Combining GANs and AutoEncoders for Efficient Anomaly Detection

Fabio Carrara[1], Giuseppe Amato[1], Luca Brombin[2], Fabrizio Falchi[1], Claudio Gennaro[1]
[1]ISTI CNR, Pisa, Italy — [2]University of Pisa

fabio.carrara@isti.cnr.it
https://github.com/fabiocarrara/cbigan-adv

ICPR 2021 - January, 14th
Milan, Italy (Virtual)

AIMH Lab
Artificial Intelligence for Media and Humanities

The Problem: Visual Anomaly Detection

GOAL — detect anomalous images (assign either ✓ or ✗ to input image)

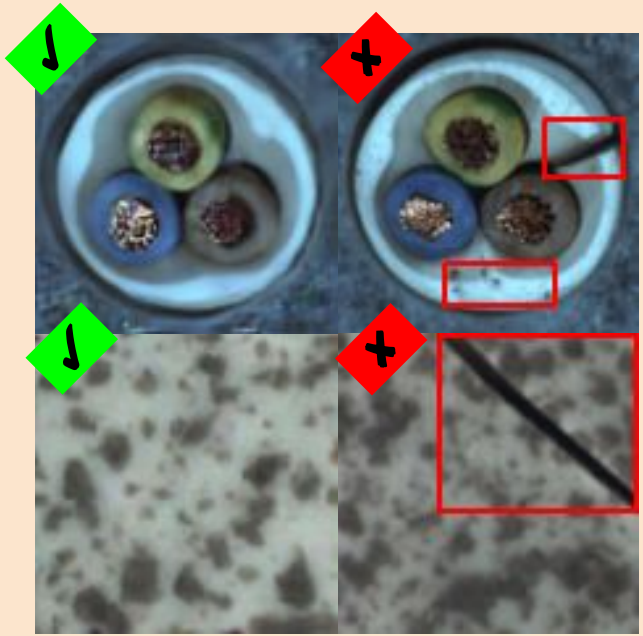

HOW — Fit *generative models* to reconstruct non-anomalous inputs (one-class training, only ✓ are commonly available in the training phase)

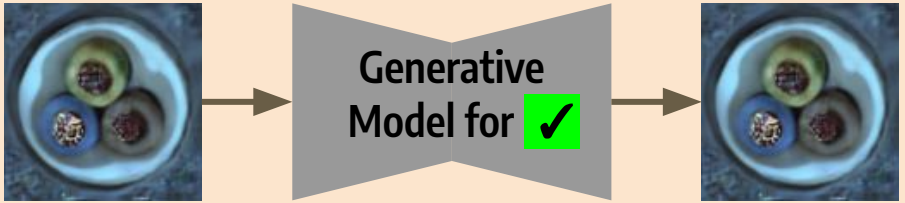

Small difference w.r.t input ⇒ ✓

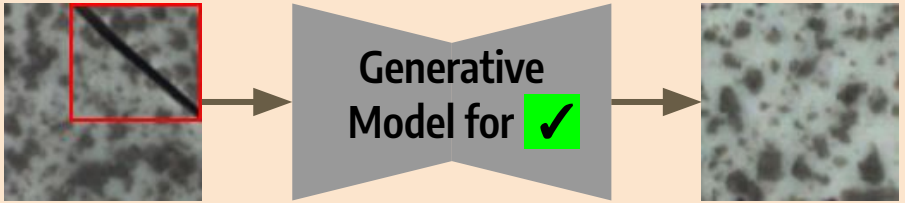

Larger difference w.r.t input (defects are missing) ⇒ ✗

The State of the Art

Generative models in reconstruction-based AD often belong to **GAN** or **AE** families.

Two main classes of methodology:

- **Iterative methods** — reconstructions are optimized via multiple iterations for each input
  - <u>better reconstructions, but very expensive</u>
  - Es: AnoGAN, VAE-grad

- **Single-pass methods** — reconstruction obtained in a single forward pass (often an Encoder-Decoder)
  - <u>more efficient, less precise reconstructions</u>
  - Es: EGBAD (BiGAN), $AE_{L2/SSIM}$, AVID, LSA, etc.

Among single-pass methods, the two commonly approaches adopted are:

- GANs

+ realistic outputs
— (s)low control on reconstruction

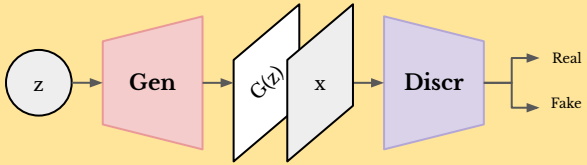

- AutoEncoders

+ preserve alignment w.r.t input
— blurry reconstructions

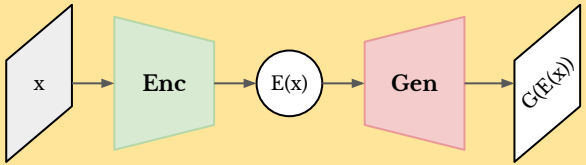

CBiGAN: Consistency BiGAN Model

We propose CBiGAN — a combined model.

+ **generalizes** both Bidirectional GANs ($\alpha = 0$) and AutoEncoders ($\alpha = 1$)

+ α can be tuned to balance the regularization

+ produces **realistic** outputs **fast** that are **consistent** with inputs

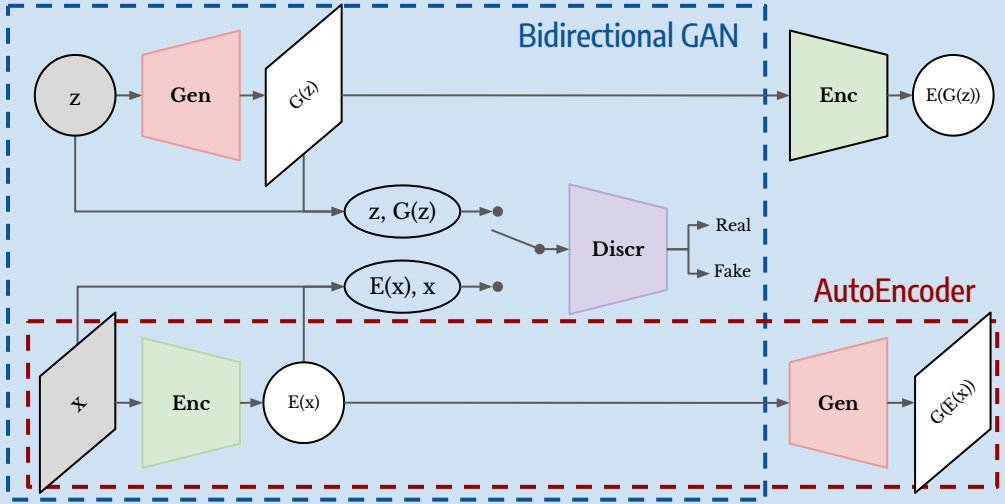

$$\mathcal{L}_{\text{total}} = (1 - \alpha) \cdot \mathcal{L}_{\text{adversarial}} + \alpha \cdot \mathcal{L}_{\text{consistency}}$$

pushes **Gen** and **Enc** to make realistic results, **Discr** to spot fakes

pushes **Gen** and **Enc** to produce outputs consistent with inputs

Evaluation: MVTec-AD Dataset

MVTec-AD — benchmark of visual AD for industrial applications

	Textures	Objects	Overall
Iterative Methods			
AnoGAN	0.54	0.56	0.55
VAE-grad	0.78	**0.76**	**0.77**
Single-pass Methods			
AE_{L2}	0.65	0.74	0.71
AVID	0.67	0.75	0.73
LSA	0.69	0.75	0.73
EGBAD (BiGAN)	0.66	0.58	0.61
CBiGAN (ours)	**0.84**	0.73	0.76

↑ Balanced Accuracy = (TPR + FPR) / 2 when using the Youden threshold. Best among single-pass methods

↓ Area Under the ROC Curve (AuROC)

	Textures	Objects	Overall
AE_{L2}	0.80	**0.74**	0.75
GeoTrans	0.59	0.71	0.67
GANomaly	0.77	0.76	0.76
EGBAD	0.66	0.57	0.60
CBiGAN (ours)	**0.85**	0.73	**0.77**

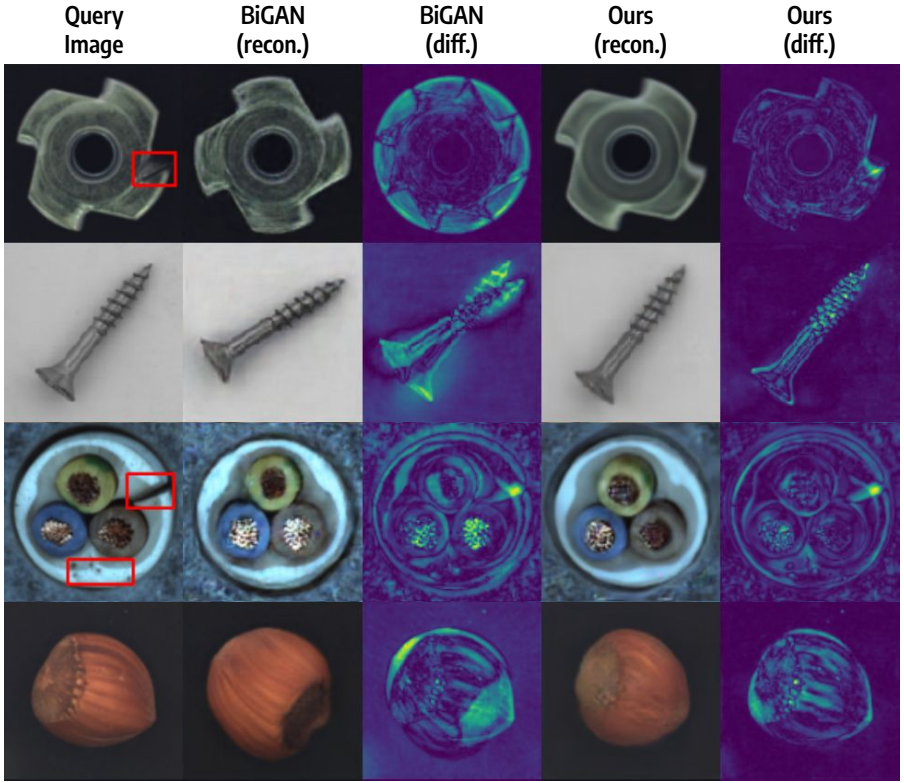

Conclusions and Future Work

- **GOAL** — **detect anomalous images:** assign either ✓ or ✘ to input image
- **CONSTRAINTS** — **one-class AD:** only good images available
- **METHOD** — **reconstruction-based AD:** difference between input and reconstruction
  - Iterative vs Single-Pass methods: effectiveness vs efficiency
- **CONTRIBUTION — we propose Consistency BiGAN for one-class single-pass AD**
  - **combines** and **generalizes** Bidirectional GANs and AutoEncoders
  - provides **reconstruction ability** of the former and **consistency** of the latter
  - efficient (single-pass) with results comparable to iterative methods
- **FUTURE WORK** — address Anomaly Localization (assign either ✓ or ✘ to each pixel)