u>	0.52	0.68

Table 3: ROCAUC scores for pixel-wise AD on the images from the Livestock dataset. For VAE and VAE-GRF model, we have $\beta = 0.1$. We used the Matern- $\frac{3}{2}$ correlation function. Best scores appear in bold and second best are underlined. Scores from the main contribution of this article are in purple.

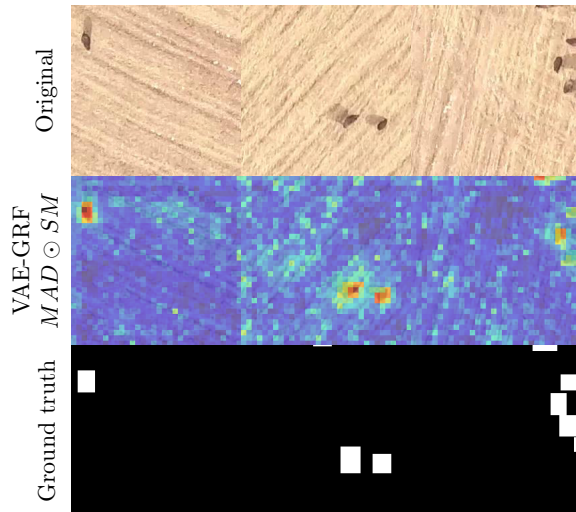


Figure 4: Selected illustrations of the detection of our model in the Livestock dataset experiment using the VAE-GRF model with Matern- $\frac{3}{2}$ correlation function. *Top row*: original image with ground truth. *Middle row*: VAE-GRF reconstruction with the anomaly map overlaid. *Bottom row*: ground truth segmentation.

Moreover, being independent of the reconstruction can be very interesting on complex datasets where reconstructions are unreliable as it will be shown in the next experiment.

4.2. Semmacape dataset

4.2.1. Presentation

The Semmacape dataset⁴ comprises 165 aerial images collected in the Gironde estuary and Pertuis sea Marine Nature Park, France, in 2020. Birds and dol-

⁴<https://semmacape.irisa.fr/>

phins have been manually annotated, the images were then subdivided into patches giving rise to 345 images with an animal and 138,544 empty images. Such a large number of empty images enables us to learn the normality (surface of the sea) and then to detect animals as anomalies. We will perform pixel-wise AD assessment with the previously introduced models.

4.2.2. Experiments & Results

The network described in Sec. 3.5 is used here; we are in the same experimental setting as before.

Tab. 4 gives the scores for the models in terms of precision, recall and F1 scores. First of all, probably because of the complexity of the dataset (many images are corrupted by the sun glare), the SM anomaly map seems particularly unreliable. Among all comparable models, the VAE-GRF with $MAD \odot SM$ metric seems to be the best performing, we argue that this is due to the fact that it operates using information from the latent space. It is thus less insensitive to the reconstructed images, which are of rather poor quality because of the noise corruption (a low SM metric was observed in practice). We also provide a comparison to the best results achieved so far on this dataset in [3] with a much more complex model combining the PaDiM approach [8] and Normalizing Flows (NF) [28]. We can conclude that our VAE-GRF is competitive with the PaDiM + NF (which is not a generative approach). We also notice that, when brought to this complex and real dataset, the performances of the approaches from [38] and [9] collapse. Fig. 5 graphically illustrates the experiment.

5. Conclusion

We have introduced a stationary GRF prior in the VAE model and shown how such a more complex prior (with full covariance matrix) can be embedded in a VAE model while preserving efficient computations. We have also demonstrated, in the context of AD, that such refinement in the modeling is relevant

	VAE	VAE-GRF	Zimmerer et al. [38]		Dehaene et al. [9]	Berg et al. [3]
	MAD	MAD	$ \nabla_{\mathbf{x}} \mathcal{E}_{\theta, \varphi}(\mathbf{x}) $	$ \nabla_{\mathbf{x}} \mathcal{E}_{\theta, \varphi}(\mathbf{x}) \odot SM$	SM grad	PaDiM + NF
F1	0.514	0.714	0.281	0.138	0.099	<u>0.530</u>
Recall	<u>0.677</u>	0.600	0.349	0.134	0.130	0.757
Precision	0.415	0.881	0.236	0.143	0.080	<u>0.408</u>

Table 4: ROCAUC scores on image-wise AD on the Semmacape dataset. Scores for PaDiM + NF are from [3]. For the VAE and the VAE-GRF model, we have $\beta = 1$ and we used the Matern- $\frac{3}{2}$ correlation function. Best scores appear in bold and second best are underlined. Scores from the main contribution of this article are in purple.

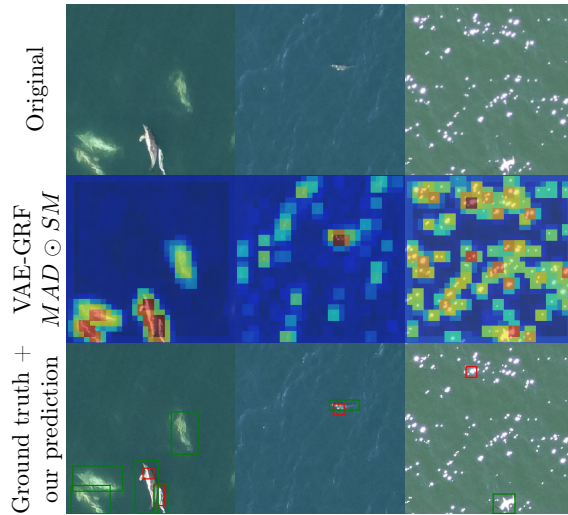


Figure 5: Selected illustrations of the detection of our model in the Semmacape dataset experiment using the VAE-GRF model with Matern- $\frac{3}{2}$ correlation function. *Top row*: original image. *Bottom row*: VAE-GRF reconstruction with the anomaly map overlaid. *Bottom row*: with ground truth bounding box (green) and prediction (red). The right column describes a typical very complex image, with a lot of sun glares which are moreover confused with the white bird (at the bottom).

for specific images, notably, for texture images. To do so, we have introduced two new metrics yielding competitive results against comparable state-of-the-art VAE models for AD tasks on several datasets. Our results suggest that the VAE-GRF might replace the VAE baseline for many tasks, as we show how the stationary assumption does not introduce any additional computational cost. Indeed, VAE-GRF might offer an efficient and relevant prior for many practical applications, especially for the processing of images that exhibit textures.

Since the stationary assumption for the GRF prior remains a strong assumption, future work might consider relaxing this assumption and we might study ways to introduce non-stationary GRF prior while preserving the tractability of the model.

Acknowledgements

This work was done as a part of the Game of Trawls project. We thank the European Maritime and Fisheries Fund (contract number 18/2216442) and France Filière Pêche (contract number 19/1000544) for funding. It was also supported by the SEMMACAPE project, which benefits from an ADEME (*Agence de la transition écologique*) grant under the “Sustainable Energies” call for research projects (2018–2019).

References

- [1] J. G. A. Barbedo, L. V. Koenigkan, T. T. Santos, and P. M. Santos. A study on the detection of cattle in UAV images using deep learning. *Sensors*, 19(24):5436, 2019.
- [2] C. Baur, S. Denner, B. Wiestler, N. Navab, and S. Albarqouni. Autoencoders for unsupervised anomaly segmentation in brain MR images: a comparative study. *Medical Image Analysis*, page 101952, 2021.

- [3] P. Berg, D. Santana Maia, M.-T. Pham, and S. Lefèvre. Weakly supervised detection of marine animals in high resolution aerial images. *Remote Sensing*, 2021.
- [4] P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger. MVTec AD—A comprehensive real-world dataset for unsupervised anomaly detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9592–9600, 2019.
- [5] L. B. Boudaoud, F. Maussang, R. Garello, and A. Chevallier. Marine bird detection based on deep learning using high-resolution aerial images. In *OCEANS 2019-Marseille*, pages 1–7, 2019.
- [6] F. P. Casale, A. V. Dalca, L. Saglietti, J. Listgarten, and N. Fusi. Gaussian process prior variational autoencoders. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 10390–10401, 2018.
- [7] T. R. Davidson, L. Falorsi, N. De Cao, T. Kipf, and J. M. Tomczak. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891*, 2018.
- [8] T. Defard, A. Setkov, A. Loesch, and R. Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, pages 475–489, 2021.
- [9] D. Dehaene, O. Frigo, S. Combrexelle, and P. Eline. Iterative energy-based projection on a normal data manifold for anomaly localization. In *International Conference on Learning Representations*, 2020.

- [10] N. Dilokthanakul, P. A. Mediano, M. Garnelo, M. C. Lee, H. Salimbeni, K. Arulkumaran, and M. Shanahan. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*, 2016.
- [11] V. Fortuin. Priors in Bayesian deep learning: A review. *International Statistical Review*, 2022.
- [12] C. Fox and R. A. Norton. Fast sampling in a linear-gaussian inverse problem. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1191–1218, 2016.
- [13] H. Gangloff, M.-T. Pham, L. Courtrai, and S. Lefèvre. Leveraging vector-quantized variational autoencoder inner metrics for anomaly detection. In *2022 26th International Conference on Pattern Recognition (ICPR)*, 2022. URL <https://hal.archives-ouvertes.fr/hal-03541964>.
- [14] P. Guttorp and T. Gneiting. Studies in the history of probability and statistics XLIX on the Matérn correlation family. *Biometrika*, 93(4):989–995, 2006.
- [15] L. Han, P. Tao, and R. R. Martin. Livestock detection in aerial images using a fully convolutional network. *Computational Visual Media*, 5(2):221–228, 2019.
- [16] D. M. Hawkins. *Identification of Outliers*. Monographs on applied probability and statistics. Chapman and Hall, 1980.
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE/CVF conference on computer vision and pattern recognition*, pages 770–778, 2016.

- [18] I. Higgins, L. Matthey, A. Pal, C. P. Burgess, X. Glorot, M. M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- [19] M. Jazbec, M. Pearce, and V. Fortuin. Factorized Gaussian process variational autoencoders. *arXiv preprint arXiv:2011.07255*, 2020.
- [20] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- [21] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- [22] W. Liu, R. Li, M. Zheng, S. Karanam, Z. Wu, B. Bhanu, R. J. Radke, and O. Camps. Towards visually explaining variational autoencoders. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8642–8651, 2020.
- [23] Y. Liu, J. Li, S. Sun, and B. Yu. Advances in Gaussian random field generation: a review. *Computational Geosciences*, 23(5):1011–1047, 2019.
- [24] P. Liznerski, L. Ruff, R. A. Vandermeulen, B. J. Franks, M. Kloft, and K.-R. Müller. Explainable deep one-class classification. In *International Conference on Learning Representations*, 2021.
- [25] G. Loaiza-Ganem and J. P. Cunningham. The continuous Bernoulli: fixing a pervasive error in variational autoencoders. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Gar-

- nett, editors, *Advances in Neural Information Processing Systems 32*, pages 13266–13276, 2019.
- [26] C. Padubidri, A. Kamilaris, S. Karatsiolis, and J. Kamminga. Counting sea lions and elephants from aerial photography using deep learning with density maps. *Animal Biotelemetry*, 9(1):1–10, 2021.
- [27] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel. Deep learning for anomaly detection: A review. *ACM Computing Surveys*, 54(2):1–38, 2021.
- [28] G. Papamakarios, I. Murray, and T. Pavlakou. Masked autoregressive flow for density estimation. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2338–2347, 2017.
- [29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Z. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035, 2019.
- [30] M. Pearce. The Gaussian process prior VAE for interpretable latent dynamics from pixels. In *Symposium on Advances in Approximate Bayesian Inference*, pages 1–12, 2020.
- [31] M. A. F. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko. A review of novelty detection. *Signal Processing*, 99:215–249, 2014.
- [32] M. A. Riedmiller and H. Braun. A direct adaptive method for

- faster backpropagation learning: the RPROP algorithm. In *Proceedings of International Conference on Neural Networks (ICNN'88)*, San Francisco, CA, USA, March 28 - April 1, 1993, pages 586–591. IEEE, 1993. doi: 10.1109/ICNN.1993.298623. URL <https://doi.org/10.1109/ICNN.1993.298623>.
- [33] H. Rue and L. Held. *Gaussian Markov random fields: theory and applications*. Chapman and Hall/CRC, 2005.
- [34] L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, and K. Müller. A unifying review of deep and shallow anomaly detection. *Proc. IEEE*, 109(5):756–795, 2021.
- [35] S. Venkataramanan, K.-C. Peng, R. V. Singh, and A. Mahalanobis. Attention guided anomaly localization in images. In *European Conference on Computer Vision*, pages 485–503. Springer, 2020.
- [36] L. Wang, D. Zhang, J. Guo, and Y. Han. Image anomaly detection using normal data only by latent space resampling. *Applied Sciences*, 10(23):8660, 2020.
- [37] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [38] D. Zimmerer, F. Isensee, J. Petersen, S. Kohl, and K. Maier-Hein. Un-supervised anomaly localization using variational auto-encoders. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 289–297, 2019.

Appendix A. Spectral computations

We now detail the formulas mentioned in Sec. 2.1.2. More details on the elements of this section can be found in [33]. Very useful properties, linked with the Fourier transform, are available for circulant and block-circulant matrices. We only recall the formulas for block-circulant matrices that are used in this study. In the following properties, DFT2 (resp. IDFT2) is the 2 dimensional (resp. inverse) Fourier transform, \odot is the element-wise matrix multiplication and \bullet is the element-wise exponentiation of the elements of the matrix. Note that in the next equations, we consider orthonormal Fourier transforms (normalized by $\frac{1}{\sqrt{l_x l_y}}$, where $l_x \times l_y$ is the matrix dimensionality).

Appendix A.1. Inverse of a block-circulant matrix

Let C be a block-circulant matrix, then we have

$$\text{base}(C^{-1}) = \frac{1}{l_x l_y} \text{IDFT2}(\text{DFT2}(\text{base}(C)) \bullet (-1)). \quad (\text{A.1})$$

Appendix A.2. Product of block-circulant matrices

Let C and D be two block-circulant matrices, then we have

$$\text{base}(CD) = \sqrt{l_x l_y} \text{IDFT2}(\text{DFT2}(\text{base}(C)) \odot \text{DFT2}(\text{base}(D))). \quad (\text{A.2})$$

Appendix A.3. Eigenvalues of a block-circulant matrix

Let C be a block-circulant matrix, then the matrix filled with the eigenvalues of C is

$$\Lambda = \sqrt{l_x l_y} \text{DFT2}(\text{base}(C)). \quad (\text{A.3})$$

Appendix A.4. Product of block-circulant matrix with vector

Let \mathbf{v} be the column major vector obtained from $n \times N$ matrix V , then $\mathbf{u} = C\mathbf{v}$ is obtained from

$$U = \sqrt{l_x l_y} \text{DFT2}(\text{DFT2}(\text{base}(C)) \odot \text{IDFT2}(V)). \quad (\text{A.4})$$

Category		VAE-GRF with exponential CF			VAE-GRF with Matern- $\frac{3}{2}$ CF			VAE-GRF with Gaussian CF		
		SM	MAD	$MAD \odot SM$	SM	MAD	$MAD \odot SM$	SM	MAD	$MAD \odot SM$
Stationary textures	Leather $\beta = 100$	0.85	0.91	0.94	<u>0.95</u>	<u>0.95</u>	0.98	0.88	0.94	<u>0.95</u>
	Tile $\beta = 0.01$	0.77	0.83	0.84	<u>0.93</u>	0.87	0.95	0.78	0.57	<u>0.77</u>
	Wood $\beta = 100$	0.73	0.75	0.78	0.78	0.86	<u>0.80</u>	0.72	0.68	0.75
Non stationary textures	Carpet $\beta = 0.01$	0.86	0.70	<u>0.90</u>	<u>0.90</u>	0.61	0.91	0.86	0.65	0.89
	Grid $\beta = 0.01$	0.82	0.64	0.85	<u>0.95</u>	0.76	0.97	0.82	0.61	0.85
	Hazelnut $\beta = 1$	0.95	0.80	<u>0.96</u>	<u>0.98</u>	0.91	0.98	0.94	0.84	<u>0.96</u>

Table B.5: ROCAUC scores for pixel-wise AD on the texture images from the MVTEC dataset. CF stands for Correlation Function. Best scores appear in bold and second best are underlined.

Appendix B. Additional experiments

Appendix B.1. Impact of the correlation function

We report in Tab. B.5 the results of the VAE-GRF model on the MVTEC dataset for different choice of correlation function use to parameterize the prior. We see that the results with the Matern-3/2 correlation function are better and more consistent across the experiments, hence our choice in the rest of the article. We also find again that the SM metric can always be improved using the MAD metric, for example, through the straightforward element-wise product that we proposed in $MAD \odot SM$.

Appendix B.2. Impact of the β hyperparameter

We report in Tab. B.6 the results of the VAE-GRF model on the MVTEC dataset for different choice of the hyperparameter β . First we see that it seems quite hard to predict the best β *a priori*. Second, the results when $\beta = 0$ (i.e. when our model becomes a deterministic AutoEncoder (AE)) are in accordance with those known in the literature for this model (see e.g. [4]). As expected, the MAD metric gives its worst results when $\beta = 0$; the AE model proscribes its usage. Indeed our hypothesis of GRF prior upon which the metric relies is not

Category		AutoEncoder ($\beta = 0$)			VAE-GRF with $\beta = 0.01$		
		SM	MAD	$MAD \odot SM$	SM	MAD	$MAD \odot SM$
Stationary textures	Leather	0.63	0.80	0.72	0.64	0.80	0.73
	Tile	<u>0.93</u>	0.64	0.92	<u>0.93</u>	0.87	0.95
	Wood	0.78	0.68	<u>0.80</u>	0.75	0.72	0.78
Non stationary textures	Carpet	<u>0.90</u>	0.38	0.87	0.90	0.61	0.91
	Grid	<u>0.96</u>	0.46	0.94	0.95	0.76	0.97
	Hazelnut	0.95	0.93	0.95	0.95	0.90	<u>0.96</u>

Category		VAE-GRF with $\beta = 1$			VAE-GRF with $\beta = 100$		
		SM	MAD	$MAD \odot SM$	SM	MAD	$MAD \odot SM$
Stationary textures	Leather	0.91	0.93	<u>0.96</u>	0.95	0.95	0.98
	Tile	0.80	0.80	0.84	0.39	0.77	0.65
	Wood	0.64	0.79	0.78	0.86	<u>0.80</u>	<u>0.80</u>
Non stationary textures	Carpet	0.86	0.73	<u>0.90</u>	0.47	0.63	0.60
	Grid	0.82	0.61	0.84	0.74	0.63	0.74
	Hazelnut	0.98	0.91	0.98	0.98	0.90	0.98

Table B.6: ROCAUC scores for pixel-wise AD on the texture images from the MVTEC dataset. The Matern- $\frac{3}{2}$ correlation function is used. Best scores appear in bold and second best are underlined.

enforced at all in this case. The lack of interpretability of the latent space is a common drawback of AE architectures.