Pattern Recognition Letters 199 (2026) 113–119

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

FsBAD: Data-efficient feature reconstruction for few-shot brain anomaly detection☆

Hussain Ahmad Madni [a], Hafsa Shujat [b], Axel De Nardin [a], Silvia Zottin [a], Gian Luca Foresti [a,*]

[a] University of Udine, Via delle Scienze, 206, Udine, 33100, Italy
[b] International Islamic University Islamabad, H10, Islamabad, 44000, Pakistan

ARTICLE INFO

Keywords:
Distribution regularization
Feature reconstruction
Few-shot learning
Medical anomaly detection

ABSTRACT

Data efficiency remains a central challenge in brain anomaly detection, where annotated datasets are often scarce. Most existing methods are tailored to single-class settings and show limited ability to generalize. We introduce FsBAD, a feature reconstruction-based approach designed for few-shot brain anomaly detection with minimal supervision. FsBAD reconstructs a nominal version of an anomalous brain scan by leveraging a small set of aligned reference samples. To enhance reconstruction quality, we propose a novel feature alignment strategy that integrates regression with distribution regularization, promoting both semantic accuracy and nominal consistency. While FsBAD is optimized for brain imaging, we evaluate its generalization capabilities on liver and retina datasets. Experiments across all three domains show that FsBAD consistently outperforms state-of-the-art methods in both image-wise classification and pixel-wise anomaly localization, even in extremely low-shot (2- to 15-shot) settings. This demonstrates FsBAD's potential as a scalable, data-efficient solution for brain anomaly detection and its robustness across medical imaging tasks.

1. Introduction

Brain anomaly detection in MRI is critical for early diagnosis and treatment, yet remains particularly challenging due to the subtle nature and variability of anomalies across individuals and modalities. Traditionally, anomaly detection in computer vision has found success in industrial inspection tasks, where anomalies typically appear in well-structured and repetitive environments [1–3]. However, translating these techniques to the medical domain, especially to brain MRI, introduces unique challenges due to the scarcity of annotated data, high intra-class variability, and ethical constraints related to privacy, data sharing, and labeling [4,5].

Most state-of-the-art anomaly detection methods rely on supervised learning paradigms similar to traditional methods [6–12], which require large volumes of well-annotated nominal data and, ideally, a diverse set of labeled anomalies. Unfortunately, such assumptions do not hold in medical scenarios, where anomalous cases are inherently rare, expensive to annotate, and sensitive in nature. Recent research has therefore shifted towards unsupervised or weakly supervised strategies, attempting to model only the distribution of healthy data and detect deviations from it [13–16]. Among these, reconstruction-based approaches using generative models (e.g., Autoencoders) aim to learn the nominal distribution through heavy training procedures, while feature-matching approaches [17,18] rely on comparing a test sample's features with a memory bank of nominal samples. Both strategies are data-hungry and lack scalability across multiple classes or imaging domains.

In this work, we tackle these limitations and introduce FsBAD, a training-free and data-efficient framework for brain anomaly detection. Our method is designed for real-world settings where only a few nominal samples are available. FsBAD reconstructs an anomaly-free version of a given query sample using a linear transformation at the feature level, based on a memory bank of few-shot nominal features. Unlike traditional reconstruction approaches, our method aligns the low-frequency (i.e., structural) content of the reconstructed image with the query, while preserving high-frequency signals that reveal anomalies. To ensure the reconstructed image conforms to the distribution of nominal samples while remaining visually close to the query image, we introduce a distribution regularization term based on classical ridge regression. This allows us to fit a Gaussian model to the nominal feature space and constrain the reconstruction accordingly. The resulting optimization problem admits a closed-form solution, enabling

☆ This article is part of a Special issue entitled: 'HCBIA' published in Pattern Recognition Letters.

* Corresponding author.
E-mail address: gianluca.foresti@uniud.it (G.L. Foresti).

https://doi.org/10.1016/j.patrec.2025.11.016
Received 9 June 2025; Received in revised form 16 September 2025; Accepted 6 November 2025
Available online 7 November 2025

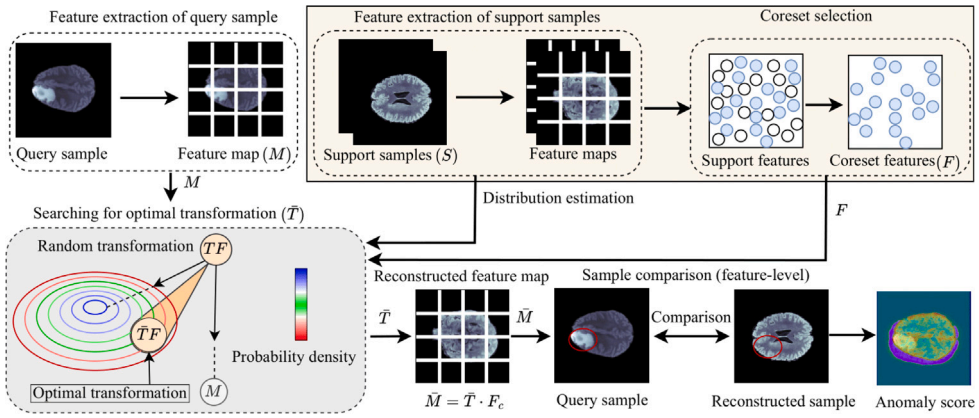

Fig. 1. Overall architecture of the proposed method. A pre-trained encoder is used to extract feature maps of support samples and each given query sample. A support feature pool is created by aggregating all the features obtained from support images. Then, down-sampling is performed on this feature pool using coreset selection F to improve the inference speed and decrease data redundancy. The query feature map M and the coreset F become input to the proposed distribution regularization with regression, as shown in the gray region. The regression is used to obtain the optimal transformation $\bar{T}$ between the query feature map M and support features F to reconstruct the sample $\bar{T}F$, represented as $\bar{M}$. The newly reconstructed sample $\bar{M}$ retains the properties of the nominal sample while sharing the similarity with the original query sample M. Finally, the comparison and alignment between them (i.e., M and $\bar{M}$) is performed to obtain the anomaly detection.

fast computation and facilitating the deployment of FsBAD in resource-constrained clinical environments. While FsBAD is designed with brain MRI in mind, we evaluate its adaptability to other medical imaging tasks including liver CT and retinal fundus images as a measure of its generalization.

We evaluate FsBAD on three different medical datasets, brain MRI [19–21], liver CT [22,23], and retinal fundus images [24], demonstrating strong performance across various few-shot settings. In particular, FsBAD achieves state-of-the-art AUROC scores in brain anomaly detection, confirming its robustness and generalization capabilities.

2. Related works

Medical anomaly detection has garnered significant attention in the computer vision community and can be broadly classified into two categories: feature-matching and reconstruction-based methods. The traditional existing reconstruction-based methods perform classification of anomalous samples using either a comparison between the query sample and the reconstructed nominal sample through the Generative Adversarial Network [25–27], or a reconstruction loss [27–29]. However, feature-matching-based methods exploit pretrained features. For example, Real-IAD [30], Dinomaly [31], PatchCore [18] and PaDiM [17] analyze anomaly detection using the MVTec [1] dataset in an unsupervised manner. Similarly, KNN-based pixel and image-wise anomaly detection [32], and feature memory comprising various feature hierarchies have been used in SPADE [33]. However, both types of models are highly dependent on large data in real-world applications.

In the medical domain, the limited availability of annotated data requires few-shot anomaly detection methods. Recently, different few-shot anomaly detection methods have been introduced [34]. For example, DiffNet [35] uses convolutional neural networks to extract features for density estimation using a normalized flow that helps estimate the distribution with a few supporting samples. In TDG [36], a hierarchical generative model has been proposed that uses every support sample to retain the multi-scale patch distribution. Another method, RegAD [37], has introduced a new perspective of few-shot anomaly detection using a meta-learning design by sharing a common model learning among multiple categories. Although the inference performance is remarkable in a few-shot setting, the migration of the model in a real-world environment is difficult due to the complex training process and enterprise applications [38].

3. Methodology

In this section, we provide a detailed overview of the proposed FsBAD method for the detection of medical anomalies in brain MRI. The proposed method introduces an efficient and scalable approach to few-shot medical anomaly detection leveraging feature reconstruction with distribution regularization. Specifically, FsBAD is composed of three key elements: (1) a feature extraction with a coreset selection module, (2) a reconstruction step, guided by a distribution regularization mechanism based on ridge regression, and (3) the anomaly detection module, which compares the reconstructed and the original samples. The overall architecture of the proposed method is shown in Fig. 1.

For the formal definition of the classification task in medical anomaly detection, we follow the few-shot learning configuration of s-shot and c-way, in which c categories are given, represented as $F = \cup_{i=1}^{c} N_i$, where N_i contains s number of nominal samples taken from a category c_i. During inference, nominal or anomalous samples are given as the query sample. The model predicts whether the query sample is nominal or anomalous at the image level and pixel level. Note that we use three different datasets from different medical domains used as different categories in the experiment settings. The detailed procedural methodology is outlined in Algorithm 1.

3.1. Feature extraction of support samples and coreset selection

Feature extraction that uses pre-trained CNN architectures has proven effective in various applications, including anomaly detection [17,18]. This approach significantly reduces computational overhead while maintaining high detection accuracy, eliminating the need to train deep models from scratch. This observation leads to the first step of the proposed method. Following the existing methods, a pre-trained CNN (i.e., ResNet-50) is used to extract hierarchical feature representations from nominal support samples, followed by a coreset selection process that optimally downsamples features to enhance computational efficiency while preserving representative information. If there $x_i \in X$ is an image in the dataset X, θ is a pre-trained model, then the hierarchy level j is represented as $\theta_{i,j}$, where $j \in 1, 2, 3, 4$. We extract intermediate-level features $j \in 2, 3$ from the network hierarchy, as these layers capture both spatial and semantic representations critical for anomaly detection. The resulting feature maps, with dimensions $H \times W \times D$, are transformed into an embedding space, where H and W denote spatial dimensions and D represents the depth of the feature

Algorithm 1: FsBAD: Data-Efficient Feature Reconstruction for Few-shot Brain Anomaly Detection

Input: Support set $S = \{x_1, x_2, \ldots, x_X\}$ (normal samples), query image $\bar{M}$, feature extractor $f(\cdot)$, ridge parameter α, threshold ζ

Output: Anomaly score and decision for query Image $\bar{M}$

Feature Extraction of Support Samples:
Extract support features: $F \leftarrow [f(x_1), f(x_2), \ldots, f(x_X)]$
Extract query features: $M \leftarrow f(M)$

Ridge Regression and Regularization:
Compute transformation via ridge regression:
$\bar{T} = \arg\min_t \|M - TF\|^2 + \alpha\|T\|^2$
Closed-form solution: $\bar{T} = MF^\tau(F^\tau F)^{-1}$
Reconstruct query: $\bar{M} \leftarrow \bar{T} \cdot F_c$
$\bar{M} \leftarrow$ Regularize$(\bar{M}, F)$

Anomaly Scoring:
score$(M) \leftarrow \|M - \bar{M}\|^2$

Decision:
if *score*$(M) > \zeta$ **then**
  return "Anomaly", score(M);
else
  return "Normal", score(M);

or the number of channels. We reshape the feature map from $H \times W \times D$ (i.e., three dimensions) to $Z \times D$ (i.e., two dimensions) for easy computation, where $Z = H \times W$. This reshaped feature map represents the embedding vectors or patch features. Then we concatenate the features of the support samples s with the dimensions $s \times Z \times D$ to achieve a support feature pool using the s-shot setting. We obtain a coreset $F \in D^{c \times d}$ of features by the coreset sampling operation, where the S-greedy algorithm is used to calculate the number of cluster centers s. To balance detection performance and computational complexity, the coverage of the coreset F is on a smaller scale, despite that it is comparable to the source support feature pools. On the other hand, the patch-level feature map M of the query sample is also extracted using the same pre-trained model. Fig. 1 shows the overall process of extracting features and selecting the coreset in the yellowish region.

3.2. Distribution regularization with regression

To reconstruct the nominal counterpart of an anomalous query sample, we introduce a transformation learning strategy using ridge regression. Our approach integrates a novel distribution regularization term, ensuring that the reconstructed features remain within the nominal feature space while preserving the original sample's anomaly-related characteristics. This ensures that the reconstructed sample remains visually and statistically consistent with normal anatomical structures while preserving the anomalous regions for effective detection. An optimal transformation weight T is searched to rebuild the query feature map M by transforming the coreset F, where $TF \approx M$ and TF have most of the characteristics of nominal samples. While finding optimal weight T solves the problem of linear least-squares [39]. Using ridge regression, the optimization is mathematically formulated as follows:

$$\bar{T} = \arg\min_t \|M - TF\|^2 + \alpha\|T\|^2, \tag{1}$$

where $\|.\|$ is the Frobenius norm, $\|T\|^2$ is the penalty term, and α represents the weight of the penalty term. In ridge regression, the penalty term controls the linearity when under- or over-constrained. However, the reconstructed sample deviates from the distribution of nominal samples as a result of the transformation derived from the ridge regression. Thus, following the transformation, the result TF is much closer to M in such a way that the features of the defect region are recreated similarly to that of M. The transformed query image whose defect regions are reconstructed has most features similar to those of the original query image. In ridge regression, to replace the original penalty term, we propose a distributional regularization term that improves the reconstruction result TF in terms of normality. We use an isotropic Gaussian density function $f\left(X \parallel \sigma, \sum\right)$ to model the distribution of normal feature maps, where support set feature maps are used to estimate σ and $\sum$, called mean and covariance, respectively. Following this, the regularization term is denoted as $\|f(TF) - f(\sigma)\|$, and the proposed regression is formulated as follows:

$$\bar{T} = \arg\min_t \|M - TF\|^2 + \alpha\|f(TF) - f(\sigma)\|, \tag{2}$$

where α is the coefficient controlling the regularization contribution. In Eq. (2), the first term $\|M - TF\|^2$ naturally drags TF to M (i.e., original query feature maps) during the search for optimal transformation T. However, in our proposed regularization $\|f(TF) - f(\sigma)\|$, the constructed TF is forced to remain closer to the distribution center in the high probability density region. The distribution regularization is visualized in the gray area of Fig. 1.

3.3. Closed-form solution and optimal transformation

In Eq. (2), the presence of a density function complicates the derivation of a closed-form solution, leading to increased computational overhead. Therefore, we propose an alternative of Eq. (2) that is approximate to it using Taylor's theorem on the distribution density function f. For $f(x)$ at γ, Taylor's expansion is represented as follows:

$$f(x) \approx f(\gamma) + \nabla f(\gamma)(x - \gamma) + \frac{1}{2}(x - \gamma)^T Hf(\gamma)(x - \gamma), \tag{3}$$

where $Hf(\gamma)$ is the Hessian matrix of $f(x)$ at γ. To approximate the density function, we use first- and second-order terms from Taylor's expansion, in which the first-order term is 0 at σ. Thus, the density function $f(TF)$ at σ is expanded using Eq. (3) as follows:

$$f(TF) \approx f(\sigma) + \nabla f(\sigma)(TF - \sigma) + \frac{1}{2}(TF - \sigma)^{\hat{T}} Hf(\sigma)(TF - \sigma) \tag{4}$$

We assume that the data distribution is isotropic Gaussian, so $\sum = \mu^2 I$ is its covariance matrix, where μ^2 the variance will be across all dimensions. Next, we take $\nabla f(\sigma)(TF - \sigma)$ as 0 and $Hf(\sigma)$ as $\sum^{-1}$, which is the inverted covariance. Finally, moving $f(\sigma)$ to the left, we have:

$$\begin{aligned} f(TF) - f(\sigma) &\approx \frac{1}{2}(TF - \sigma)^{\hat{T}} \sum^{-1}(TF - \sigma) \\ &= \frac{1}{2\mu^2}(TF - \sigma)^{\hat{T}} I(TF - \sigma) \\ &= \frac{1}{2\mu^2}\|TF - \sigma\|^2, \end{aligned} \tag{5}$$

where I denotes the identity matrix. To reformulate the approximate version of Eq. (2), $\frac{1}{2\mu^2}$ is merged into α, given as follows:

$$\bar{T} = \arg\min_t \|M - TF\|^2 + \alpha\|TF - \sigma\|^2 \tag{6}$$

If the optimization goal of least square error (LSE) is $T' = \arg\min_t \|M - TF\|^2$, then the closed-form solution is expressed as follows:

$$T' = MF^\tau(F^\tau F)^{-1} \tag{7}$$

If M_* and F_* are assumed as column-wise augmented versions of M and F, respectively, then Eq. (6) is transformed to fit into Eq. (7) by variable substitution as follows:

$$F_* = [F\sqrt{\alpha}F], M_* = [M\sqrt{\alpha}\sigma] \tag{8}$$

where the coefficient α of the regularization terms is always positive. As the equality holds in $\|M_* - TF_*\|^2 = \|M - TF\|^2 + \alpha\|\sigma - TF\|^2$, then Eq. (6) is reformulated as follows:

$$\bar{T} = \arg\min_t \|M_* - TF_*\|^2 \tag{9}$$

Finally, the optimization solution by concluding Eq. (7) is represented as follows:

$$\bar{T} = M_* F_*^{\tau}(F_*^{\tau} F_*)^{-1} = (M F^{\tau} + \alpha \cdot \sigma F^{\tau})(F^{\tau} F + \alpha \cdot F F^{\tau})^{-1} \quad (10)$$

3.4. Feature reconstruction and anomaly detection

Once the query feature map $\bar{M}$ is reconstructed, its alignment with the original sample M is performed to compute the anomaly segmentation map for pixels and the anomaly score for images. Thus, the anomaly score A_i for a query image with the pixel i is computed as follows:

$$A_i = \|M_i - \bar{M}_i\|^2, i \in \mathcal{P}, \quad (11)$$

where $\mathcal{P}$ represents the pixel indices of M, and M_i is a feature vector for a pixel i. We compute the maximum distance d between pixels in the image to calculate the image-level score for classification. It is mathematically expressed as follows:

$$d = max\{\|M_i - \bar{M}^2\| \mid i \in \mathcal{P}\} \quad (12)$$

In this step, the segmentation map is up-scaled to align with the original input resolution with bi-linear interpolation as used in [17,18]. Finally, a smooth segmentation map is created using a Gaussian kernel of width 4.

4. Experiments

4.1. Datasets

We evaluate FsBAD on two brain MRI datasets, as well as two additional datasets focusing on liver CT and retinal RESC to test the generalization capabilities of the proposed model. In particular, we relied on the brain MRI dataset extracted from the 2022 Medical Segmentation Decathlon [40], as well as the one provided by the BMAD benchmark [19] to assess our model capabilities on Brain MRI data, Furthermore, we selected the two remaining BMAD [41] datasets, each representing distinct medical imaging modalities, namely liver CT (i.e., organ anomalies) [22,23], and retinal RESC (i.e., vascular abnormalities) [24].

For each dataset, we use 375 training samples and consider the anomaly detection task in a one-class setting, where the test sets include both nominal and anomalous samples. Specifically, the brain MRI dataset includes 640 nominal and 3075 anomalous test images, the liver CT dataset includes 833 nominal and 660 anomalous test images, and the retinal RESC dataset comprises 1041 nominal and 764 anomalous test images.

Fig. 2 illustrates representative nominal and anomalous examples from each dataset. The top row shows nominal samples, while the bottom row presents anomalous cases, with the regions of interest highlighted using red circles.

Furthermore, to assess the generalization capability of our approach beyond the BMAD datasets, we also evaluate FsBAD on a brain tumor dataset introduced in [40].

4.2. Technical details

All the experiments are performed using the PyTorch module on a Windows-based system assembled with an NVIDIA GeForce RTX 8090 GPU with 24 GB of RAM and an i7-8700 CPU with memory of 32 GB. We use pre-trained ResNet-50 [42] as the backbone network for feature extraction, where feature maps from two middle-level layers are extracted. For both the training and testing phases, we reshape each query sample to 368 × 368. Moreover, the sampling rate in coreset selection is set to 0.15, and the coefficient α of the regularization term is 0.3 for all experiments.

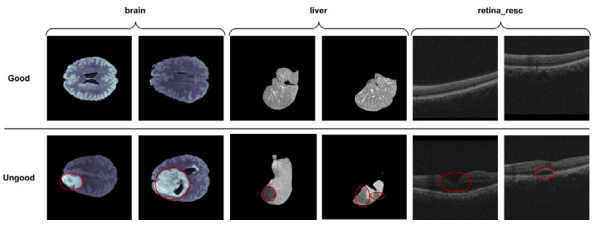

Fig. 2. Some samples from the brain, liver, and retina_resc datasets. The upper row shows the nominal samples from the given datasets, while the lower row shows the anomalous samples from the datasets where anomalies are represented by red canvas.

4.3. Results and comparison with state-of-the-art methods

Anomaly maps representing the feature reconstruction capabilities of the proposed method are shown in Fig. 3. We can observe how the approach effectively detects anomalies by reconstructing the features of the query samples. These reconstructed samples, which retain the essential properties of the original samples, are decoded using a trained decoder. The reconstructed and original query samples are then compared to calculate the anomaly score.

We benchmark FsBAD against state-of-the-art few-shot anomaly detection methods, including DiffNet [35], PatchCore [43], RegAD [37], TDG [36], PaDiM [17], and FastRcon [44]. As evaluation metrics, we rely on the popular Area Under Receiver Operator Curve (AUROC) metric for the pixel-wise anomaly localization and image-wise anomaly detection, ensuring a comprehensive assessment. In Table 1, a comparison between our approach and the other state-of-the-art models is provided for the individual classes that characterize the four selected datasets. As we can see, FsBAD consistently shows the highest performance in both pixel- and image-wise AUROC across all evaluation settings compared to existing methods.

Furthermore, we evaluate the proposed method in terms of training and testing time on a brain tumor dataset introduced in [40] using the same evaluation metrics. We compare the performance of the proposed method with other baseline methods based on both pixel- and image-wise anomaly prediction, as summarized in Table 2. These results further substantiate the effectiveness and superiority of the proposed method over existing approaches.

4.4. Ablation study

In this section, we investigate the impact of the core components of the proposed approach. Specifically, we verify how hyper-parameters such as data-resolution, coefficient of regularization term, and sampling ratio in coreset selection affect the final performance of the framework. A visual summary of the performed ablation study is shown in Fig. 4.

4.4.1. Speed performance

In the proposed FsBAD, the training stage is comprised of feature extraction, subsampling of features, and computation of the transformation matrix. The speed comparison for both the training and testing phases is given in Table 2. FsBAD achieves the fastest training and inference times among all evaluated methods, demonstrating its computational efficiency for real-world deployment.

4.4.2. Impact of image resolution

We reshape each image into a resolution $r \times r$ regardless of its original resolution. The resolution of the image affects the feature points. We present the AUROC performance against image resolution r in Fig. 4, where FsBAD performs well when image resolution is equal to 368. The performance of the proposed method is higher for both image-level and pixel-levels at this image resolution.

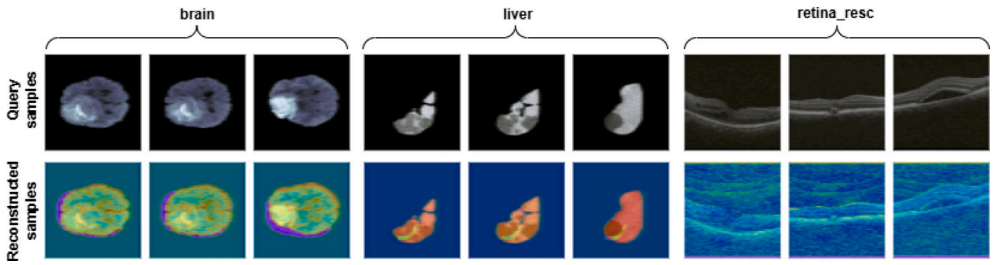

Fig. 3. Visualization of the feature-reconstruction ability of the proposed method. The upper row contains original query samples, while the lower row contains the samples of reconstructed feature maps.

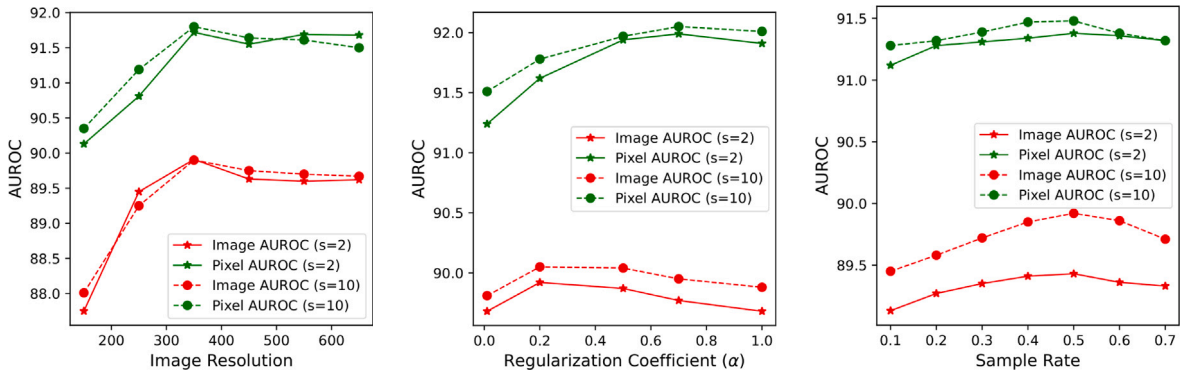

Fig. 4. Impact of main hyperparameters including image resolution r, regularization coefficient α, and sample rate on image-wise and pixel-wise AUROC. (left) AUROC performance with different image resolution r. (center) AUROC performance against different value of coefficient α of regularization term. (right) AUROC performance with different sampling rate.

Table 1
Comparisons of FsBAD with existing methods on brain MRI, liver CT, and retina_RESC in terms of AUROC computed for both pixel-wise localization and image-wise anomaly detection. s represents the number of support samples (i.e., few-shot) in a few-shot setting. The highest performance is shown in bold text.

Dataset	s	Image-wise AUROC									Pixel-wise AUROC						
		PatchCore [43]	DiffNet [35]	FastRcon [44]	PaDiM [17]	TDG [36]	RegAD [37]	Real-IAD [30]	Dinomaly [31]	FsBAD (Our)	FastRcon [44]	RegAD [37]	PatchCore [43]	PaDiM [17]	Real-IAD [30]	Dinomaly [31]	FsBAD (Our)
Brain MRI [19]	2	78.45	82.63	88.82	71.87	70.87	85.64	85.00	89.30	**90.94**	92.13	91.57	92.41	89.09	86.30	92.36	**94.18**
	5	79.52	82.97	89.71	72.45	71.38	86.74	85.28	89.57	**91.09**	92.92	92.18	93.48	90.80	87.58	92.68	**95.48**
	10	81.85	83.36	90.28	73.28	72.07	87.22	86.79	89.88	**91.83**	93.55	92.98	94.29	91.72	89.22	93.70	**96.07**
	15	82.57	83.92	91.42	74.32	72.67	88.33	88.14	90.15	**92.83**	94.67	93.44	95.30	92.30	90.29	94.38	**97.83**
Brain tumor [40]	2	78.35	82.15	86.88	75.06	75.52	87.39	87.38	89.41	**92.11**	92.13	88.27	88.70	88.95	88.80	90.83	**94.19**
	5	79.12	82.48	87.25	75.18	76.15	87.95	87.53	89.65	**92.24**	92.82	88.68	89.25	89.58	89.25	91.26	**94.78**
	10	79.85	83.23	87.93	75.37	76.64	88.46	88.08	90.22	**92.78**	93.65	88.91	89.90	90.49	89.45	91.89	**95.29**
	15	80.12	83.42	88.68	75.58	77.25	88.78	88.68	91.70	**93.22**	94.13	89.23	90.52	91.02	89.46	92.40	**95.81**
Liver CT [22] [23]	2	72.68	83.32	88.35	75.11	70.12	86.95	86.45	88.88	**91.86**	90.17	88.73	88.76	87.27	87.62	89.54	**92.17**
	5	73.77	84.77	88.98	75.94	71.26	87.27	86.77	89.32	**92.21**	90.85	89.44	89.69	88.49	87.71	90.27	**92.79**
	10	75.02	85.90	89.53	76.52	72.27	88.06	87.48	90.18	**92.44**	91.25	90.19	90.22	89.57	88.42	91.25	**93.58**
	15	75.49	87.11	90.78	77.18	73.17	88.76	87.98	90.85	**93.19**	92.12	90.96	91.41	90.72	88.72	92.08	**95.11**
Retina_RESC [24]	2	75.88	72.15	82.16	72.14	72.68	81.14	82.58	83.33	**85.13**	84.66	82.24	83.24	83.87	84.72	84.10	**86.79**
	5	76.12	74.28	82.68	73.42	72.44	81.49	82.75	83.68	**85.56**	85.79	82.68	84.12	85.59	85.18	84.57	**87.76**
	10	77.45	75.65	83.24	74.15	72.82	81.78	83.02	84.13	**86.91**	86.64	83.10	84.82	86.48	85.60	85.06	**89.24**
	15	78.26	77.22	84.09	75.46	73.19	82.08	83.20	84.60	**88.78**	87.23	83.81	85.54	87.94	86.02	85.70	**89.91**

Table 2
Speed comparison of both training and testing stages in terms of time in seconds. Additional training on aggregated features is required by the RegAD. The aggregated testing speed is measured as average time per query sample. The minimum time for training and testing time is shown as bold text.

Method	Support samples (s)	Train Time	Test Time
PatchCore [43]	–	1.05	0.86
RegAD [37]	2	865.15	0.06
RegAD [37]	10	2449.87	0.06
FastRecon [44]	2	0.71	**0.05**
FastRecon [44]	10	2.12	**0.05**
FsBAD (Ours)	2	**0.58**	**0.05**
FsBAD (Ours)	10	**1.96**	**0.05**

4.4.3. Impact of regularization coefficient

As shown in Fig. 4, the proposed FsBAD achieves better AUROC for both pixel-wise localization and image-wise anomaly detection. This indicates the importance of distributed regularization and the performance of the proposed approach against different coefficients α of the regularization term for 2-shot and 10-shot settings. The AUROC performance is at the peak value when α is set to 0.5 and 0.7 for 2-shot and 10-shot respectively.

4.4.4. Impact of sample rate

Fig. 4 represents the impact of AUROC performance against different sampling rates. In the support feature pool, the sample rate affects the number of feature points. The best AUROC performance is reflected when the sampling rate is set to 0.50.

4.5. Limitations

Although FsBAD achieves strong performance in few-shot anomaly detection, it has some limitations. The method depends on pretrained feature extractors and representative support samples, which may not fully capture diverse normal patterns in medical data. It can also be sensitive to hyperparameters and may struggle with subtle or highly complex anomalies. Addressing these issues through domain-adaptive features, more robust optimization, and the use of contextual or temporal cues presents promising directions for future work.

5. Conclusion

In this paper, we propose FsBAD, a novel few-shot medical anomaly detection method that reconstructs the nominal version of a query sample using a minimal set of nominal reference features. By integrating distribution regularization with ridge regression, FsBAD ensures that the reconstructed sample retains nominal characteristics while effectively highlighting anomalies. Our approach significantly reduces the reliance on large labeled datasets, making it a scalable and computationally efficient solution for brain anomaly detection, with proven ability to generalize across domains such as liver and retina imaging. Comparison with existing state-of-the-art models shows improved performance in all experimental settings considered, demonstrating its robustness to different training scenarios, data amounts, and imaging modalities.

CRediT authorship contribution statement

Hussain Ahmad Madni: Writing – original draft, Methodology, Formal analysis, Data curation, Conceptualization. **Hafsa Shujat:** Writing – review & editing, Validation, Investigation. **Axel De Nardin:** Writing – review & editing, Validation, Formal analysis. **Silvia Zottin:** Writing – review & editing, Validation, Formal analysis. **Gian Luca Foresti:** Supervision, Resources, Project administration, Investigation, Funding acquisition.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Gian Luca Foresti reports was provided by University of Udine. Gian Luca Foresti reports a relationship with University of Udine that includes: employment, funding grants, and non-financial support. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This paper was partially supported by: FVG project “Supporting the diagnosis of rare diseases through AI” (2023–26) (Project A with CUP: F53C22001770002 and Project B with CUP F53C22001780002); “Piano Stretegico Dipartimentale on Artificial Intelligence” (2022–25) project at the University of Udine; Piano Nazionale di Ripresa e Resilienza (PNRR) DD 3277 del 30 dicembre 2021 (PNRR Missione 4, Componente 2, Investimento 1.5) - iNEST.

Data availability

Data will be made available on request.

References

[1] P. Bergmann, M. Fauser, D. Sattlegger, C. Steger, MVTec AD–A comprehensive real-world dataset for unsupervised anomaly detection, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 9592–9600.

[2] T. Matsubara, R. Tachibana, K. Uehara, Anomaly machine component detection by deep generative model with unregularized score, in: 2018 International Joint Conference on Neural Networks, IJCNN, IEEE, 2018, pp. 1–8.

[3] A. De Nardin, P. Mishra, G.L. Foresti, C. Piciarelli, Masked transformer for image anomaly localization, Int. J. Neural Syst. 32 (07) (2022) 2250030.

[4] D. Veritti, L. Rubinato, V. Sarao, A. De Nardin, G.L. Foresti, P. Lanzetta, Behind the mask: a critical perspective on the ethical, moral, and legal implications of AI in ophthalmology, Graefe’s Arch. Clin. Exp. Ophthalmol. 262 (3) (2024) 975–982.

[5] H.A. Madni, R.M. Umer, S. Zottin, C. Marr, G.L. Foresti, FL-W3S: Cross-domain federated learning for weakly supervised semantic segmentation of white blood cells, Int. J. Med. Informatics 195 (2025) 105806, http://dx.doi.org/10.1016/j.ijmedinf.2025.105806.

[6] V. Sarao, D. Veritti, A. De Nardin, M. Misciagna, G. Foresti, P. Lanzetta, Explainable artificial intelligence model for the detection of geographic atrophy using colour retinal photographs, BMJ Open Ophthalmol. 8 (1) (2023) e001411.

[7] S.H. Kim, S. Schramm, J. Wihl, P. Raffler, M. Tahedl, J. Canisius, I. Luiken, L. Endrös, S. Reischl, A. Marka, et al., Boosting LLM-assisted diagnosis: 10-minute LLM tutorial elevates radiology residents’ performance in brain MRI interpretation, Neuroradiology (2025) 1–13.

[8] Y. Zhou, L. Song, J. Shen, Improving medical large vision-language models with abnormal-aware feedback, 2025, arXiv preprint arXiv:2501.01377.

[9] C.Y. Li, K.J. Chang, C.F. Yang, H.Y. Wu, W. Chen, H. Bansal, L. Chen, Y.P. Yang, Y.C. Chen, S.P. Chen, et al., Towards a holistic framework for multimodal LLM in 3D brain CT radiology report generation, Nat. Commun. 16 (1) (2025) 2258.

[10] S.H. Kim, J. Wihl, S. Schramm, C. Berberich, E. Rosenkranz, L. Schmitzer, K. Serguen, C. Klenk, N. Lenhart, C. Zimmer, et al., Human-AI collaboration in large language model-assisted brain MRI differential diagnosis: a usability study, Eur. Radiol. (2025) 1–12.

[11] Y. Zhou, L. Song, J. Shen, MAM: Modular multi-agent framework for multi-modal medical diagnosis via role-specialized collaboration, 2025, arXiv preprint arXiv:2506.19835.

[12] J. Yang, L. Fan, Q. Sun, X. Gao, Bidirectional two-dimensional supervised multiset canonical correlation analysis for multi-view feature extraction, Pattern Recognit. Lett. (2025).

[13] B. Zong, Q. Song, M.R. Min, W. Cheng, C. Lumezanu, D. Cho, H. Chen, Deep autoencoding gaussian mixture model for unsupervised anomaly detection, in: International Conference on Learning Representations, 2018, p. null.

[14] D. Gong, L. Liu, V. Le, B. Saha, M.R. Mansour, S. Venkatesh, A.v.d. Hengel, Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1705–1714.

[15] H.A. Madni, H. Shujat, A. De Nardin, S. Zottin, G.L. Foresti, Unsupervised brain MRI anomaly detection via inter-realization channels, Int. J. Neural Syst. (ja) null, http://dx.doi.org/10.1142/S0129065725500479.

[16] Z. Guo, T.J.J. Wang, S. Pehlivan, A. Radman, M. Cao, J. Laaksonen, Prompt-based weakly-supervised vision-language pre-training, Pattern Recognit. Lett. (2025).

[17] T. Defard, A. Setkov, A. Loesch, R. Audigier, Padim: a patch distribution modeling framework for anomaly detection and localization, in: International Conference on Pattern Recognition, Springer, 2021, pp. 475–489.

[18] J. Yi, S. Yoon, Patch svdd: Patch-level svdd for anomaly detection and segmentation, in: Proceedings of the Asian Conference on Computer Vision, 2020.

[19] U. Baid, S. Ghodasara, S. Mohan, M. Bilello, E. Calabrese, E. Colak, K. Farahani, J. Kalpathy-Cramer, F.C. Kitamura, S. Pati, et al., The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification, 2021, arXiv preprint arXiv:2107.02314.

[20] S. Bakas, H. Akbari, A. Sotiras, M. Bilello, M. Rozycki, J.S. Kirby, J.B. Freymann, K. Farahani, C. Davatzikos, Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features, Sci. Data 4 (1) (2017) 1–13.

[21] B.H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al., The multimodal brain tumor image segmentation benchmark (BRATS), IEEE Trans. Med. Imaging 34 (10) (2014) 1993–2024.

[22] B. Landman, Z. Xu, J. Igelsias, M. Styner, T. Langerak, A. Klein, Miccai multi-atlas labeling beyond the cranial vault–workshop and challenge, in: Proc. MICCAI Multi-Atlas Labeling beyond Cranial Vault—Workshop Challenge, vol. 5, 2015, p. 12.

[23] P. Bilic, P. Christ, H.B. Li, E. Vorontsov, A. Ben-Cohen, G. Kaissis, A. Szeskin, C. Jacobs, G.E.H. Mamani, G. Chartrand, et al., The liver tumor segmentation benchmark (lits), Med. Image Anal. 84 (2023) 102680.

[24] J. Hu, Y. Chen, Z. Yi, Automated segmentation of macular edema in OCT using deep neural networks, Med. Image Anal. 55 (2019) 216–227.

[25] V. Zavrtanik, M. Kristan, D. Skočaj, Reconstruction by inpainting for visual anomaly detection, Pattern Recognit. 112 (2021) 107706.
[26] T. Schlegl, P. Seeböck, S.M. Waldstein, U. Schmidt-Erfurth, G. Langs, Unsupervised anomaly detection with generative adversarial networks to guide marker discovery, in: International Conference on Information Processing in Medical Imaging, Springer, 2017, pp. 146–157.
[27] J. An, S. Cho, Variational autoencoder based anomaly detection using reconstruction probability, Spec. Lect. IE 2 (1) (2015) 1–18.
[28] H. Zenati, M. Romain, C.S. Foo, B. Lecouat, V. Chandrasekhar, Adversarially learned anomaly detection, in: 2018 IEEE International Conference on Data Mining, ICDM, IEEE, 2018, pp. 727–736.
[29] C. Zhou, R.C. Paffenroth, Anomaly detection with robust deep autoencoders, in: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2017, pp. 665–674.
[30] C. Wang, W. Zhu, B.B. Gao, Z. Gan, J. Zhang, Z. Gu, S. Qian, M. Chen, L. Ma, Real-iad: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 22883–22892.
[31] J. Guo, S. Lu, W. Zhang, F. Chen, H. Li, H. Liao, Dinomaly: The less is more philosophy in multi-class unsupervised anomaly detection, in: Proceedings of the Computer Vision and Pattern Recognition Conference, 2025, pp. 20405–20415.
[32] E. Eskin, A. Arnold, M. Prerau, L. Portnoy, S. Stolfo, A geometric framework for unsupervised anomaly detection: Detecting intrusions in unlabeled data, Appl. Data Min. Comput. Secur. (2002) 77–101.
[33] N. Cohen, Y. Hoshen, Sub-image anomaly detection with deep pyramid correspondences, 2020, arXiv preprint arXiv:2005.02357.
[34] J. Xiao, R. Wang, C. He, X. Chen, Cross-domain few-shot 3D point cloud semantic segmentation, Pattern Recognit. Lett. (2025).
[35] M. Rudolph, B. Wandt, B. Rosenhahn, Same same but differnet: Semi-supervised defect detection with normalizing flows, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2021, pp. 1907–1916.
[36] S. Sheynin, S. Benaim, L. Wolf, A hierarchical transformation-discriminating generative model for few shot anomaly detection, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 8495–8504.
[37] C. Huang, H. Guan, A. Jiang, Y. Zhang, M. Spratling, Y.F. Wang, Registration based few-shot anomaly detection, in: European Conference on Computer Vision, Springer, 2022, pp. 303–319.
[38] S. Barra, C. Bisogni, M. De Marsico, S. Ricciardi, Visual question answering: Which investigated applications? Pattern Recognit. Lett. 151 (2021) 325–331.
[39] D. Wertheimer, L. Tang, B. Hariharan, Few-shot classification with feature map reconstruction networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 8012–8021.
[40] M. Antonelli, A. Reinke, S. Bakas, K. Farahani, A. Kopp-Schneider, B.A. Landman, G. Litjens, B. Menze, O. Ronneberger, R.M. Summers, et al., The medical segmentation decathlon, Nat. Commun. 13 (1) (2022) 4128.
[41] J. Bao, H. Sun, H. Deng, Y. He, Z. Zhang, X. Li, Bmad: Benchmarks for medical anomaly detection, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 4042–4053.
[42] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
[43] K. Roth, L. Pemula, J. Zepeda, B. Schölkopf, T. Brox, P. Gehler, Towards total recall in industrial anomaly detection, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 14318–14328.
[44] Z. Fang, X. Wang, H. Li, J. Liu, Q. Hu, J. Xiao, Fastrecon: Few-shot industrial anomaly detection via fast feature reconstruction, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 17481–17490.