

Looking at the whole picture: constrained unsupervised anomaly segmentation

Julio Silva-Rodríguez
jjsilva@upv.es

Institute of Transport and Territory
Universitat Politècnica de València
Valencia, Spain

Valery Naranjo
vnaranjo@dc.com.upv.es

Institute of Research and Innovation in
Bioengineering
Universitat Politècnica de València
Valencia, Spain

Jose Dolz
jose.dolz@etsmtl.ca

LIVIA Laboratory
École de Technologie Supérieure (ETS)
Montreal, Canada

Abstract

Current unsupervised anomaly localization approaches rely on generative models to learn the distribution of normal images, which is later used to identify potential anomalous regions derived from errors on the reconstructed images. However, a main limitation of nearly all prior literature is the need of employing anomalous images to set a class-specific threshold to locate the anomalies. This limits their usability in realistic scenarios, where only normal data is typically accessible. Despite this major drawback, only a handful of works have addressed this limitation, by integrating supervision on attention maps during training. In this work, we propose a novel formulation that does not require accessing images with abnormalities to define the threshold. Furthermore, and in contrast to very recent work, the proposed constraint is formulated in a more principled manner, leveraging well-known knowledge in constrained optimization. In particular, the equality constraint on the attention maps in prior work is replaced by an inequality constraint, which allows more flexibility. In addition, to address the limitations of penalty-based functions we employ an extension of the popular log-barrier methods to handle the constraint. Comprehensive experiments on the popular BRATS'19 dataset demonstrate that the proposed approach substantially outperforms relevant literature, establishing new state-of-the-art results for unsupervised lesion segmentation.

1 Introduction

Under the supervised learning paradigm, deep learning models have achieved astonishing performance in a wide range of applications. Nevertheless, a main limitation of these models is the large amount of labeled data required for training. Obtaining such curated labeled datasets is a cumbersome process prone to subjectivity, which makes access to sufficient training data difficult in practice. This problem is further magnified in the context of medical image segmentation, where labeling involves assigning a category to each image pixel

resulting impractical when volumetric data is involved. In addition, even if annotated images are available, there exist some applications, such as brain lesion detection, where large intraclass variations are not captured during training, failing to cover the broad range of abnormalities that might be present in a scan. This makes that, in a fully-supervised setting, deep models might have difficulties when learning from such class-imbalanced training sets. Thus, considering the scarcity and the diversity of target objects in these scenarios, lesion segmentation is typically modeled as an anomaly localization task, which is trained in an unsupervised manner. In particular, the training dataset contains only *normal* images and *abnormal* images are not accessible during training.

A common strategy for unsupervised anomaly segmentation is to model the distribution of normal images, for which generative models, such as generative adversarial networks (GANs) [2, 7, 51, 53, 54, 56] and variational auto-encoders (VAEs) [11, 12, 28, 32, 40] have been widely employed. To achieve this, input images are compared to their reconstructed normal counterparts, which are recovered from the learned distribution, and anomalies are identified from the reconstruction error. Nevertheless, these methods require to learn a threshold to estimate the pixel-wise difference between the input and its reconstructed image in order to localize abnormalities. As the threshold needs to be computed based on abnormal training images, this limits their usability in realistic scenarios, where only normal data is provided. Inspired by the observations that attention-based supervision can alleviate the need of large labeled training data [23], class-activation maps have been integrated in the training. In particular, [57] leverage the generated attention maps as an additional supervision cue, enforcing the network to provide attentive regions covering the whole context in normal images. This term was formulated as an equality constraint with the form of a L_1 penalty over each individual pixel. Nevertheless, we found that explicitly forcing the network to produce maximum attention values across each pixel does not achieve satisfactory results in the context of brain lesion segmentation. In addition, recent literature in constrained optimization for deep neural networks suggests that simple penalties –such as the function used in [57]– might not be the optimal solution to constraint the output of a CNN [20].

Based on these observations, we propose a novel formulation for unsupervised semantic segmentation of brain lesions in medical images. The key contributions of our work can be summarized as follows:

- A novel constrained formulation for unsupervised anomaly localization, which integrates an auxiliary size-constrained loss to force the network to generate class activation masks (CAMs) that cover the whole context in normal images.
- In particular, size information is imposed through inequality constraints on the region proportion of generated CAMs, which give more flexibility than the pixel-wise equality constraint in [57]. In addition, to address the limitations of penalty-based functions, we resort to an extended version of the standard log-barrier.
- Furthermore, while our method yields significant improvements when anomalous images are used to define a class-specific threshold to locate anomalies –following the literature–, our formulation still outperforms existing approaches without accessing to anomalous images, which contrasts to most prior works.
- We benchmark the proposed model against a relevant body of literature on the popular BRATS challenge dataset. Comprehensive experiments demonstrate the superior performance of our model, establishing a new state-of-the-art for this task.

2 Related Work

Unsupervised Anomaly Segmentation. Unsupervised anomaly segmentation aims at identifying abnormal pixels on test images, containing, for example, lesions on medical images [2, 11], defects in industrial images [10, 24, 57] or abnormal events in videos [1, 51]. A main body of the literature has explored unsupervised deep (generative) representation learning to learn the distribution from normal data. The underlying assumption is that a model trained on normal data will not be able to reconstruct anomalous regions, and the reconstructed difference can therefore be used as an anomaly score. Under this learning paradigm, generative adversarial networks (GAN) [5] and variational auto-encoders (VAE) [22] are typically employed. Nevertheless, even though GAN and VAE model the latent variable, the manner in which they approximate the distribution of a set of samples differs. GAN-based approaches [2, 1, 51, 53, 54, 56] approximate the distribution by optimizing a generator to map random samples from a prior distribution in the latent space into data points that a trained discriminator cannot distinguish. On the other hand, data distribution is approximated in VAE by using variational inference, where an encoder approximates the posterior distribution in the latent space and a decoder models the likelihood [11, 22, 13, 28, 52, 40]. In the context of medical images, and more relevant to this work, several works have proposed different improvements on VAEs [11, 12, 28], to overcome specific limitations. For example, to handle the lack of consistency in the learned latent representation on prior works, [11] included a constraint that helps mapping an image containing abnormal anatomy close to its corresponding healthy image in the latent space. A detailed survey on unsupervised anomaly localization in medical imaging can be found in [9]. Nevertheless, a main limitation of these approaches is that the threshold to estimate the pixel-wise anomaly score has to be computed in images with anomalies, which might not be available in practice. To alleviate this issue, [57] propose to integrate attention maps from Grad-CAM [55] during the training as supervisory signals. In particular, in addition to standard learning objectives, authors employ an auxiliary loss that tries to maximize the attention maps on normal images by including an equality constraint with the form of a L_1 penalty over each individual pixel.

Constrained segmentation. Imposing global constraints on the output predictions of deep CNNs has gained attention recently, particularly in weakly supervised segmentation. These constraints can be embedded into the network outputs in the form of direct loss functions, which guide the network training when fully labeled images are not accessible. For example, a popular scenario is to enforce the softmax predictions to satisfy a prior knowledge on the size of the target region. Jia et al. [18] employed a L_2 penalty to impose equality constraints on the size of the target regions in the context of histopathology image segmentation. In [53], authors leverage the target properties by enforcing the label distribution of predicted images to match an inferred label distribution of a given image, which is achieved with a KL-divergence term. Similarly, Zhou et al. [59] proposed a novel loss objective in the context of partially labeled images, which integrated an auxiliary term, based on a KL-divergence, to enforce that the average output size distributions of different organs approximates their empirical distributions, obtained from fully-labeled images. While the equality-constrained formulations proposed in these works are very interesting, they assume exact knowledge of the target size prior. In contrast, inequality constraints can relax this assumption, allowing much more flexibility. In [27], authors imposed inequality constraints on a latent distribution –which represents a “fake” ground truth– instead of the network output, to avoid the computational complexity of directly using Lagrangian-dual optimization. Then, the net-

work parameters are optimized to minimize the KL divergence between the network softmax probabilities and the latent distribution. Nevertheless, their formulation is limited to linear constraints. More recently, inequality constraints have been tackled by augmenting the learning objective with a penalty-based function, e.g., L_2 penalty, which can be imposed within a continuous optimization framework [6, 19, 20], or in the discrete domain [29]. Despite these methods have demonstrated remarkable performance in weakly supervised segmentation, they require that prior knowledge, *exact* or *approximate*, is given. This contrasts with the proposed approach, which is trained on data without anomalies, and hence the size of the target is zero.

3 Methods

An overview of our method is presented in Fig. 1, and we describe each component below.

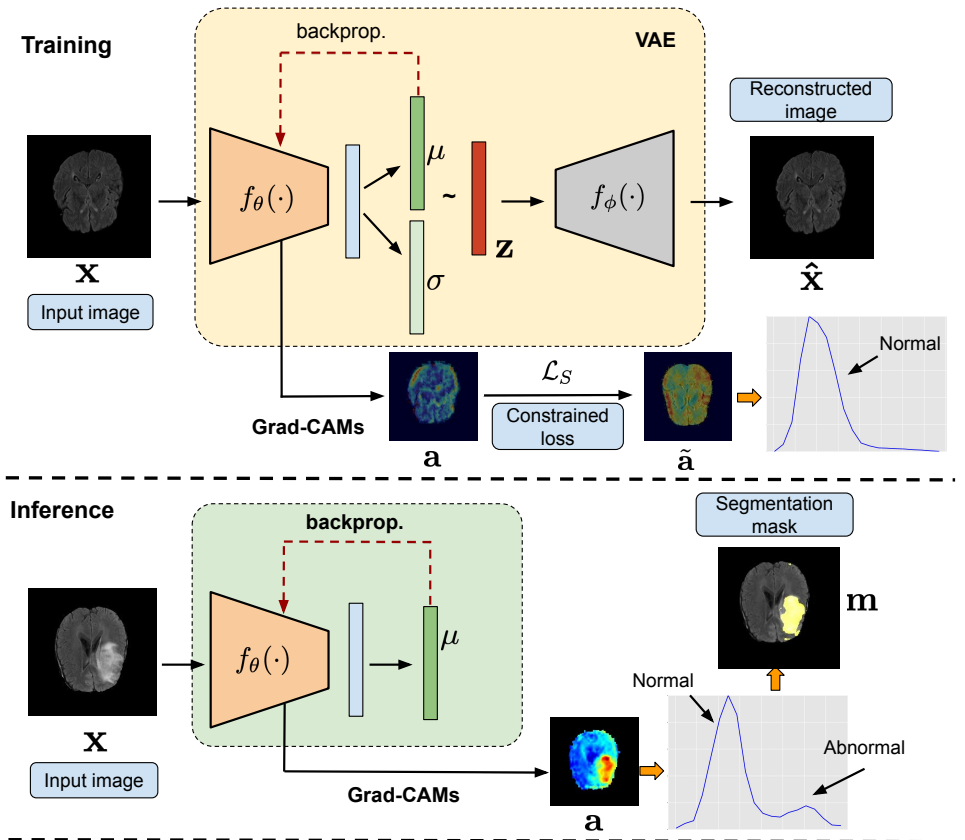


Figure 1: **Method overview.** Following the standard literature, the VAE is optimized to maximize the evidence lower bound (ELBO), which satisfies Eq. 1. In addition, we include a size regularizer (in the form of a size-constrained loss \mathcal{L}_S) on the attention maps \mathbf{a} , to force the network to search in the whole image. At inference, the attention map is thresholded to obtain the final segmentation mask \mathbf{m} .

Preliminaries. Let us denote the set of unlabeled training images as $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$, where $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^{\Omega_i}$ represents the i^{th} image and Ω_i denotes the spatial image domain. This dataset contains only normal images, e.g., healthy images in the medical context. We also define an encoder, $f_\theta(\cdot) : \mathcal{X} \rightarrow \mathcal{Z}$, parameterized by θ , which is optimized to project normal data points in \mathcal{D} into a manifold represented by a lower dimensionality d , $\mathbf{z} \in \mathcal{Z} \subset \mathbb{R}^d$. Furthermore, a decoder $f_\phi(\cdot) : \mathcal{Z} \rightarrow \mathcal{X}$ parameterized by ϕ tries to reconstruct an input image $\mathbf{x} \in \mathcal{X}$ from $\mathbf{z} \in \mathcal{Z}$, which results in $\hat{\mathbf{x}} = f_\phi(f_\theta(\mathbf{x}))$.

3.1 Vanilla VAE

A Variational Autoencoder (VAE) is an encoder-decoder style generative model, which is currently the dominant strategy for unsupervised anomaly location. Training a VAE consists on minimizing a two-term loss function, which is equivalent to maximize the evidence lower-bound (ELBO) [22]:

$$\mathcal{L}_{VAE} = \mathcal{L}_R(\mathbf{x}, \hat{\mathbf{x}}) + \beta \mathcal{L}_{KL}(q_\theta(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \quad (1)$$

where \mathcal{L}_R is the reconstruction error term between the input and its reconstructed counterpart. The right-hand term is the Kullback-Leibler (KL) divergence (weighted by β) between the approximate posterior $q_\theta(\mathbf{z}|\mathbf{x})$ and the prior $p(\mathbf{z})$, which acts as a regularizer, penalizing approximations for $q_\theta(\mathbf{z}|\mathbf{x})$ that differ from the prior.

3.2 Size regularizer via VAE attention

Similar to very recent literature [5], we integrate supervision on attention maps during training. In particular, attention maps $\mathbf{a} \in \mathbb{R}^{\Omega_i}$ are generated from the latent mean vector \mathbf{z}_μ , by using Grad-CAM [5] via backpropagation to an encoder block output $f_\theta^s(\mathbf{x})$, at a given network depth s . Thus, for a given input image \mathbf{x}^j its corresponding attention map is computed as $\mathbf{a}^j = \sigma(\sum_k^K \alpha_k f_\theta^s(\mathbf{x}^j)_k)$, where K is the total number of filters of that encoder layer, σ a sigmoid operation, and α_k are the generated gradients such that: $\alpha_k = \frac{1}{|\mathbf{a}^j|} \sum_{l \in \Omega_T} \frac{\partial z_\mu}{\partial \mathbf{a}_{k,l}^j}$, where Ω_T is the spatial features domain.

In [5], authors leveraged the attention maps by enforcing them to cover the whole normal image. To achieve this, their loss function was augmented with an additional term, referred to as expansion loss, which takes the form of: $\mathcal{L}_s = \frac{1}{|\mathbf{a}^j|} \sum_{l \in \Omega_i} (1 - \mathbf{a}_l^j)$. We can easily observe that this term resembles to an equality constraint, forcing the class activation maps to be maximum at the whole image in a pixel-wise manner (i.e., it penalizes each single pixel individually). Contrary to this work, we integrate supervision on attention maps by enforcing inequality constraints on its global target size, which allows much more flexibility, particularly when integrating the notion of expected target size, as we describe below. Indeed, as demonstrated by our results, relaxing the assumption that CAMs in normal images should cover the whole image brings a substantial performance gain. Hence, we aim at minimizing the following constrained optimization problem:

$$\min_{\theta, \phi} \mathcal{L}_{VAE}(\theta, \phi) \quad \text{s.t.} \quad f_c(\mathbf{a}^n) \leq 0, \quad n = 1, \dots, N \quad (2)$$

where $f_c(\mathbf{a}^j) = (1 - \frac{1}{|\Omega_i|} \sum_{l \in \Omega_i} \mathbf{a}_l^j) - p$ is the constraint over the attention map from the j -th image, which enforces the generated attention map to cover the whole image, relaxed by a certain margin p (in our context, p defines the proportion of pixels over an entire image).

To better highlight the advantages over [57], let us take for example the case where the size proportion is $p = 0.1$ (i.e., desired target size equal to 0.9, or 90% of the image). Following the formulation in [57], this is achieved when *all the pixel predictions are equal to 0.9*, resulting in a region covering the whole image once the class activation map is thresholded. In contrast, our formulation can yield to multiple solutions, as we do not constrain individual pixels to have a 0.9 value. For instance, 90% of the pixels having a prediction close to 1 and 10% of the pixels with a close to 0 prediction would be a valid solution, as it satisfies the global constraint. Furthermore, note that the gradients from both terms are also different. The term in [57] leads to different gradients at each logit, while our term backpropagates the same gradient value through all the logits, based on the global target size difference. Thus, both terms are fundamentally different and lead to different solutions.

From eq. 2 we can derive an approximate unconstrained optimization problem by employing a penalty-based method, which takes the hard constraint and moves it into the loss function as a penalty term ($\mathcal{P}(\cdot)$): $\min_{\theta, \phi} \mathcal{L}_{VAE}(\theta, \phi) + \lambda \mathcal{P}(f_c(\mathbf{a}))$. Thus, each time that the constrained $f_c(\mathbf{a}^n) \leq 0$ is violated, the penalty term $\mathcal{P}(f_c(\mathbf{a}^n))$ increases.

3.3 Extended log-barrier as an alternative to penalty-based functions

Despite having demonstrated a good performance in several applications [46, 48, 49, 27] penalty-based methods have several drawbacks. First, these unconstrained minimization problems have increasingly unfavorable structure due to ill-conditioning [4, 25], which typically results in an exceedingly slow convergence. And second, finding the optimal penalty weight is not trivial. To address these limitations, we replace the penalty-based functions by the approximation of log-barrier¹ presented in [20], which is formally defined as:

$$\tilde{\psi}_t(z) = \begin{cases} -\frac{1}{t} \log(-z) & \text{if } z \leq -\frac{1}{t^2} \\ tz - \frac{1}{t} \log(\frac{1}{t^2}) + \frac{1}{t} & \text{otherwise,} \end{cases} \quad (3)$$

where t controls the barrier over time, and z is the constraint $f_c(\mathbf{a}^n)$. Thus, by taking into account the approximation in 3, we can solve the following unconstrained problem by using standard Gradient Descent:

$$\min_{\theta, \phi} \underbrace{\mathcal{L}_{VAE}(\theta, \phi)}_{\text{Standard VAE loss}} + \lambda \underbrace{\sum_{n=1}^N \tilde{\psi}_t\left(\left(1 - \frac{1}{|\Omega_i|} \sum_{l \in \Omega_i} \mathbf{a}_l^n\right) - p\right)}_{\mathcal{L}_S: \text{Size regularizer}} \quad (4)$$

In this scenario, for a given t , the optimizer will try to find a solution with a good compromise between minimizing the loss of the VAE and satisfying the constraint $f_c(\mathbf{a}^n)$.

3.4 Inference

During inference, we use the generated attention as an anomaly saliency map. In order to avoid saturation caused by large activations, we replaced the sigmoid operation by a minimum-maximum normalization. Finally, the map is thresholded to locate the anomalous patterns in the image (different strategies to set the threshold are discussed later). Instead of inverting the generated attention maps, as in [57], we found that anomalous regions actually produce stronger gradients, which aligns with recent observations on natural [24] and brain

¹Note that this function is convex, continuous and twice-differentiable.

MRI [8] images. We believe that these larger gradients result in strongest activations, as CAM based on GradCAM are weighted by the input gradient.

4 Experiments

4.1 Experimental setting

Datasets. The experiments described in this work were carried out using the popular BraTS 2019 dataset [9, 4, 26], which contains 335 multi-institutional multi-modal MR scans with their corresponding Glioma segmentation masks. Following [8], from every patient, 10 consecutive axial slices of FLAIR modality of resolution 224×224 pixels were extracted around the center to get a pseudo MRI volume. Then, the dataset is split into training, validation and testing groups, with 271, 32 and 32 patients, respectively. Following the standard literature, during training only the slices without lesions are used as normal samples. For validation and testing, scans with less than 0.01% of tumour are discarded. A summary of the dataset used is presented in Table 1 in Supplemental Materials.

Evaluation Metrics. We resort to standard metrics for unsupervised brain lesion segmentation, as in [8]. Concretely, we compute the dataset-level area under precision-recall curve (AUPRC) at pixel level, as well the area under receptive-operative curve (AUROC). From the former, we obtain the operative point (OP) as threshold to generate the final segmentation masks. Then, we compute the best dataset-level DICE-score ([DICE]) and intersection-over-union ([IoU]) over these segmentation masks. Finally, we compute the average DICE over single scans. For each experiment, the metrics reported are the average of three consecutive repetitions of the training, to account for the variability of the stochastic factors involved in the process.

Implementation Details. The VAE architecture used in this work is based on the recently proposed framework in [5]. Concretely, the convolution layers of ResNet-18 [17] are used as the encoder, followed by a dense latent space $\mathbf{z} \in \mathbb{R}^{32}$. For image generation, a residual decoder is used, which is symmetrical to the encoder. It is noteworthy to mention that, even though several methods have resorted to a spatial latent space [6, 5], we observed that a dense latent space provided better results, which aligns to the recent benchmark in [8]. The VAE was trained during 400 iterations with eq. (1) to stabilize the convergence using $\beta = 1$. Then, the proposed regularizer was integrated (equation 4) with $t = 20$ and $\lambda = 10$, applied to the Grad-CAMs obtained from the first convolutional block of the encoder. We use a batch size of 32 images, and a learning rate of $1e-4$ with ADAM optimizer. The reconstruction loss, \mathcal{L}_R , in eq. (1) is the binary cross-entropy, and p in eq. (2) is set empirically to 0.2. Ablation experiments reported in the experimental section, which are performed on the validation set, empirically validate these choices. The code and trained models are publicly available on (https://github.com/cvblab/anomaly_localization_vae_gcams).

Baselines. In order to compare our approach to state-of-the-art methods, we implemented prior works and validated them on the dataset used, under the same conditions. First, we use residual-based methods to match the recently benchmark on unsupervised lesion localization in [8]. We also include recently proposed methods that integrate CAMs to locate anomalies. For both strategies, the training hyper-parameters and AE/VAE architectures were similar to the implementation of the proposed method. **Residual methods**, given an anomalous sample, aim to use the AE/VAE to reconstruct its normal counterpart. Then, they obtain an anomaly

localization map using the residual between both images such that $\mathbf{m} = |\mathbf{x} - \hat{\mathbf{x}}|$, where $|\cdot|$ indicates the absolute value. On the AE/VAE scenario, we include methods which propose modifications over vanilla versions, including context data augmentation in Context AE [40], Bayesian AEs [23], Restoration VAEs [12], an adversarial-based VAEs, AnoVAEGAN [6] and a recent GAN-based approach, F-anoGAN [34]. For methods including adversarial learning, DC-GAN [50] is used as discriminator. During inference, residual maps are masked using a slight-eroded brain mask, to avoid noisy reconstructions along the brain borderline. **CAMs-based:** we use Grad-CAM VAE [24], which obtains regular Grad-CAMs on the encoder from the latent space \mathbf{z}_μ of a trained vanilla VAE. Concretely, we include a disentanglement variant of CAMs proposed in this work, which computes the combination of individually-calculated CAMs from each dimension in \mathbf{z}_μ , referred to as Grad-CAM_D VAE. We also use the recent method in [57] (CAVGA), which applies a L1 penalty on the generated CAM to maximize the attention. In contrast to our model and [24], the anomaly mask in [57] is generated by focusing on the regions not activated on the saliency map such that $\mathbf{a} = 1 - \text{CAM}$, hypothesizing that the network has learnt to focus only on normal regions. Then, \mathbf{a} is thresholded with 0.5 to obtain the final anomaly mask $\mathbf{m} \in \mathbb{R}^{\Omega_i}$. For both methods, the network layer to obtain the Grad-CAMs is the same as in our method.

How the attention maps are thresholded? Nearly all prior approaches resort to anomalous images to define the threshold to obtain the final segmentation masks. In particular, these methods look at the AUPRC on the anomalous images, which is then used to compute the threshold value. Having access to images with anomalies is unrealistic in practice, and the value found might be biased towards the images employed. To alleviate this issue, we also evaluate our model in the scenario where the threshold is simply set to 0.5.

4.2 Results

Comparison to the literature. The quantitative results obtained by the proposed model and baselines on the test cohort are presented in Table 1. Results from the baselines range between [0.056-0.511](AUPRC) and [0.188-0.525] (DICE), which are in line with previous literature [8]. We can observe that the proposed methodology (*last row*) outperforms previous approaches by a large margin, with a substantial increase of $\sim 24\%$ and $\sim 18\%$ in terms of AUPRC and DICE, respectively, compared to the best prior model, i.e., F-anoGAN.

Threshold computed from normal or abnormal images? Furthermore, we can observe that while employing images with anomalies to select the optimal threshold yields the best result, this might be unrealistic in a fully unsupervised scenario. Nevertheless, an interesting property of our approach is that *it can still achieve large performance gains without having access to anomalous images to define the threshold*, unlike prior works.

Image vs. pixel-level constraint. The following experiment demonstrates the benefits of imposing the constraint on the whole image rather than in a pixel-wise manner, as in [57]. In particular, we compare the two strategies when the constraint is enforced via a L2-penalty function, whose results are presented in Table 2. These results illustrate the superiority of our method, which is consistent across every p value.

Extended log-barrier vs. penalty-based functions. To motivate the choice of employing the extended log-barrier over standard penalty-based functions in the constrained optimization problem in eq. (2), we compare them in Table 3. First, it can be observed that across

Method	AUROC	AUPRC	[DICE]	[IoU]	DICE ($\mu \pm \sigma$)
CAVGA [5]*	0.726	0.056	0.188	0.104	0.182 \pm 0.096
Bayesian VAE [28]	0.922	0.193	0.342	0.206	0.329 \pm 0.115
AnoVAEGAN [8]	0.925	0.232	0.359	0.221	0.349 \pm 0.115
Bayesian AE [28]	0.940	0.279	0.389	0.242	0.375 \pm 0.130
AE	0.937	0.261	0.397	0.248	0.386 \pm 0.125
Grad-CAM _D VAE [24]	0.941	0.312	0.400	0.250	0.361 \pm 0.164
Restoration VAE [12]	0.934	0.352	0.403	0.252	0.345 \pm 0.186
Context VAE [11]	0.939	0.271	0.406	0.255	0.394 \pm 0.126
Context AE [11]	0.940	0.278	0.411	0.259	0.399 \pm 0.126
VAE [8, 11]	0.940	0.273	0.411	0.259	0.399 \pm 0.127
F-anoGAN [24]	0.946	0.511	0.525	0.369	0.494 \pm 0.151
Proposed*(th=0.5)	0.981	0.753	0.606	0.439	0.582 \pm 0.238
Proposed	0.981	0.753	0.704	0.543	0.665 \pm 0.200

Table 1: Comparison to prior literature. (*) indicates that the class-specific threshold to locate the anomaly is set without accessing to abnormal images. Best results in bold.

Regularization	Size (proportion) term p						
	0	0.05	0.10	0.15	0.20	0.25	0.30
L2 (pixel-level)	0.489	0.288	0.275	0.329	0.288	0.264	0.201
L2 (image-level)	0.576	0.589	0.648	0.594	0.666	0.553	0.531

Table 2: Quantitative comparison, in terms of AUPRC, between enforcing the constraint at pixel-level (i.e., [8]) or at image-level (i.e., proposed approach) across different p values.

different p values, imposing the constraint with the extended log-barrier consistently outperforms the L_2 -penalty, with substantial performance gains. Furthermore, we empirically observe that despite any analyzed value of p outperforms current sota, setting $p = 0.2$ brings the largest performance gain.

Regularization	Size (proportion) term p						
	0	0.05	0.10	0.15	0.20	0.25	0.30
L2 (Penalty)	0.576	0.589	0.648	0.594	0.666	0.553	0.531
	0.581	0.581	0.631	0.589	0.632	0.557	0.521
Extended Log Barrier [21]	0.682	0.664	0.646	0.641	0.710	0.640	0.610
	0.653	0.638	0.625	0.623	0.661	0.625	0.600

Table 3: Impact of the type of regularization and the proportion term p on our method. We report the AUPRC (*top row*) and the best [DICE] (*bottom row*). Best results in bold.

In the Supplemental Materials, we provide comprehensive ablation experiments to validate several elements of our model, and motivate the choice of the values employed in our formulation, as well as our experimental setting.

Qualitative evaluation. Visual results of the proposed and existing methods are depicted in Figure 2. We can observe that our approach identifies as anomalous more complete regions of the lesions, whereas existing methods are prone to produce a significant amount of false positives (*top and bottom rows*) and fail to discover many abnormal pixels (*top row*).




Figure 2: Qualitative evaluation of the proposed and existing high-performing methods.

5 Conclusions

We proposed a novel constrained formulation for the task of unsupervised segmentation of brain lesions. In particular, we resort to generated CAMs to identify anomalous regions, which contrasts with most existing works that rely on the pixel-wise reconstruction error. Our formulation integrates a size-constrained loss that enforces the CAMs to cover the whole image in normal images. In contrast to very recent work, we tackle this problem by imposing inequality constraints on the whole target CAMs, which allows more flexibility than equality constraints over each single pixel. Last, and to overcome the limitations of penalty-based methods, we resort to an extension of standard log-barrier methods. Quantitative and qualitative results demonstrate that our model significantly outperforms prior literature on unsupervised lesion segmentation, without the need of accessing to anomalous images.

Acknowledgements

J. Silva-Rodríguez work was supported by the Spanish Government under FPI Grant PRE2018-083443. The DGX-A100 used in this work was partially funded by Generalitat Valenciana / European Union through the European Regional Development Fund (ERDF) of the Valencian Community (IDIFEDER/2020/030).

References

- [1] Davide Abati, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Latent space autoregression for novelty detection. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 481–490, 2019.
- [2] Simon Andermatt, Antal Horváth, Simon Pezold, and Philippe Cattin. Pathology segmentation using distributional differences to images of healthy origin. In *International MICCAI Brainlesion Workshop*, pages 228–238, 2018.

- [3] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S. Kirby, John B. Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing The Cancer Genome Atlas glioma MRI collections with expert segmentation labels and radiomic features. *Scientific Data*, 4(September):1–13, 2017.
- [4] Spyridon Bakas et al. Identifying the Best Machine Learning Algorithms for Brain Tumor Segmentation, Progression Assessment, and Overall Survival Prediction in the BRATS Challenge. 2018.
- [5] Mathilde Bateson, Hoel Kervadec, Jose Dolz, Hervé Lombaert, and Ismail Ben Ayed. Constrained domain adaptation for segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 326–334, 2019.
- [6] Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Deep autoencoding models for unsupervised anomaly segmentation in brain MR images. In *International MICCAI Brainlesion Workshop*, pages 161–169, 2019.
- [7] Christoph Baur, Robert Graf, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. SteGANomaly: Inhibiting CycleGAN steganography for unsupervised anomaly detection in brain MRI. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 718–727, 2020.
- [8] Christoph Baur, Stefan Denner, Benedikt Wiestler, Nassir Navab, and Shadi Albarqouni. Autoencoders for unsupervised anomaly segmentation in brain MR images: A comparative study. *Medical Image Analysis*, 69(8):1–16, 2021. ISSN 13618423. doi: 10.1016/j.media.2020.101952.
- [9] Christoph Baur, Stefan Denner, Benedikt Wiestler, Nassir Navab, and Shadi Albarqouni. Autoencoders for unsupervised anomaly segmentation in brain mr images: A comparative study. *Medical Image Analysis*, page 101952, 2021.
- [10] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. In *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*, 2019.
- [11] Xiaoran Chen and Ender Konukoglu. Unsupervised detection of lesions in brain MRI using constrained adversarial auto-encoders. In *International Conference on Medical Imaging with Deep Learning*, 2018.
- [12] Xiaoran Chen, Suhang You, Kerem Can Tezcan, and Ender Konukoglu. Unsupervised lesion detection via image restoration with a normative prior. *Medical Image Analysis*, 64, 2020.
- [13] David Dehaene, Oriel Frigo, Sébastien Combretelle, and Pierre Eline. Iterative energy-based projection on a normal data manifold for anomaly localization. In *International Conference on Learning Representations*, 2019.
- [14] Anthony V Fiacco and Garth P McCormick. *Nonlinear programming: sequential unconstrained minimization techniques*. SIAM, 1990.

- [15] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- [16] Frank S He, Yang Liu, Alexander G Schwing, and Jian Peng. Learning to play in a day: Faster deep reinforcement learning by optimality tightening. In *5th International Conference on Learning Representations, ICLR 2017*, 2017.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2016.
- [18] Zhipeng Jia, Xingyi Huang, I Eric, Chao Chang, and Yan Xu. Constrained deep weak supervision for histopathology image segmentation. *IEEE transactions on medical imaging*, 36(11):2376–2388, 2017.
- [19] Hoel Kervadec, Jose Dolz, Éric Granger, and Ismail Ben Ayed. Curriculum semi-supervised segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 568–576, 2019.
- [20] Hoel Kervadec, Jose Dolz, Meng Tang, Eric Granger, Yuri Boykov, and Ismail Ben Ayed. Constrained-CNN losses for weakly supervised segmentation. *Medical image analysis*, 54:88–99, 2019.
- [21] Hoel Kervadec, Jose Dolz, Jing Yuan, Christian Desrosiers, Eric Granger, and Ismail Ben Ayed. Constrained Deep Networks: Lagrangian Optimization via Log-Barrier Extensions. pages 1–23, 2019. URL <http://arxiv.org/abs/1904.04205>.
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2013.
- [23] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 9215–9223, 2018.
- [24] Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyang Wu, Bir Bhanu, Richard J Radke, and Octavia Camps. Towards visually explaining variational autoencoders. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 8642–8651, 2020.
- [25] David G Luenberger. *Introduction to linear and nonlinear programming*, volume 28. Addison-wesley Reading, MA, 1973.
- [26] Bjoern Menze et al. The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2015.
- [27] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 1796–1804, 2015.

- [28] Nick Pawlowski, MC Lee, Martin Rajchl, Steven McDonagh, Enzo Ferrante, Konstantinos Kamnitsas, Sam Cooke, Susan Stevenson, Aneesh Khetani, Tom Newman, et al. Unsupervised lesion detection in brain CT using bayesian convolutional autoencoders. *Medical Imaging with Deep Learning*, 2018.
- [29] Jizong Peng, Hoel Kervadec, Jose Dolz, Ismail Ben Ayed, Marco Pedersoli, and Christian Desrosiers. Discretely-constrained deep network for weakly supervised segmentation. *Neural Networks*, 130:297–308, 2020.
- [30] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *International Conference on Learning Representations, ICLR*, 2016.
- [31] Mahdyar Ravanbakhsh, Enver Sangineto, Moin Nabi, and Nicu Sebe. Training adversarial discriminators for cross-channel abnormal event detection in crowds. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 1896–1904, 2019.
- [32] Mohammad Sabokrou, Masoud Pourreza, Mohsen Fayyaz, Rahim Entezari, Mahmood Fathy, Jürgen Gall, and Ehsan Adeli. Avid: Adversarial visual irregularity detection. In *Asian Conference on Computer Vision*, pages 488–505, 2018.
- [33] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging*, pages 146–157, 2017.
- [34] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019.
- [35] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal of Computer Vision*, 128(2): 336–359, 2020.
- [36] Liyan Sun, Jiexiang Wang, Yue Huang, Xinghao Ding, Hayit Greenspan, and John Paisley. An adversarial learning approach to medical image synthesis for lesion detection. *IEEE journal of biomedical and health informatics*, 24(8):2303–2314, 2020.
- [37] Shashanka Venkataramanan, Kuan-Chuan Peng, Rajat Vikram Singh, and Abhijit Mahalanobis. Attention guided anomaly localization in images. In *European Conference on Computer Vision*, pages 485–503. Springer, 2020.
- [38] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2020–2030, 2017.
- [39] Yuyin Zhou, Zhe Li, Song Bai, Chong Wang, Xinlei Chen, Mei Han, Elliot Fishman, and Alan L Yuille. Prior-aware neural network for partially-supervised multi-organ segmentation. In *Proceedings of the International Conference on Computer Vision*, pages 10672–10681, 2019.

- [40] David Zimmerer, Simon Kohl, Jens Petersen, Fabian Isensee, and Klaus Maier-Hein. Context-encoding variational autoencoder for unsupervised anomaly detection. In *International Conference on Medical Imaging with Deep Learning—Extended Abstract Track*, 2019.
- [41] David Zimmerer, Fabian Isensee, Jens Petersen, Simon Kohl, and Klaus Maier-Hein. Abstract: Unsupervised anomaly localization using variational auto-encoders. *Informatik aktuell*, page 199, 2020.