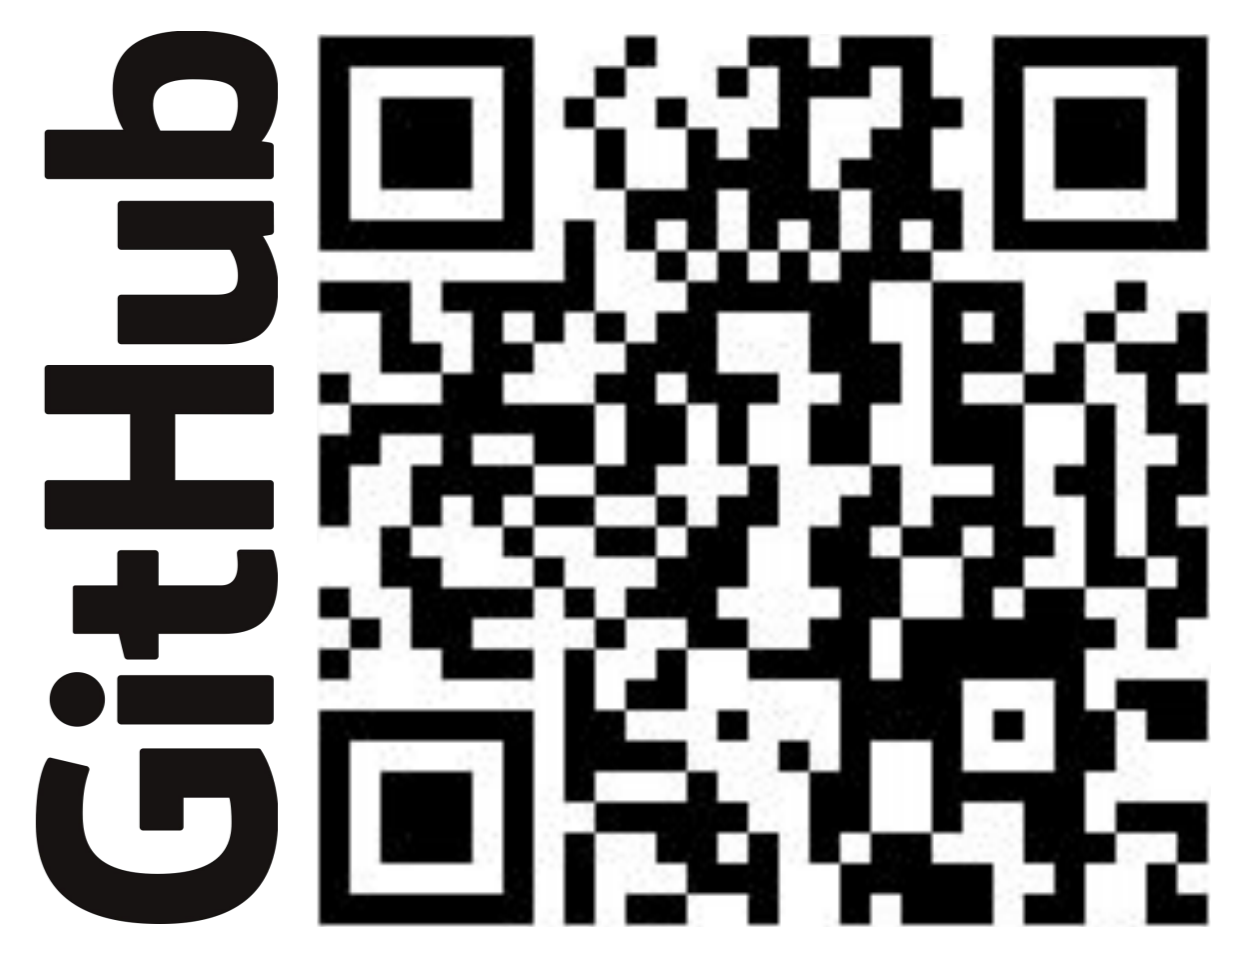


Combining GANs and AutoEncoders for Efficient Anomaly Detection



Fabio Carrara, Giuseppe Amato, Luca Brombin, Fabrizio Falchi, Claudio Gennaro

✉ fabio.carrara@isti.cnr.it

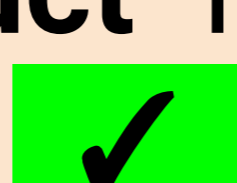
🌐 <https://github.com/fabiocarrara/cbigan-ad>

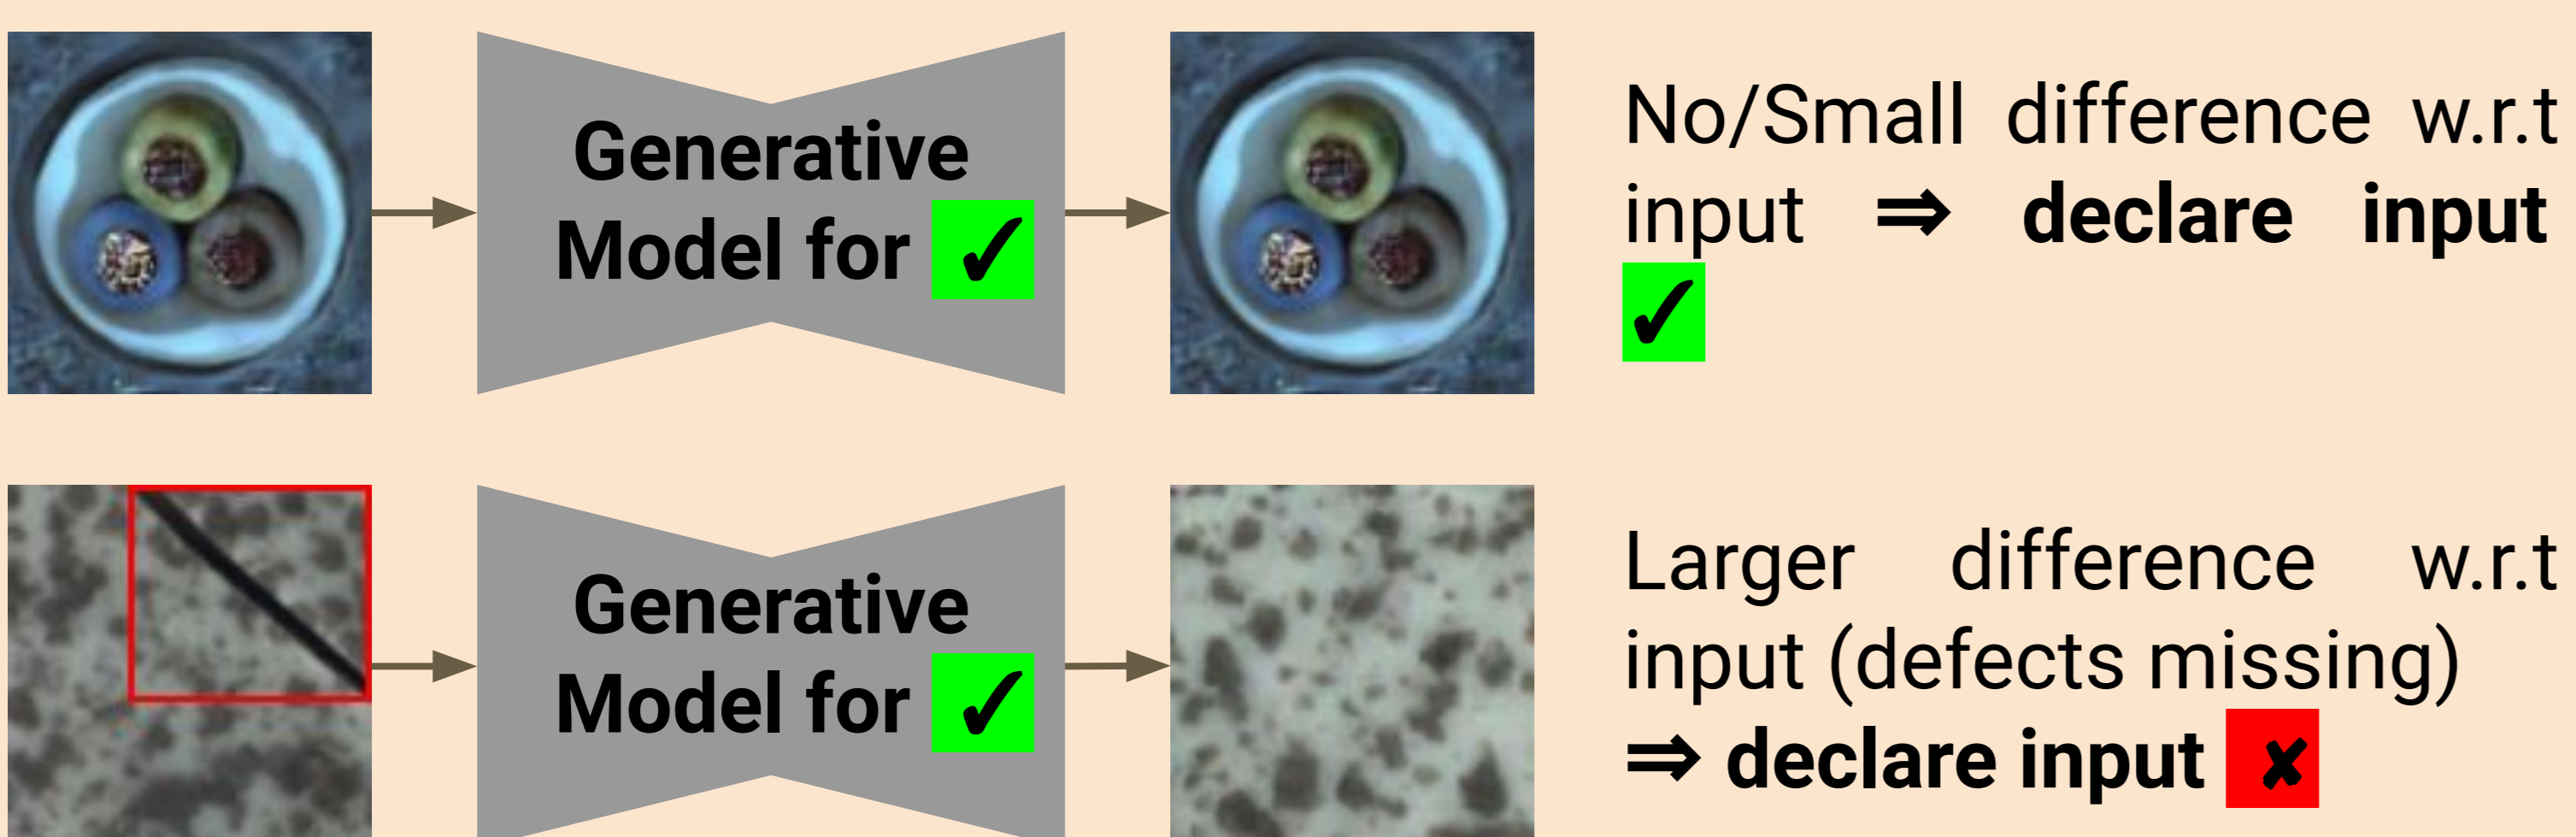


The Problem: Visual Anomaly Detection

GOAL: detect anomalous images (assign either  or  to input image)



HOW: Fit a *generative model* to **reconstruct** non-anomalous inputs (one-class training, only  are commonly available in the training phase)

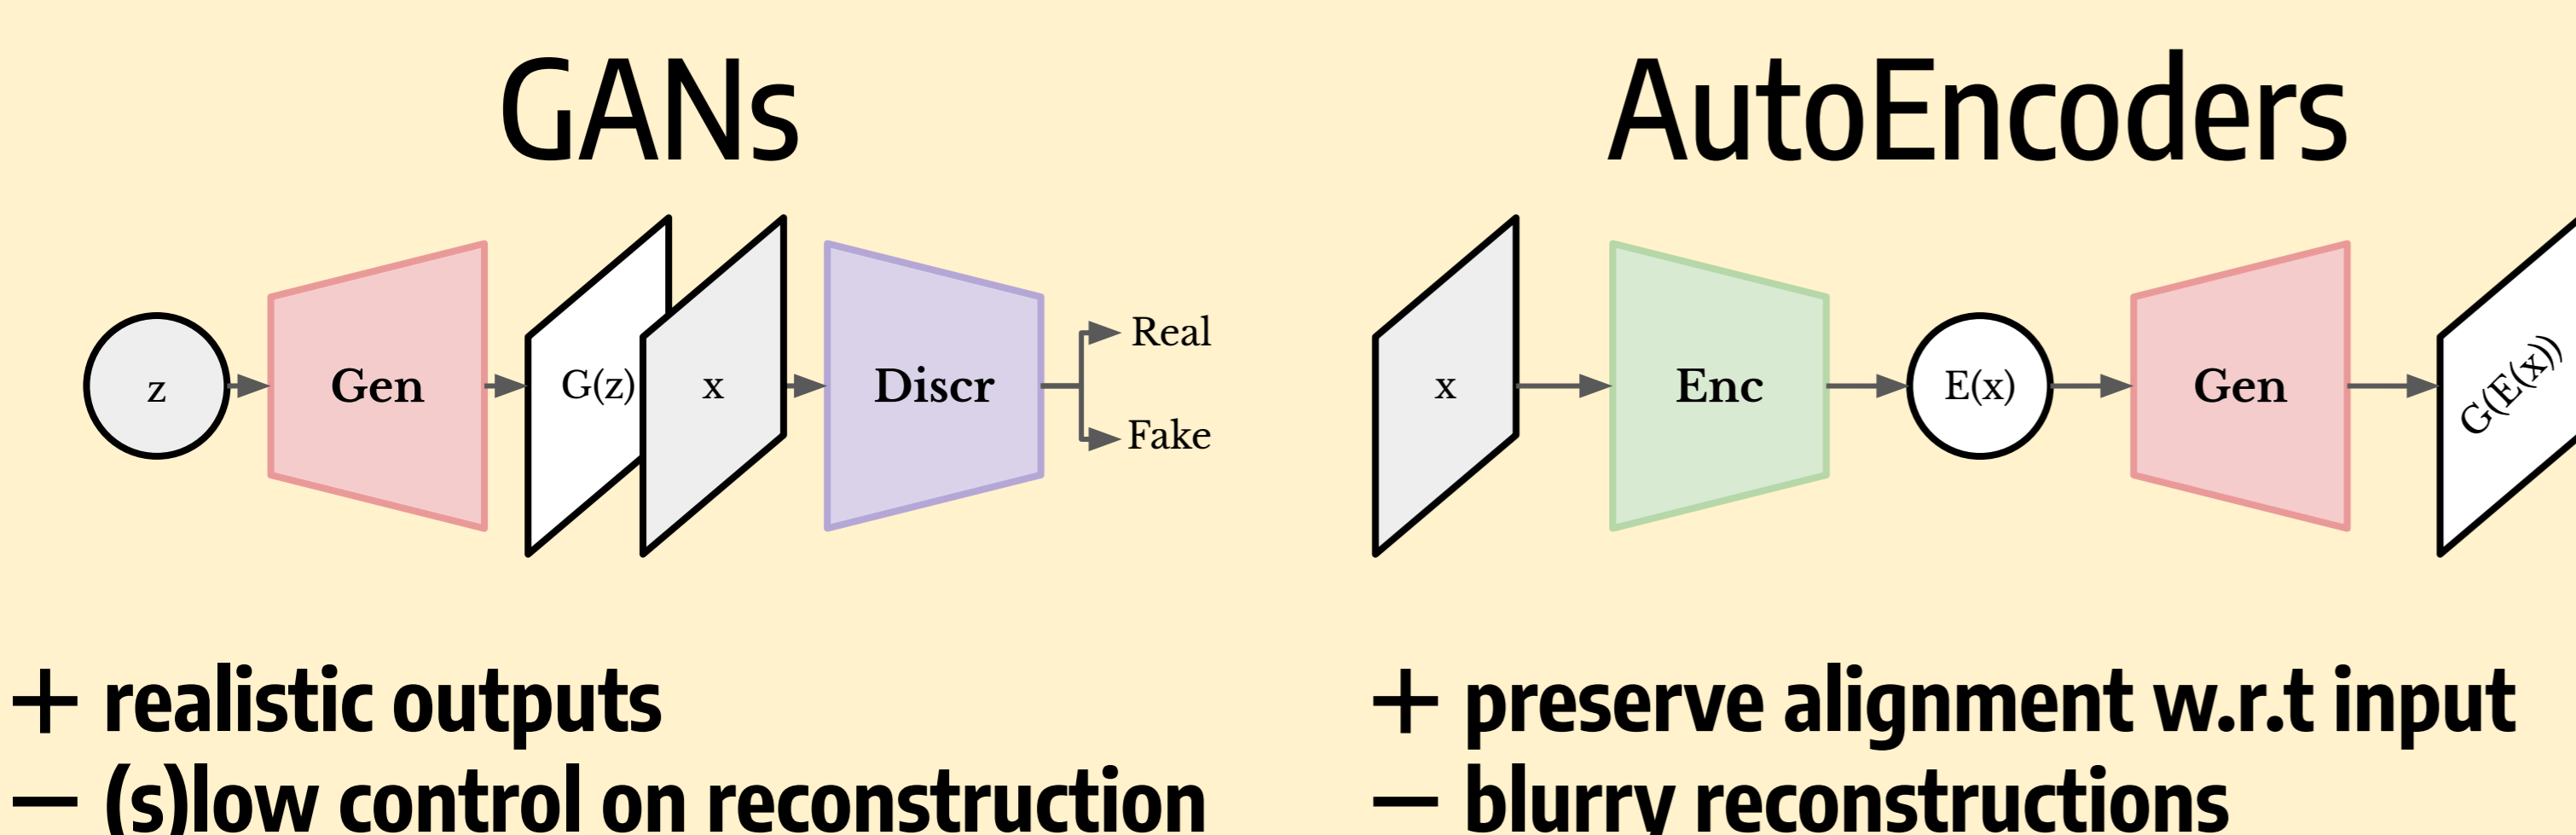


The State of The Art

Generative models used in reconstruction-based methods often are GANs and AE (vanilla & VAEs)

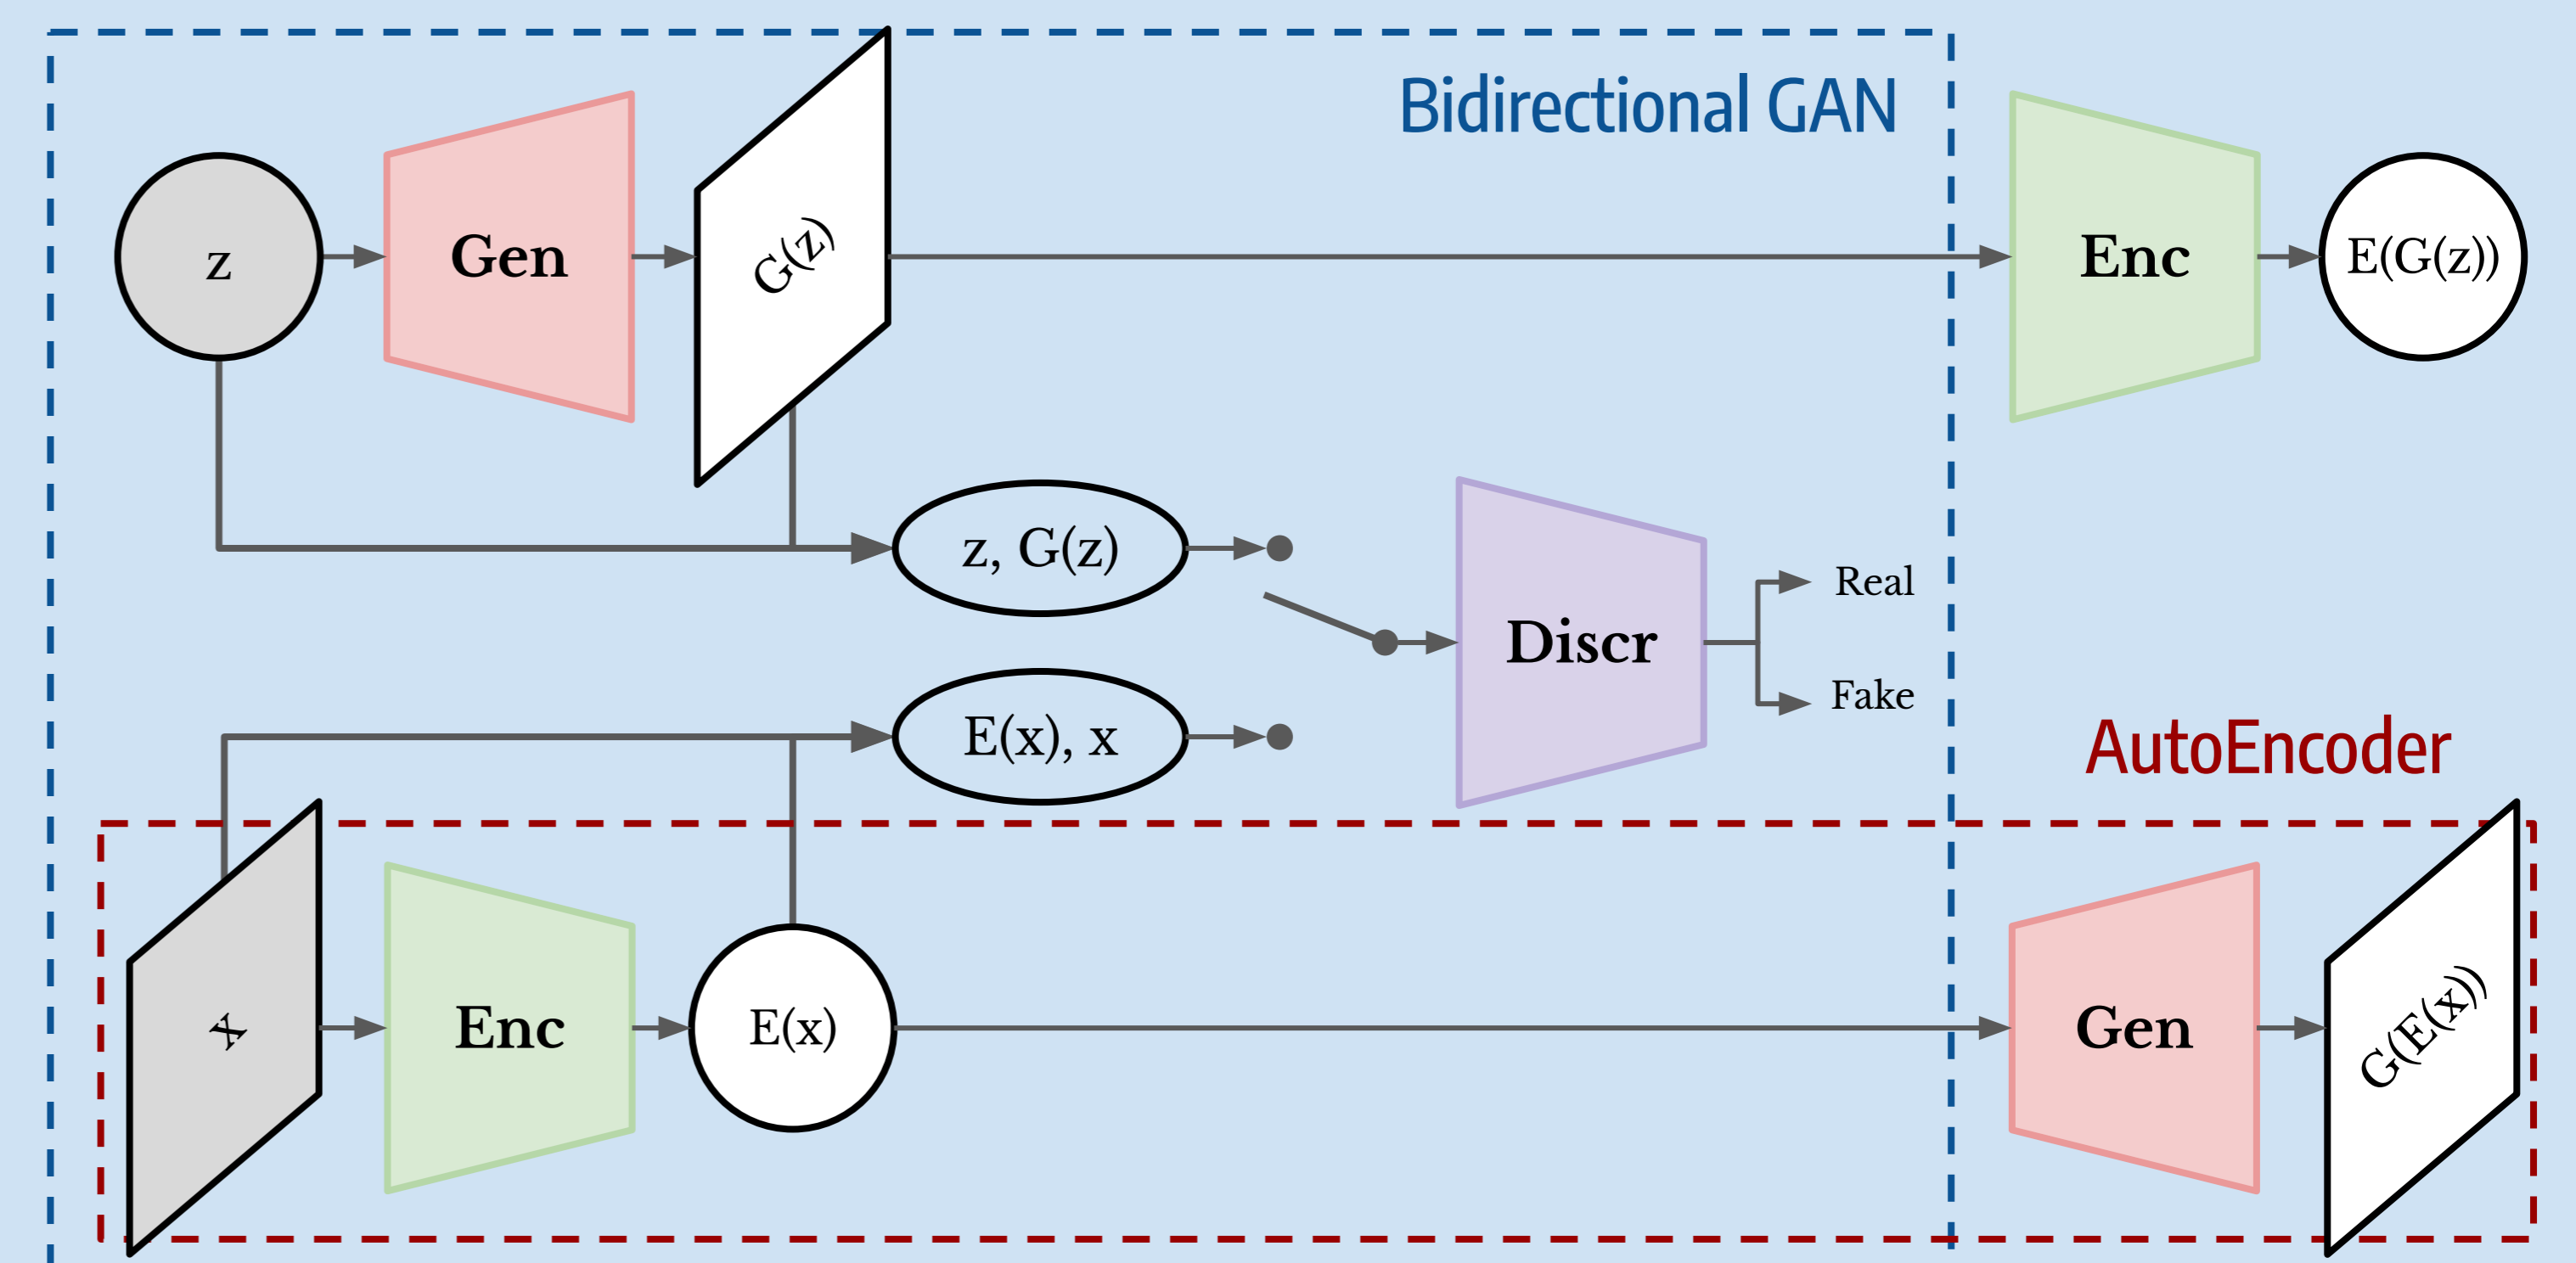
- **Iterative methods:** reconstructions are optimized via multiple iterations for each input
 - better reconstructions, but very expensive
 - Es: AnoGAN, VAE-grad
- **Single-pass methods:** reconstruction obtained in a single forward pass (often an Encoder-Decoder)
 - more efficient, less precise reconstructions
 - Es: EGBAD (BiGAN), AE L2/SSIM, AVID, LSA, etc.

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



CBiGAN: Consistency BiGAN Model

We propose a *combined* model named **CBiGAN**



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

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

pushes Gen and Enc to produce outputs consistent with inputs

- + generalizes both BiGANs ($\alpha = 0$) and AEs ($\alpha = 1$), α can be tuned
- + produces realistic outputs fast that are consistent with inputs

Evaluation: MVTec-AD Dataset

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

	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

↑ Area Under the Curve (AUC)

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

