Generative Models and Quality Constraints for Anomaly Detection

BOUGAHAM, Arnaud

Publication date:
2024

Document Version
Peer reviewed version

Generative Models and Quality Constraints for Anomaly Detection : Application to Industrial and Medical Images

Arnaud Bougaham[1], Benoît Frénay[1] and Isabelle Linden[2]

[1]Faculty of Computer Science, NaDI Institute, University of Namur, Rue Grandgagnage 21, Namur, 5000, Belgium.
[2]Department of Management Sciences, NaDI Institute, University of Namur, Rempart de la Vierge 8, Namur, 5000, Belgium.

Context

Advanced anomaly detection :

- Crucial quality control step
- Important part of Industry 4.0 opportunities

Traditional algorithms suffer from practical drawbacks :

- High false positive rate
- Limited Regions of Inspection

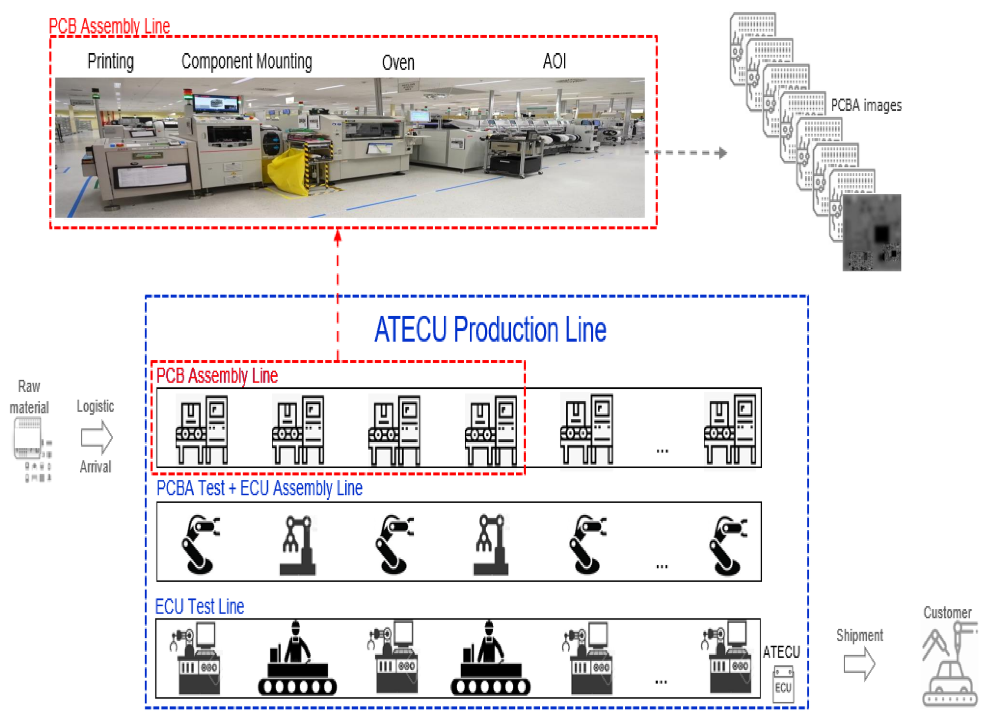

Figure 1: PCBA Manufacturing Process Flow

Research Questions

- What are the best deep learning techniques to detect anomalies that exist in real-world industrial datasets? (unsupervised learning, high-resolution images, imbalanced datasets, etc.)
- How to integrate the business constraints (full TPR, acceptable inference time, worker interactions, explainable decisions, etc.) into a binary normal/abnormal classification algorithm?

Methods & Results

[1][2] VQGanoDIP = VQGAN Reconstruction + Metrics Classification for Anomaly Detection

- Train an Encoder, a Codebook, a Generator and a Discriminator through a GAN framework, in an autoencoder architecture
- Get statistics on the residual image and in the networks losses to quantify how the image is different from the normality
- Train a binary extra tree classifier to discriminate between normal and abnormal products

$$\mathcal{Q}^* = \underset{E,G,Z}{argmin}\ \underset{D}{max}\ \underset{x\sim p(x)}{\mathbb{E}}\left[\mathcal{L}_{VQ}(E,G,Z) + \lambda\mathcal{L}_{GAN}(\{E,G,Z\},D)\right]$$

$$\begin{cases} \mathcal{L}_{GAN}(\{E,G,Z\},D) = \left[\log D(x) + \log\left(1 - D\left(G\left(Z\left(E(x)\right)\right)\right)\right)\right], \\ \mathcal{L}_{VQ}(E,G,Z) = \left\|x - G\left(Z\left(E(x)\right)\right)\right\|^2 + \|sg[E(x)] - z_q\|_2^2 + \|sg[z_q] - E(x)\|_2^2, \\ \lambda = \frac{\nabla_{GL}[\mathcal{L}_{rec}(\{E,G,Z\})]}{\nabla_{GL}[\mathcal{L}_{GAN}(\{E,G,Z\},D)]+\delta}, \end{cases}$$

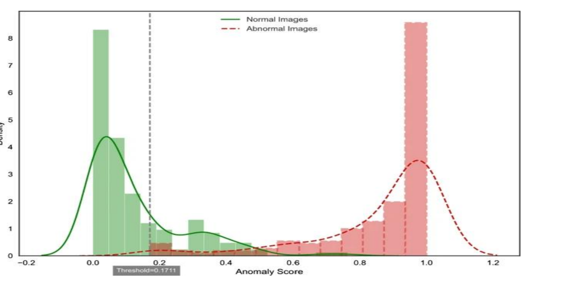

Dataset (Classifier)	Accuracy(%)↑ STD	ZFN
PCBA (ET)	95.69	87.93
Cable (XGBoost)	76.82	57.94
Carpet (LR)	85.6	50.21
Grid (LR)	95.98	85.43
Hazelnut (LGBM)	98.95	98.25
Leather (XGBoost)	92.17	90.43
Screw (ADA)	93	83.67
Transistor (LGBM)	88.7	49.15
Zipper (LGBM)	92.55	81.57

Quality: 37%
False Positives saved

Performance: 72%
Inspection time saved

Figure 2: VQGanoDIP Training and Inference Architecture, Loss Functions, Anomaly Localisation, Qualitative and Quantitative Results

[3] Cycle-Consistent Adversarial Networks for Industrial and Medical Anomaly Detection

- Take the few abnormal data available into consideration, to train a cycle-GAN
- Evaluate on both Industrial and medical datasets

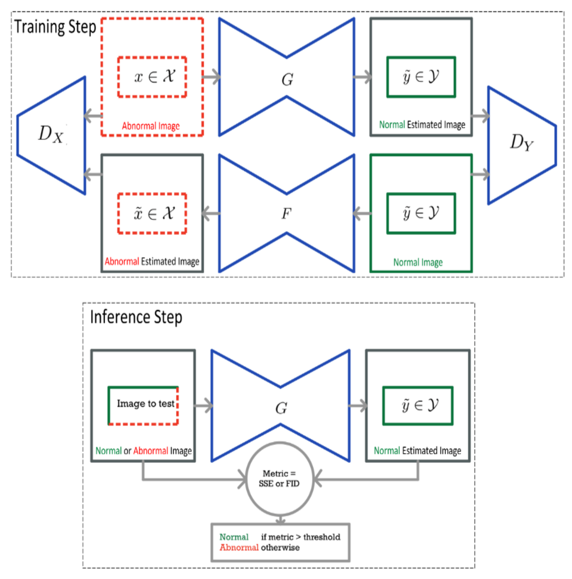

		CycleGAN-AD-256 (ours) ZFN	ACC	AUC	CycleGAN-AD-64 (ours) ZFN	ACC	AUC	Ganomaly [3] ZFN	ACC	AUC	Padim [10] ZFN	ACC	AUC	PatchCore [25] ZFN	ACC	AUC
Hazelnut	FID	98.00 ± 2.14	99.14 ± 0.70	99.89 ± 0.12	74.86 ± 12.89	94.57 ± 2.10	97.32 ± 0.91	51.14 ± 1.67	66.57 ± 3.33	63.31 ± 4.73	/	/	/	/	/	/
	SSE	96.29 ± 2.94	98.29 ± 1.07	99.67 ± 0.25	95.43 ± 2.29	98.29 ± 0.57	99.71 ± 0.14	80.29 ± 10.32	87.14 ± 7.17	92.16 ± 4.90	53.43 ± 1.71	95.71 ± 0.00	92.02 ± 0.29	54.29 ± 0.00	95.71 ± 0.00	92.16 ± 0.00
Screw	FID	52.03 ± 2.18	57.63 ± 4.25	54.34 ± 8.10	50.51 ± 0.68	57.46 ± 3.36	53.08 ± 5.96	51.19 ± 1.15	57.63 ± 1.42	54.04 ± 3.14	/	/	/	/	/	/
	SSE	52.37 ± 3.53	57.97 ± 5.46	52.81 ± 6.29	52.03 ± 3.24	57.63 ± 2.89	51.31 ± 5.06	62.71 ± 10.39	77.12 ± 5.22	80.86 ± 3.49	55.59 ± 4.57	66.27 ± 5.42	54.02 ± 11.91	59.32 ± 0.00	90.68 ± 0.00	84.83 ± 0.00
Tile	FID	91.43 ± 7.88	98.33 ± 1.21	99.30 ± 0.73	58.81 ± 6.37	78.81 ± 1.90	83.53 ± 2.61	57.62 ± 3.07	70.48 ± 2.05	72.47 ± 1.56	/	/	/	/	/	/
	SSE	78.10 ± 7.85	89.76 ± 3.42	95.40 ± 2.52	52.86 ± 2.21	75.24 ± 3.64	78.32 ± 2.67	50.48 ± 0.58	53.57 ± 2.61	42.39 ± 5.50	61.90 ± 0.00	88.10 ± 0.00	81.86 ± 0.00	61.90 ± 0.00	88.10 ± 0.00	81.86 ± 0.00
Wood	FID	91.33 ± 6.78	97.00 ± 1.94	99.04 ± 0.79	71.00 ± 14.85	88.67 ± 1.63	92.82 ± 2.09	53.33 ± 3.80	56.00 ± 2.71	43.62 ± 3.76	/	/	/	/	/	/
	SSE	97.00 ± 3.71	97.67 ± 2.91	98.89 ± 1.37	92.33 ± 6.96	96.33 ± 2.87	98.48 ± 1.51	61.00 ± 4.29	71.67 ± 7.67	75.09 ± 7.24	94.67 ± 10.67	95.33 ± 9.33	95.02 ± 9.96	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
Brain MRI	FID	78.57 ± 4.65	87.76 ± 3.35	93.19 ± 2.56	73.27 ± 4.05	79.39 ± 2.99	80.85 ± 3.92	54.90 ± 2.08	60.20 ± 3.48	58.88 ± 4.83	/	/	/	/	/	/
	SSE	84.49 ± 3.95	86.94 ± 3.73	91.50 ± 2.92	84.08 ± 4.81	88.37 ± 3.63	91.96 ± 3.16	61.02 ± 4.63	68.37 ± 1.12	69.98 ± 2.78	50.41 ± 0.50	92.45 ± 13.06	91.58 ± 12.78	50.41 ± 0.50	98.98 ± 0.00	97.98 ± 0.02
Breast Ultrasound	FID	83.23 ± 6.53	91.38 ± 1.23	95.45 ± 0.91	84.92 ± 3.29	85.85 ± 3.46	89.62 ± 3.52	57.69 ± 2.48	70.15 ± 4.83	74.41 ± 5.29	/	/	/	/	/	/
	SSE	85.23 ± 2.98	89.38 ± 3.31	92.53 ± 2.41	86.46 ± 3.85	87.85 ± 3.01	91.23 ± 2.31	61.23 ± 5.19	68.62 ± 2.45	71.81 ± 2.29	50.62 ± 0.31	96.77 ± 4.92	97.27 ± 2.42	50.62 ± 0.31	99.23 ± 0.00	98.48 ± 0.01
Retina OCT	FID	50.73 ± 0.59	97.23 ± 0.05	98.81 ± 0.08	50.19 ± 0.08	92.10 ± 0.22	96.87 ± 0.09	50.09 ± 0.07	69.92 ± 5.23	76.25 ± 6.61	/	/	/	/	/	/
	SSE	50.29 ± 0.26	96.74 ± 0.07	98.49 ± 0.10	51.02 ± 0.75	96.45 ± 0.04	98.33 ± 0.07	50.37 ± 0.40	79.33 ± 1.34	86.86 ± 1.47	50.01 ± 0.01	93.90 ± 3.03	98.24 ± 1.51	50.01 ± 0.01	99.95 ± 0.05	99.97 ± 0.00
MEAN		79.89	89.93	91.43	74.31	86.25	88.05	62.03	74.89	78.83	59.52	89.79	87.14	60.94	96.09	93.61

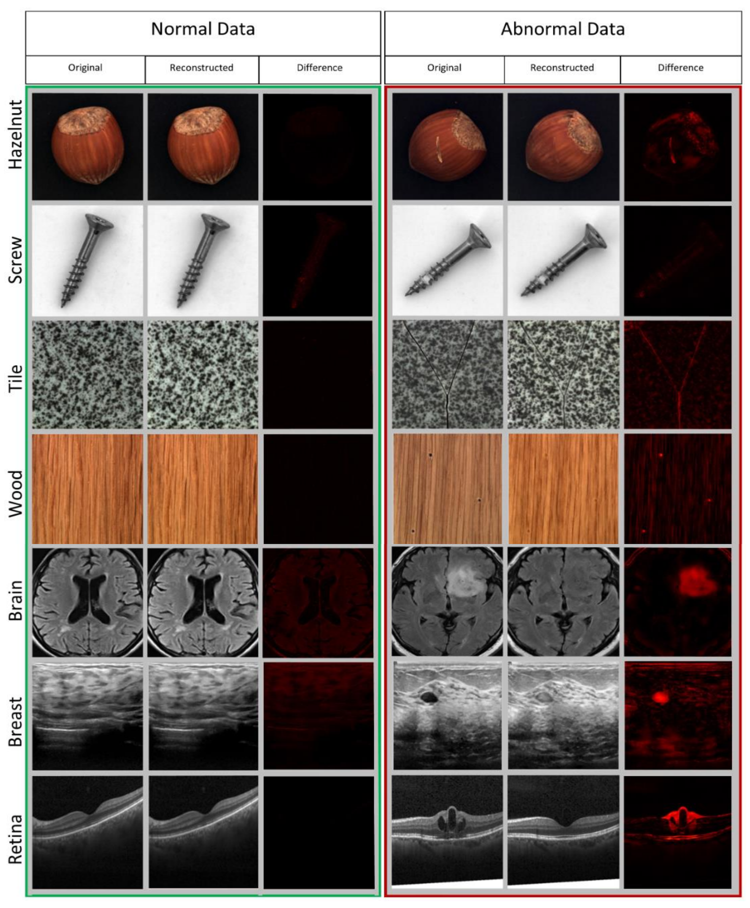

Figure 3: CycleGAN AD Training and Inference Architecture , Anomaly Localisation, Qualitative and Quantitative Results

Future Works

1st part of the approach is an image reconstruction through a GAN model. To optimally reconstruct the input image (and focus on the real anomalies), other generative model architectures will be compared.

- ViT in GAN : Replace all or parts of the CNNs by Visual Transformers.

CNNs of GAN → Vision Transformer

- Diffusion Model : Consider an autoencoder with latent diffusion process.

GAN in autoencoder → Diffusion Model

2nd part of the approach is the metrics classification. A customized function loss is being designed to reflect a full sensitivity constraint.

- A partial AUC [4] is formulated, so that only the Full TPR part is targeted while minimizing the FPR.
- This pAUC is then approximated by a Wilcoxon-Mann-Whitney Statistic loss [5], well fitted for a neural network classifier.

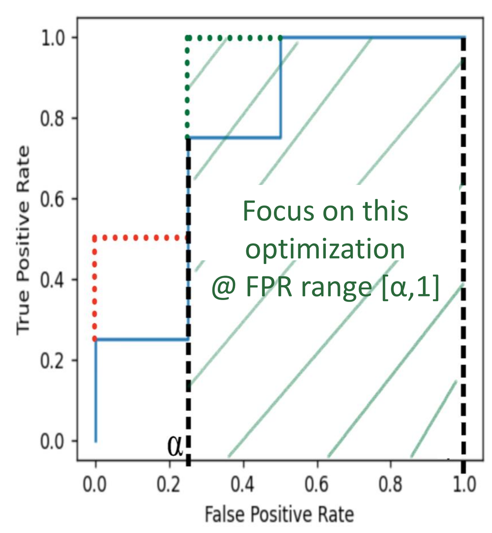

$$\mathcal{L}_{WMW_{approx}\,pAUC} = \begin{cases} \frac{1}{S^+ S^-_\alpha}\sum_{i=1}^{S^+}\sum_{j=1}^{S^-_\alpha}(s_i^+ + \gamma - s_j^-)^2 & \text{if } s_i^+ + \gamma < s_j^- \\ 0 & \text{otherwise} \end{cases}$$

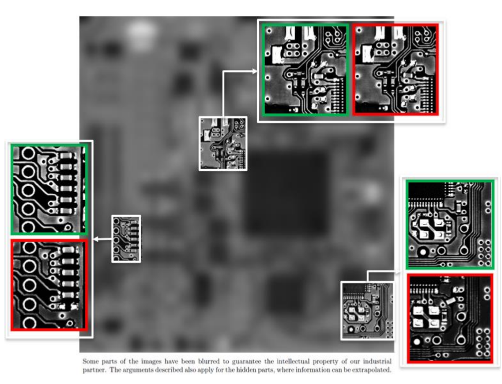

On-site industrial integration is performed, highly focused on the worker habits, the inference time and quality constraints.

Another medical use case is being studied, on activation detection during alternative rest/stimulation phases for fPET sinograms.

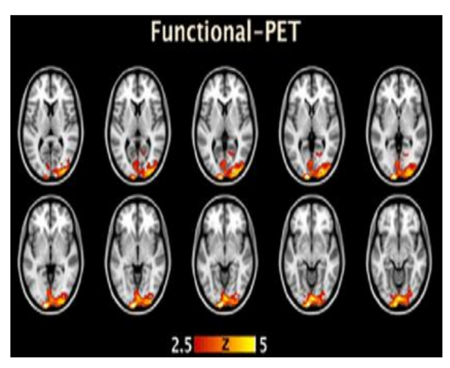

References

[1] Bougaham, A. et al. (2021). GanoDIP - GAN Anomaly Detection through Intermediate Patches: a PCBA Manufacturing Case. Proceedings of the Third International Workshop LIDTA, PMLR.
[2] Bougaham, A. et al. (2023). Composite Score for Anomaly Detection in Imbalanced Real-World Industrial Dataset. Machine Learning Journal, Springer.
[3] Bougaham, A. et al. (2023). Industrial and Medical Anomaly Detection Through Cycle-Consistent Adversarial Networks. arXiv preprint.
[4] Dodd, LE. And Pepe, MS. (2003), Partial AUC estimation and regression. Biometrics.
[5] Yan, L. et al. (2003), Optimizing classifier performance via an approximate to the Wilcoxon-Mann-Whitney statistic. Proceedings of ICML.

Arnaud Bougaham
Ph.D. Candidate AI applied to Industry
arnaud.bougaham@unamur.be

humalearn.info.unamur.be
directory.unamur.be/staff/abougaha

LinkedIn

NADI Namur Digital Institute

UNIVERSITÉ DE NAMUR

HuMaLearn