

PANDA : Perceptually Aware Neural Detection of Anomalies

Jack W. Barker¹, Toby P. Breckon^{1,2}

Department of Computer Science¹ | Department of Engineering²
Durham University, UK

Abstract—Semi-supervised methods of anomaly detection have seen substantial advancement in recent years. Of particular interest are applications of such methods to diverse, real-world anomaly detection problems where anomalous variations can vary from the visually obvious to the very subtle. In this work, we propose a novel fine-grained VAE-GAN architecture trained in a semi-supervised manner in order to detect both visually distinct and subtle anomalies. With the use of a residually connected dual-feature extractor, a fine-grained discriminator and a perceptual loss function, we are able to detect subtle, low inter-class (anomaly vs. normal) variant anomalies with greater detection capability and smaller margins of deviation in AUC value during inference compared to prior work whilst also remaining time-efficient during inference. We achieve state-of-the-art anomaly detection results when compared extensively with prior semi-supervised approaches across a multitude of anomaly detection benchmark tasks including trivial leave-one-out tasks (CIFAR-10 - $AUPRC_{avg}$: 0.91; MNIST - $AUPRC_{avg}$: 0.90) in addition to challenging real-world anomaly detection tasks (plant leaf disease - AUC: 0.776; threat item X-ray - AUC: 0.51), video frame-level anomaly detection (UCSDPed1 - AUC: 0.95) and high frequency texture with object anomalous defect detection (MVTEC - AUC_{avg} : 0.83).

I. INTRODUCTION

Anomaly detection is the task of recognising samples of a given dataset which deviate significantly from established normality and as such, represent unexpected eventualities or outliers in the scope of a given task [1]. Anomaly detection is a challenging task because of the broad range of variational forms which anomalies may present, representing an unbounded (open set) distribution of possible deviations from normality.

The ultimate application challenge of anomaly detection models is effective and efficient applications to critical real-world tasks [2]–[4] with little to no human intervention at training time. However, by contrast, the *modus operandi* of anomaly detection evaluation in the literature is to demonstrate model performance across trivial and unrealistic ‘leave one out’ tasks on general datasets such as MNIST [5] or CIFAR-10 [6] in which one class from the dataset is labelled as anomalous and all other classes as normal. This evaluation methodology is highly unrealistic as not only are these datasets not intended for anomaly detection, but the act of directly comparing the classes present in the datasets in this way is unlikely to occur in a real-world anomaly detection application. Anomalies occurring within real-world problems may be subtle, localised to a small sub-region of the image, exhibit



Fig. 1. **Top:** Leaves from Plant Village [7] featuring visible diseases. **Bottom:** Anomalous instance segmentation masks generated by PANDA for the respective diseased leaves.

high variance and even be the result of subterfuge by an adversary [7]–[10].

Anomaly detection methods have had varying degrees of success across real-world tasks including, but not limited to: retinal diagnosis [11], [12], airport security scanning [8], [13]–[15] and factory line inspection [10]. However, these methods can often attribute their limited success to being domain-specific and are not applied across multiple, diverse (multi-spectral; cross-domain) datasets. Such prior anomaly detection methods are overly focused on more general features to aid in categorising visually obvious anomalies [11], [14] akin to the flawed evaluation methodology of ‘leave-one-out tasks’, meaning they do not perform overly well with visually subtle anomalies in tasks such as [7], [11], [12].

Whilst supervised methods [8], [13] can obtain superior performance across anomaly detection benchmark tasks, often by following a simplistic anomaly detection by classification paradigm with discrete classes, they require large, labeled-datasets for training. These can be both expensive to obtain, unbalanced in nature, and will always struggle to provide sufficient coverage of rare, low-occurrence anomalies given the potential open-ended scope of the anomalous class space. These challenges of training data adequacy could lead to potential adversarial example attacks against such methods [16], [17].

By contrast, semi-supervised methods [12], [14], [15], [18], [19] overcome this issue by learning a close approximation to the true distribution manifold exclusively over the non-

anomalous (normal) data samples [11]. Such techniques use generative methods in order to approximate the distribution [14], [15]. Furthermore, methods such as [11], [12], [15] suffer from slow inference times which can hinder real-world applicability in scenarios where high-throughput processing is required. Methods such as [12], [14] exhibit vastly differing accuracy with each iteration over the same dataset leading to a sparse confidence interval as demonstrated in our experiments, impeding real-world applicability due to unpredictable detection behaviour at inference.

In this work, we propose the Perceptually Aware Neural Detection of Anomalies (PANDA-GAN), a Variational Autoencoder Generative Adversarial Network (VAE-GAN) based architecture to combat the task of detecting subtle fine-grained anomalies present in real-world anomaly detection applications whilst also retaining time-efficiency at inference. PANDA includes three novel proposals: (1) a Fine-Grained Visual Categorisation Discriminator Network (FGVC) to ease the problem of detecting visually subtle, low inter-class variance anomalies present in anomaly detection problems and to provide a harsher critic during training for the GAN generator module; (2) a residually connected dual-feature extractor implementation within our generator module that carries lower-level features in given images forward and combines them residually with higher-level, later features in the architecture; (3) a perceptual loss function based on feature error instead of raw pixel-error [20] that in turn obtains higher-fidelity images, but has not yet been applied to the task of generative anomaly detection. This work represents the first instance of these techniques being jointly applied to semi-supervised anomaly detection.

II. RELATED WORK

Many works have addressed the problem of generative semi-supervised anomaly detection. Initial methods propose the use of Variational Autoencoder (VAE) architectures [21], in which a latent representation z is learned from the image space X through the use of an encoder. A second module (decoder) then maps from z , back to the image space to produce \hat{X} . The encoder and decoder can be trained using the reconstruction error between the original image $x \in X$ and the produced image $\hat{x} \in \hat{X}$. Early implementations of VAE [21], however do not capture the distribution of the data $p(X)$ well due to the oversimplification of the learned prior probability $p(z|X)$. VAE are only capable of learning a uni-modal distribution, which fails to capture complex distributions that are commonplace in real world anomaly detection scenarios.

Generative Adversarial Networks (GAN), first proposed by Goodfellow *et al.* [22], combat this simplification by forcing a Generator, G , to model complex distributions in data from random noise in order to generate representative image samples drawn from this distribution. The learning objective of G is to reduce the confidence of the Discriminator, D , to assign an effective probability to a presented image; whether or not it is original or a generated image. The zero-sum end-game of a GAN is the Nash Equilibrium whereby the Generator network

and Discriminator network have saturated learning such that the probability of the Discriminator to distinguish between real and synthetically generated images converges.

AnoGAN [11] is the first GAN-based, semi-supervised method of anomaly detection. The model is trained only on non-anomalous data to learn the manifold z of normal samples. When an anomaly x_a is processed by the Generator network, it produces a non-anomalous image x'_a . Taking an l_2 reconstruction error will outline anomalies present in the image. Although this method proved that it was possible to use GAN for the task of anomaly detection, the computational performance is incredibly slow hence limiting real world applicability. GANomaly [14] overcomes these issues by training a Generator network and a secondary encoder in order to map the generated samples into a second latent space \hat{z} which is then used to better learn the original latent priors z , mapping between latent values efficiently at the same time as the Generator learns the distribution manifold over data x . Follow-on work, Skip-GANomaly [15], introduces the notion of skip connections into the network architecture in order to preserve image information detail across the encoder-decoder structure, greatly improving performance.

Concurrently, the author of AnoGAN produced the Fast-AnnoGAN [12] method of anomaly detection which is similar to their previous work, but replaces the Deep Convolutional GAN (DCGAN) [23] with a Wasserstein Generative model [24]. It uses the trained Generator from the GAN training to train an encoder which maps images to the latent space. This enables the overall Fast-AnnoGAN architecture to avoid the computationally expensive operation of obtaining a latent representation at inference. A similar approach was adopted by Houssam *et al.* [25] which is based on the BiGAN architecture [26]. They use an approach similar to the one indicated in [26] to solve the optimisation problem $\min_{G,E} \max_D V(D, G, E)$ where the features of X are learned by the network E to produce the pair $(x, E(x))$. The Generator network learns the pair of $(z, G(z))$ from the real features of X [27]. This simultaneous learning of the pairs forces the network to learn the mapping from not only image data to latent space, but from latent space back to image data.

Although GANs have risen to prominence and gained significant results, producing high-fidelity images, they suffer from volatile training issues such as mode collapse and instability, leading to over-fitting during training. VAE-GAN [28] are VAE which are trained in an adversarial manner. The first instance of these being utilised for anomaly detection is Baur *et al.* [18] which utilises a common VAE architecture, but applies a Discriminator network in order to determine whether the images are real, or reconstructed. Similar work was performed in ADAE [19] which uses a dual, parallel model in which the primary VAE is the Generator and the second VAE is the Discriminator. As well as reconstructing the images from the latent representation, the authors also compute the error between the data distributions between the dual networks in order to gain a closer fit for the manifold over real examples.

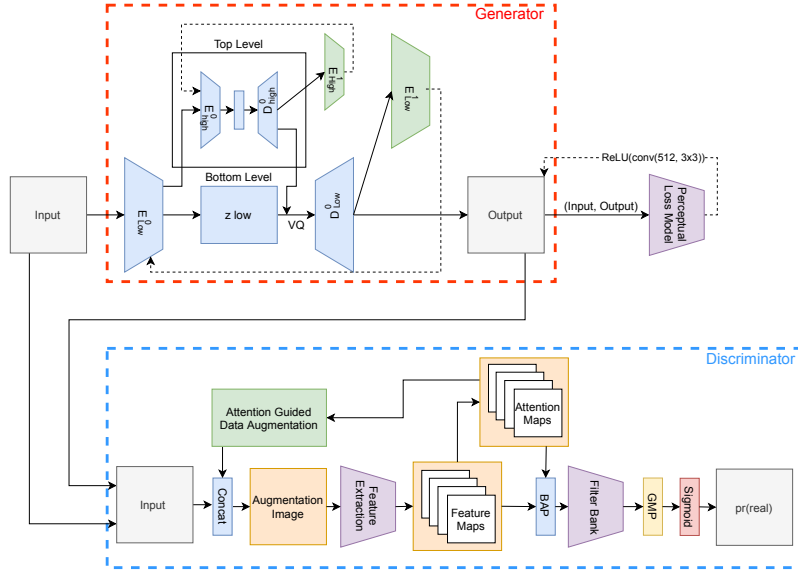


Fig. 2. Proposed model architecture featuring our PANDA-GAN architecture with the Generator network (upper) and the Discriminator network (lower).

III. APPROACH

Our proposed method applies a unique VAE-GAN architecture to the task of anomaly detection. Our Generator network (Section III-A) uses skip-connected dual-feature extractor encoders ($\{E_{high}^0, E_{low}^0\}$) to extract both a top and bottom level latent space which residually combines early, low-level information with high-level features further in the architecture. This network is trained adversarially with a fine-grained Discriminator network (Section III-B) which is optimised to assign a true probability that a presented image conforms to the distribution of normality and is not synthetically generated and used to detect subtle discriminating features from input images during inference. A visual overview of our approach is shown in Figure 2.

A. Generator Network

Our Generator VAE model is trained adversarially to create a VAE-GAN which does not exhibit mode collapse or vanishing gradient training difficulties present in traditional GAN architectures [22]. Similar to a U-net model architecture, input images x are first fed into a low-level encoder E_{low}^0 which maps to a latent representation $z_{low} = p_{\theta}(z_{low}|x) \sim N(\bar{x}, \sigma)$. We then further encode z_{low} into a higher-order latent representation $z_{high} = p_{\theta}(z_{high}|z_{low}) \sim N(\bar{z}_{low}, \sigma) \cdot p_{\theta}(z_{low}|x)$ which captures higher-level image features. Latent representation $p_{\theta}(z_{low}|x)$ is then combined residually through skip-connections to $p_{\theta}(z_{high}|z_{low})$; meaning that low-level information is preserved upon decoding back into image space x' via $q_{\phi}(x'|z)$. We use only one higher-order latent representation due to memory constraints and to keep our method more efficient during inference. We show that our method obtains state-of-the-art performance by utilising just one higher-order latent representation.

Additionally, we also utilise two secondary Encoders, $\{E_{high}^1, E_{low}^1\}$ during training exclusively, to re-encode the decoded respective latent representations $\{q_{\phi}(z_{high}|x)|z_{high}, q_{\phi}(z_{low}|x)\}$ back into latent space z_{high} and z'_{low} respectively. This better models and constrains the latent priors of $\{E_{high}^0, E_{low}^0\}$ to not be entirely dependent on the input image. This approach was employed in GANomaly [14] and Skip-GANomaly [15] to yield better performance. $\{E_{high}^1$ and $E_{low}^1\}$ are not enabled during inference to increase network efficiency.

Overall our VAE-GAN learning objective seeks to minimise over $\forall x \in X$:

$$L_{VAE} = L_{rec} + C(x, x') + L_{z[0]} + L_{z[1]} \quad (1)$$

where $L_{rec} = \|x - x'\|_2$ and x' is generated from the Generator Network: $D_{low}^0(D_{high}^0(E_{high}^0(E_{low}^0(x))) \cdot E_{low}^0(x))$. C is the critic, or commonly named Discriminator Network (Section III-B). Latent loss $L_{z[i]} = \|E_{[i]}^1 - E_{[i]}^0\|_2, i = \{high, low\}$.

We also introduce the notion of perceptual loss (PL) to calculate feature error rather than pixel-wise error. Previously introduced to Style Transfer *et al.* [20] we introduce usage into the problem of Anomaly Detection instead of using conventional Pixel-Wise Loss (PWL). While PWL computes pixel differences between x and x' on low-level pixel information, PL takes the advantage of taking the error between the high-level activation features [20]. Feeding the pair $(x_i, x'_i), \forall x_i \in X$ through a pre-trained conventional Convolutional Neural Network ($f(\cdot)$) obtains differing activations $(f(x_i), f(x'_i))$ of a given convolutional feature extraction layer. PL is then calculated as $\|f(x_i), f(x'_i)\|_2$. We use a VGG19 [29] network as our Perceptual Loss model and take the error between the activations of the 14th layer. We utilise two variants on perceptual loss:

- General perceptual loss ($PL_g()$) : Weights obtained by pre-training a CNN trained across ImageNet [30].
- Problem-specific perceptual loss ($PL_{ps}()$) : Weights obtained from CNN which was pre-trained in a self-supervised fashion over non-anomalous samples from the specific anomaly detection task dataset.

B. Discriminator Network

In contrast to prior work in anomaly detection which utilise conventional Discriminators [11], [12], [14], [15], [31], in this work we incorporate a Fine-grained Visual Categorisation (FGVC) Discriminator. In general, FGVC is for use in obtaining specific sub-class classification of objects (e.g. species of bird or model of car) [32]. Typical FGVC datasets are inherently difficult to classify due to highly localised and visually subtle distinguishing features between classes. Within real-world anomaly detection problems, there exist varying levels of anomaly ranging from visually obvious to negligibly subtle. Our FGVC discriminator is optimised to detect these subtle anomalies during inference by recognising the discriminating regions within presented images. It also acts as a harsher critic to our Generator module during training, promoting emphasis on generation of the objects themselves rather than the background context within images.

Our discriminator is inspired from the Weakly Supervised Data Augmentation Network (WS-DAN) architecture [33], a proven method in FGVC which obtains superior categorisation performance in the task of FGVC [33].

The WS-DAN architecture contains attention layers which allow the network to focus upon both detailed features and key discriminative object parts during inference when categorising anomalous data. This mechanism also allows attention guided data augmentation within the network leading to higher information gain and optimised augmentation of non-anomalous samples during training. The resulting attention maps are combined with feature representations via Bilinear Attention Pooling (BAP). This combined feature representation is then fed into a discriminative filter bank of 1×1 convolutions followed by a Global Max Pooling (GMP) [34] layer on the resulting feature matrix to reduce dimensionality in the output, and results in a 1×1 patch in the output which is the area of highest discrimination for the discriminator network. This allows our Generator to refine these areas in the next iteration and thus enable the overall PANDA-GAN architecture to reduce the reconstruction error substantially, gaining a better fit on the manifold over X . We use a Sigmoid activation function which issues a continuous probability score for a presented image on whether it is an element of the real dataset or a synthetic image produced by the Generator network. To prevent vanishing gradients, which is common with logistic functions, we use the residual network, ResNet-50 [35] as our main backbone architecture. The pair of real, non-anomalous data examples (x) and the generated, synthetic examples (x') from x are fed into the discriminator to obtain a probability score that each of the images is an element of X .

Overall, the Discriminator seeks to optimise:

$$L_C = -\log(C(x)) - \log(1 - C(x')) \quad (2)$$

where C represents the discriminator, or critic network. The pair (x, x') obtains probability outputs $C(x)$ and $C(x')$ respectively. L_c represents $Pr(x \in X|(x, x'))$.

C. Anomaly Scoring

Anomaly scoring is the process of categorising samples as anomalous or non-anomalous based on the knowledge that the network has obtained via the gained approximation to the manifold over X (normal samples) during training.

Anomalous samples will be reconstructed by the Generator model from this approximated manifold producing a normal appearing sample output upon generation. This allows us to infer a distribution of reconstruction error N_{recon} over both normal and anomalous samples. When combined with the two Discriminator scores across both the input sample and the synthetically generated sample distributions $N_{C(x)}$ and $N_{C(x')}$ respectively, these offer more information during anomaly scoring. Due to the independence of these normally distributed random variables, we can combine them into a new distribution using summation $N_{com} \sim N(\Sigma(\mu_i, \sigma_i^2)) = N_{com} = N_{recon} + N_{C(x)} + N_{C(x')}$. We then normalise N_{com} to values $0 < x < 1$ via Equation 3 as follows:

$$N_{com} = \forall x \in [N_{com}], \frac{x'_i - x_{min}}{x_{max} - x_{min}} \quad (3)$$

Once we obtain the distribution N_{com} , an anomaly score is given to each sample presented to the network. A boundary between $N_{com}(x_{normal})$ and $N_{com}(x_{anomalous})$ determines whether the presented sample is anomalous on the basis of maximum likelihood.

IV. EVALUATION

A. Experimental Setup

The experimental setup comprises of the following dataset configurations:

- **Plant Village** [7]: comprises of 11 genera of agricultural plant for food with 26 diseases. The present diseases range from those which are subtle in appearance (Powdery Mildew) to diseases such as Isariopsis and Black Measles which are visually obvious.
- **X-ray Security Electronics Dataset** [8]: comprises large consumer electronics items (e.g., laptops) with and without intra-object anomaly concealment present. Anomaly concealments consist of replica Improvised Explosive Devices (IED), metal screws, metal plates, knife blades and similar hidden inside the electronic items.
- **UCSDped1** [9]: comprises surveillance video of pedestrian-only zone featuring pedestrians walking, running, standing, or any such similar activity. Anomalies are pedestrians riding or driving vehicles (i.e. bicycles, skateboards, scooters, vans, cars, etc) in the zone.
- **MNIST** [5]: comprises a handwritten numeric digit (0-9) leave-one-out classification task (one select digit

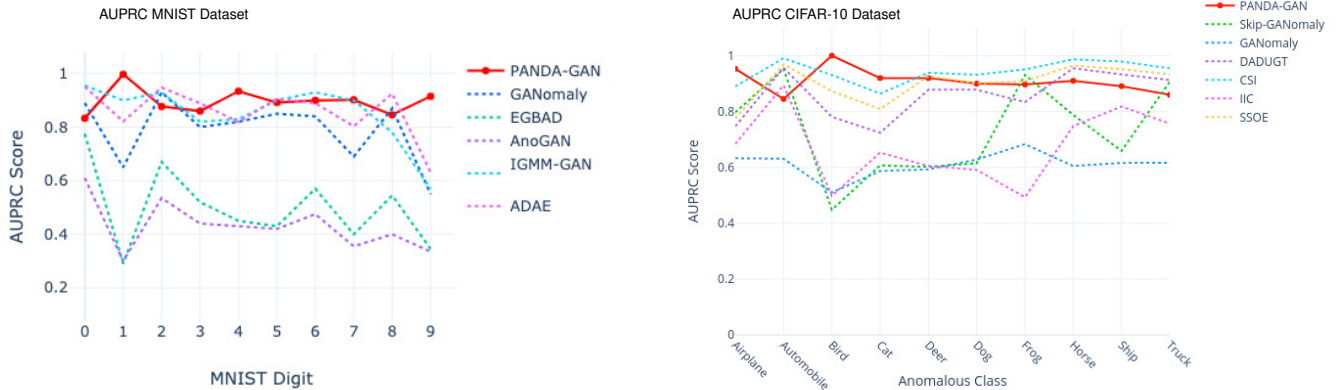


Fig. 3. AUPRC performance comparison of PANDA-GAN against prior state-of-the-art methods of anomaly detection across 'leave-one-out' tasks across **Left:** MNIST handwritten digit dataset **Right:** CIFAR-10 object dataset.

class is anomalous, all remaining digit classes are non-anomalous) across all digits from 0 to 9.

- **CIFAR-10** [6]: comprises 10 classes of tiny objects with dimension 32×32 such that a leave-one-out anomaly classification task (one select object class is anomalous, all remaining object classes are non-anomalous) is posed.
- **MVTEC** [10]: comprises of high-resolution industrial inspection data comprising of 15 object and texture classes of typical factory-line objects. Sub-classes of each object contain both visual defects of the respective object and corresponding defect-free objects.

We use the dataset split for $\{train:validate:test\}$ as following: $\{13,593:2,589:12,661\}$ for Plant Village [7], $\{229:25:125\}$ for X-ray Security Electronics [8]. The $\{train:test\}$ split for MNIST and CIFAR-10 is $\{80\%:20\%\}$ across both datasets [14], [25]. The hyper-parameters and data configurations are fine-tuned by systematic grid search in order to obtain the best results across the problems presented in this work. Pixel values in input images are normalised to a mean and a standard deviation of 0.5. All models use ADAM momentum [36] except our Perceptual Loss model which uses Stochastic Gradient Descent (SGD) with momentum 0.9. Learning rates used: 7×10^{-6} - Generator, 1×10^{-5} - Discriminator, and 1×10^{-4} - Perceptual Loss model. Training is performed on a Nvidia Titan X GPU using a batch size of 15.

B. Results and Discussion

The AUPRC statistical score across the classical 'leave-one-out' anomaly detection tasks (MNIST / CIFAR-10) are outlined in Figure 3 where it can be observed that our approach (PANDA-GAN) outperforms prior state-of-the-art approaches on these seminal, albeit unrealistic benchmark tasks.

As is evident from Figure 3, further direct comparison of model performance solely across these 'leave-one-out' MNIST / CIFAR-10 based anomaly detection tasks is becoming decreasingly informative due to potential performance saturation

among competing approaches. Across the MNIST task (left) it can be seen that our method obtains state-of-the-art results across 40% of classes and obtains close performance to the other prior methods while exhibiting uniform performance across all classes. Most noticeably is the result across the digit 9 whereby PANDA-GAN is over 0.2 AUPRC higher than prior methods. Across CIFAR-10 (right), our method obtains state-of-the-art in 30% of classes and matches closely with other such methods (DADUGT [46], CSI [47], and SSOE [48]) while also obtaining close to uniform performance across all classes.

By contrast, Table I outlines quantitative results across the challenging real-world benchmark datasets of Plant Village [7], Laptop X-ray [8] and UCSDPed1 [9] providing numerous statistical comparatives including Area Under Curve (AUC), the 95% confidence interval of the AUC, inference time (I/t , ms) per image. Datasets [7], [8] feature particularly subtle anomalies by nature and as such pose as challenging tasks for semi-supervised anomaly detection models.

PANDA-GAN obtains the highest AUC value across both image based datasets (Plant Village: 0.776- using Pixel-Wise Loss (PWL); Laptop X-ray: 0.51- using Problem Specific Perceptual Loss(PL_{ps})) in comparison to leading state-of-the-art methods [11], [12], [14], [15], [21], [25], [29] (Table I). Over multiple evaluations, PANDA-GAN also obtains tighter confidence-intervals compared to prior semi-supervised work illustrating our PANDA method can produce more stable and reliable results across the same dataset while other such approaches can suffer from sporadic performance at inference (Table I). Observing the I/t (ms), the PANDA method is significantly faster than prior methods.

In Table II, the quantitative AUC results across the MVTEC dataset can be observed. This is a challenging dataset due to the large variation in appearance of anomalies present in textures and objects. Some objects (carpet, hazelnut, screw) exhibit visually distinct and obvious anomalies, but other

TABLE I

RESULTS OF MODELS ACROSS LEAF DISEASE [7] AND X-RAY LAPTOP ANOMALY DETECTION [8] IMAGE DATASETS AS WELL AS RESULTS ACROSS UCSDPED1 [9] PEDESTRIAN DETECTION AND CROWD CONTROL VIDEO DATASET USING FRAME-LEVEL COMPARISON [37].

Model	Loss	Image Dataset												
		Plant Village [7]					Laptop X-ray [8]							
		AUC	95% CI (AUC)	Average Rec_Err	Average Adv_Err	I/t(ms)	AUC	95% CI (AUC)	Average Rec_Err	Average Adv_Err	I/t(ms)			
VAE [21]	-	0.65	0.60<x<0.70	0.56	-	6.9	0.21	0.19<x<0.23	0.80	-	9.4			
AnoGAN [11]	-	0.65	0.65<x<0.66	0.45	0.88	7151	0.41	0.39<x<0.42	0.4	0.92	7223			
EGBAD [25]	-	0.70	0.65<x<0.67	0.40	0.92	87	0.47	0.42<x<0.43	0.41	0.94	89			
GANomaly [14]	-	0.73	0.68<x<0.73	0.39	0.75	28	0.49	0.41<x<0.51	0.34	0.78	273			
f-AnoGAN [12]	-	0.765	0.65<x<0.78	0.12	0.72	65	0.50	0.49<x<0.53	0.1	0.72	86			
Skip-GANomaly [15]	-	0.771	0.74<x<0.77	0.13	0.74	123	0.51	0.48<x<0.58	0.11	0.68	112			
PANDA-GAN	PWL	0.776	0.77<x<0.78	0.012	0.994	15.2	0.42	0.30<x<0.48	0.052	0.987	16.8			
	PL _g ()	0.74	0.73<x<0.75	0.40	0.993	20	0.45	0.29<x<0.52	0.015	0.658	36			
	PL _{ps} ()	0.75	0.76<x<0.78	0.20	0.986	20.8	0.51	0.48<x<0.55	0.045	0.782	30			
Model	Loss	Video Dataset												
		UCSDPed1 [9]												
		AUC	EER											
		SF [38]	-	0.675	31									
		MPPCA [39]	-	0.7696	40									
		MDT [40]	-	0.818	25									
		SRC [41]	-	0.86	19									
		AMDN [42]	-	0.921	16									
		PCA-NET GMM [37]	-	0.926	11.2									
		AED-GAN [43]	-	0.974	8									
		PANDA-GAN	PWL	0.945	35									
PL _g ()	0.95		75											
PL _{ps} ()	0.93		96											

TABLE II
AUPRC RESULTS ACROSS MVTEC [10] DATASET.

Model	Classes															AUC _{avg}
	Bottle	Cable	Capsule	Carpet	Grid	Hazelnut	Leather	Metal Nut	Pill	Screw	Tile	Toothbrush	Transistor	Wood	Zipper	
AnoGAN [11]	0.8	0.477	0.442	0.337	0.871	0.259	0.451	0.284	0.711	1	0.401	0.439	0.692	0.567	0.715	0.563
GANomaly [14]	0.794	0.711	0.721	0.821	0.743	0.874	0.808	0.694	0.671	1	0.72	0.7	0.808	0.92	0.744	0.782
Skip-GANomaly [15]	0.937	0.674	0.718	0.795	0.657	0.906	0.908	0.79	0.758	1	0.85	0.689	0.814	0.919	0.663	0.805
DA-GAN [44]	0.983	0.665	0.687	0.903	0.867	1	0.944	0.815	0.768	1	0.961	0.95	0.794	0.979	0.781	0.873
U-Net [45]	0.863	0.636	0.673	0.774	0.857	0.996	0.87	0.676	0.78	1	0.964	0.811	0.674	0.958	0.75	0.819
PANDA-GAN	0.826	0.68	0.98	0.95	0.95	0.922	0.75	0.79	0.95	1	0.85	0.66	0.9	0.68	0.62	0.834

TABLE III
ABLATION STUDY OF PANDA-GAN ACROSS PLANT VILLAGE [7] AND LAPTOP ANOMALY [8].

Model	Dataset											
	Plant Village						Laptop Anomaly					
	Loss			Network Architecture			Loss			Network Architecture		
	PWL	PL(g)	PL(ps)	$E_{high}^0 \cap D_{high}^0$	E_{high}^1	E_{low}^1	PWL	PL(g)	PL(ps)	$E_{high}^0 \cap D_{high}^0$	E_{high}^1	E_{low}^1
PANDA-GAN	0.75	0.754	0.762	✗	-	✗	0.38	0.42	0.434	✗	-	✗
	0.752	0.74	0.746	✗	-	✓	0.419	0.442	0.446	✗	-	✓
	0.751	0.738	0.744	✓	✗	✗	0.464	0.468	0.476	✓	✗	✗
	0.764	0.751	0.762	✓	✗	✓	0.462	0.442	0.43	✓	✗	✗
	0.771	0.764	0.775	✓	✓	✗	0.496	0.515	0.504	✓	✓	✗
	0.776	0.741	0.769	✓	✓	✓	0.42	0.451	0.512	✓	✓	✓
PANDA-GAN DCGAN Discriminator	0.769	0.751	0.74	✓	✓	✓	0.413	0.424	0.469	✓	✓	✓

objects (cable, metal nut, toothbrush) feature subtle anomalies which are hard to detect. This is reflected in the results of various methods across this dataset with PANDA-GAN obtaining superior AUC performance across 6 classes.

Our ablation study (Table III) produces evaluation over individual components to our novel architecture with respect to variations in both loss function, our network architecture components ($E_{high}^0 \cap D_{high}^0$, E_{high}^1 , E_{low}^1) and our choice of discriminator architecture across two of the more challenging real-world anomaly detection task datasets. For comparison we include the DCGAN [23] discriminator architecture from

GANomaly / Skip-GANomaly [14], [15], which is the next best performing approach in terms of AUC across the same datasets (Table I to compare against our FGVC-based discriminator architecture choice.

From the results of Table III, it can be observed that synergy exists between components of our generator network obtaining the highest AUC value only when all three of our novel components are activated. Generally we see that the more components we activate in our architecture, the better the performance obtained during our ablation study. Overall, the problem-specific perceptual loss (PL_{ps}) performs better across

the Laptop X-ray dataset by a clear margin from the other loss functions tested against. Across the Plant Village dataset, there is negligible difference between the pixel-wise loss and the problem specific perceptual loss. Both performed almost identically and gained a clear advantage over the general perceptual loss function (PL_g).

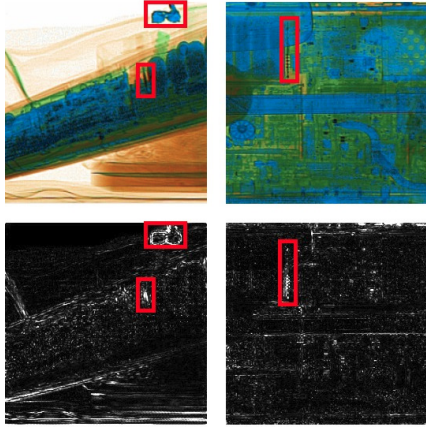


Fig. 4. **Top:** X-ray images of large electronic items (Anomalous region ground truth outlined with red bounding box). **Bottom:** Non-threshold instance segmentation masks produced by PANDA-GAN outlining anomalous artifacts within X-ray scans.

V. CONCLUSION

We propose the semi-supervised method of Perceptually Aware Neural Detection of Anomalies (PANDA-GAN), a Variational Autoencoder Generative Adversarial Network (VAE-GAN) based architecture. PANDA-GAN includes three novel proposals: a Fine-Grained Visual Categorisation Discriminator Network (FGVC) to ease the problem of low-inter class variance present in anomaly detection problems, a dual-latent space implementation that carries higher-level features in given images forward in the architecture and the use of a perceptual loss function to compute feature error which enables higher aptitude at detecting anomalies within the given problems presented in this work.

Our PANDA-GAN architecture obtains superior results across: (CIFAR-10, $AUPRC_{avg}$: 0.91; MNIST, $AUPRC_{avg}$: 0.90) and challenging real-world datasets (Plant Leaf Disease, AUC: 0.776, Threat Item X-ray, AUC: 0.51, MVTEC, AUC_{avg} : 0.83), in addition to video based frame-level anomaly detection (UCSDPed1, AUC: 0.95). Outperforming prior work by [9], [11], [12], [14], [15], [19], [25], [37]–[39], [41], [42], [46], [49].

REFERENCES

- [1] S. Akçay and T. Breckon, “Towards automatic threat detection: A survey of advances of deep learning within x-ray security imaging,” *ArXiv*, vol. abs/2001.01293, 2020.
- [2] V. Narayan and B. Shaju, *Malware and Anomaly Detection Using Machine Learning and Deep Learning Methods*, pp. 104–131, 01 2020.
- [3] X. Hoque and S. Sharma, *Ensembled Deep Learning Approach for Maritime Anomaly Detection System*, pp. 862–869, 09 2019.
- [4] X. Fang, W. Guo, Q. Li, J. Zhu, Z. Chen, J. Yu, B. Zhou, and H. Yang, “Sewer pipeline fault identification using anomaly detection algorithms on video sequences,” *IEEE Access*, vol. 8, pp. 39574–39586, 2020.

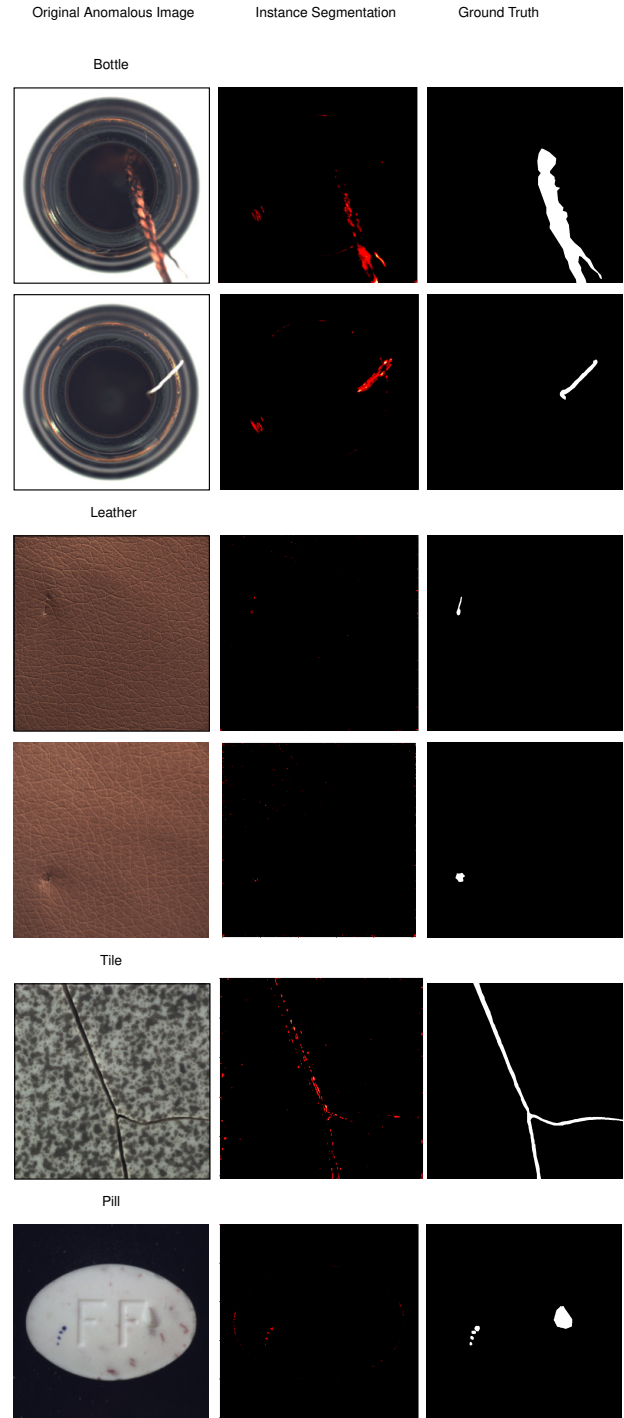


Fig. 5. Anomalous regions obtained from PANDA-GAN across the MVTEC dataset together with the original image and the ground truth for the anomalous region.

- [5] Y. LeCun, C. Cortes, and C.J. Burges, “Mnist handwritten digit database,” *ATT Labs [Online]. Available: <http://yann.lecun.com/exdb/mnist>*, vol. 2, 2010.
- [6] A. Krizhevsky and G. Hinton, “Learning multiple layers of features from tiny images,” *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- [7] David P. Hughes and M. Salathé, “An open access repository of images on plant health to enable the development of mobile disease

- diagnostics through machine learning and crowdsourcing,” *ArXiv*, vol. abs/1511.08060, 2015.
- [8] N. Bhowmik, Y.F.A. Gaus, S. Akçay, J. W. Barker, and T. P. Breckon, “On the impact of object and sub-component level segmentation strategies for supervised anomaly detection within x-ray security imagery,” in *18th IEEE International Conference on Machine Learning and Applications (ICMLA 2019)*. December 2019, IEEE.
 - [9] V. Mahadevan, W. X. Li, V. Bhalodia, and N. Vasconcelos, “Anomaly detection in crowded scenes,” in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 1975–1981.
 - [10] P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, “Mvtec ad — a comprehensive real-world dataset for unsupervised anomaly detection,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9584–9592.
 - [11] T. Schlegl, P. Seeböck, W. Philipp, U. Schmidt-Erfurth, and G. Langs, “Unsupervised anomaly detection with generative adversarial networks to guide marker discovery,” 03 2017, pp. 146–157.
 - [12] T. Schlegl, P. Seeböck, S.M. Waldstein, G. Langs, and U. Schmidt-Erfurth, “f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks,” *Medical Image Analysis*, vol. 54, pp. 30–44, 2019.
 - [13] Y. F. A. Gaus, N. Bhowmik, S. Akçay, P. M. Guillén-García, J. W. Barker, and T. P. Breckon, “Evaluation of a dual convolutional neural network architecture for object-wise anomaly detection in cluttered x-ray security imagery,” in *2019 International Joint Conference on Neural Networks (IJCNN)*, 2019, pp. 1–8.
 - [14] S. Akçay, A. Atapour-Abarghouei, and T. P. Breckon, “Ganomaly : semi-supervised anomaly detection via adversarial training,” in *14th Asian Conference on Computer Vision*, number 11363 in Lecture notes in computer science, pp. 622–637. Springer, 2019.
 - [15] S. Akçay, A. Atapour-Abarghouei, and T.P. Breckon, “Skip-GANomaly: Skip Connected and Adversarially Trained Encoder-Decoder Anomaly Detection,” *Proceedings of the International Joint Conference on Neural Networks*, vol. 2019-July, jan 2019.
 - [16] I.J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” in *International Conference on Learning Representations*, 2015.
 - [17] A. Paudice, L. Muñoz-González, A. Gyorgy, and E.C Lupu, “Detection of adversarial training examples in poisoning attacks through anomaly detection,” *arXiv preprint arXiv:1802.03041*, 2018.
 - [18] C. Baur, B. Wiestler, S. Albarqouni, and N. Navab, “Deep Autoencoding Models for Unsupervised Anomaly Segmentation in Brain MR Images,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11383 LNCS, pp. 161–169, apr 2018.
 - [19] Ha. S. Vu, D. Ueta, K. Hashimoto, K. Maeno, S. Pranata, and S. M. Shen, “Anomaly Detection with Adversarial Dual Autoencoders,” 2019.
 - [20] J. Johnson, A. Alahi, and L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *European Conference on Computer Vision*, 2016.
 - [21] D.P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations*, 2014.
 - [22] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in Neural Information Processing Systems*, 2014, vol. 3, pp. 2672–2680.
 - [23] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” in *4th International Conference on Learning Representations*, November 2016.
 - [24] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein generative adversarial networks,” in *Proceedings of the 34th International Conference on Machine Learning*. 06–11 Aug 2017, vol. 70 of *Proceedings of Machine Learning Research*, pp. 214–223, PMLR.
 - [25] H. Zenati, C.S. Foo, B. Lecouat, G. Manek, and V.R. Chandrasekhar, “Efficient gan-based anomaly detection,” *arXiv*, vol. abs/1802.06222, 2018.
 - [26] J. Donahue, T. Darrell, and K. Philipp, “Adversarial feature learning,” in *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*. May 2019, International Conference on Learning Representations, ICLR.
 - [27] F. Di Mattia, P. Galeone, M. De Simoni, and E. Ghelfi, “A survey on gans for anomaly detection,” *ArXiv*, vol. abs/1906.11632, 2019.
 - [28] A.B.L. Larsen, S.K.Sønderby, H. Larochelle, and O. Winther, “Auto-encoding beyond pixels using a learned similarity metric,” in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, 2016, p. 1558–1566.
 - [29] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *International Conference on Learning Representations*, 2015.
 - [30] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
 - [31] A. Razavi, A.V.D. Oord, and O. Vinyals, “Generating diverse high-fidelity images with vq-vae-2,” in *Advances in Neural Information Processing Systems*, 2019, vol. 32, pp. 14866–14876.
 - [32] Y. Wang and Z. Wang, “A survey of recent work on fine-grained image classification techniques,” *Journal of Visual Communication and Image Representation*, vol. 59, pp. 210 – 214, 2019.
 - [33] T. Hu and H. Qi, “See better before looking closer: Weakly supervised data augmentation network for fine-grained visual classification,” *CoRR*, vol. abs/1901.09891, 2019.
 - [34] Y. Wang, V. Morariu, and L. Davis, “Learning a discriminative filter bank within a cnn for fine-grained recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4148–4157.
 - [35] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
 - [36] D.P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR*, Y. Bengio and Y. LeCun, Eds., 2015.
 - [37] H. Shaonian, H. Dongjun, and Z. Xinmin, “Learning multimodal deep representations for crowd anomaly event detection,” *Mathematical Problems in Engineering*, vol. 2018, pp. 1–13, 01 2018.
 - [38] R. Mehran, A. Oyama, and M. Shah, “Abnormal crowd behavior detection using social force model,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 935–942.
 - [39] J. Kim and K. Grauman, “Observe locally, infer globally: a space-time mrf for detecting abnormal activities with incremental updates,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
 - [40] L. Weixin, V. Mahadevan, and N. Vasconcelos, “Anomaly detection and localization in crowded scenes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 1, pp. 18–32, 2014.
 - [41] C. Yang, Y. Junsong, and L. Ji, “Abnormal event detection in crowded scenes using sparse representation,” *Pattern Recognition*, vol. 46, no. 7, pp. 1851–1864, 2013.
 - [42] D. Xu, E. Ricci, Y. Yan, J. Song, and N. Sebe, “Learning deep representations of appearance and motion for anomalous event detection,” in *The British Machine Vision Conference (BMVC)*, 2015.
 - [43] M. Ravanbakhsh, M. Nabi, E. Sanginetto, L. Marcenaro, C. Regazzoni, and N. Sebe, “Abnormal event detection in videos using generative adversarial nets,” *2017 IEEE International Conference on Image Processing (ICIP)*, pp. 1577–1581, 2017.
 - [44] T. W. Tang, W. H. Kuo, J. H. Lan, C. F. Ding, H. Hsu, and H. T. Young, “Anomaly detection neural network with dual auto-encoders gan and its industrial inspection applications,” *Sensors*, vol. 20, pp. 3336, 06 2020.
 - [45] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2015, vol. 9351 of LNCS, pp. 234–241, Springer.
 - [46] I. Golan and R. El-Yaniv, “Deep anomaly detection using geometric transformations,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, p. 9781–9791.
 - [47] J. Tack, S. Mo, J. Jeong, and J. Shin, “Csi: Novelty detection via contrastive learning on distributionally shifted instances,” in *Advances in Neural Information Processing Systems*, 2020.
 - [48] D. Hendrycks, M. Mazeika, S. Kadavath, and D. Song, “Using self-supervised learning can improve model robustness and uncertainty,” in *Advances in Neural Information Processing Systems*, 2019, vol. 32, pp. 15663–15674.
 - [49] D. Smolyak, K. Gray, S. Badirli, and G. Mohler, “Coupled igmm-gans with applications to anomaly detection in human mobility data,” *ACM Trans. Spatial Algorithms Syst.*, vol. 6, no. 4, June 2020.