TURBO: The Swiss Knife of Auto-Encoders

Guillaume Quétant*, Yury Belousov, Vitaliy Kinakh, Slava Voloshynovskiy*

UNIVERSITÉ DE GENÈVE

*guillaume.quetant@unige.ch, *svolos@unige.ch

Introduction

A rigorous interpretation of the concepts and principles behind machine learning methods, unifying as many models as possible. Two complementary methods:

- Bounded Information Bottleneck Auto-Encoder (BIB-AE)
- Two-way Uni-directional Representations by Bounded Optimisation (TURBO)

	BIB-AE	TURBO
Paradigm	**Minimising** the mutual information between the input space and the latent space, while maximising the mutual information between the latent space and the output space **One-way** encoding Data and latent space distributions are considered **independently**	**Maximising** the mutual information between the input space and the latent space, and maximising the mutual information between the latent space and the output space **Two-way** encoding Data and latent space distributions are considered **jointly**
Targeted tasks	• Data compression, privacy, classification • Representation learning	• Linking relevant modalities • Transcoding/translation between modalities
Advantages	• Theoretical basis for both supervised and unsupervised tasks • Allows for easy sampling	• Interpretable latent space • Seamlessly handles paired, unpaired and partially paired data • The encoder can represent a physical system, while the decoder can represent a learnable model
Drawbacks	• Not suited for data translation • Enforces a distribution for the latent space • Struggles to map discontinuous data distributions to continuous latent space distributions	• More hyperparameters to tune • More modules increases training complexity
Particular cases	VAE, GAN, VAE/GAN	AAE, GAN, pix2pix, SRGAN, CycleGAN, Flows
Related models	InfoVAE, CLUB	ALAE

TURBO

Maximising four lower bounds to mutual information gives rise to the following two-way auto-encoder setup.

Direct

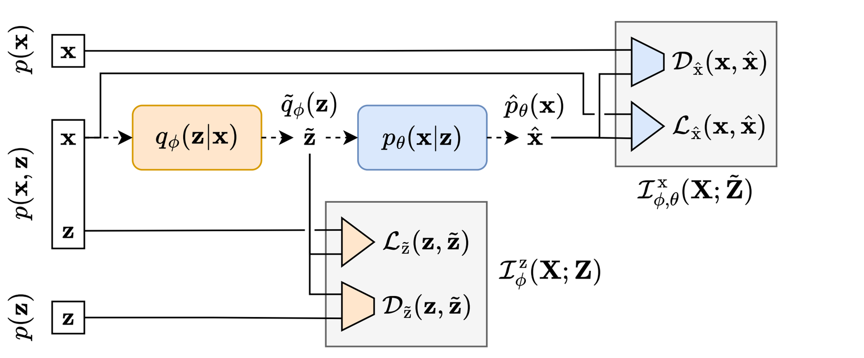

Reverse

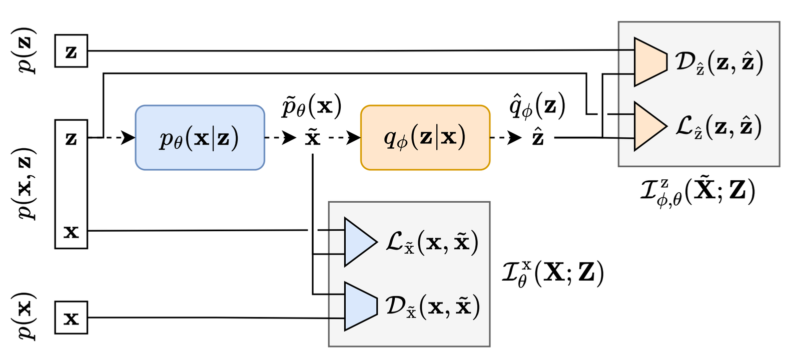

Generalisation of Many Models

Turning on and off the TURBO loss terms lead to many well-known machine learning models.

AAE

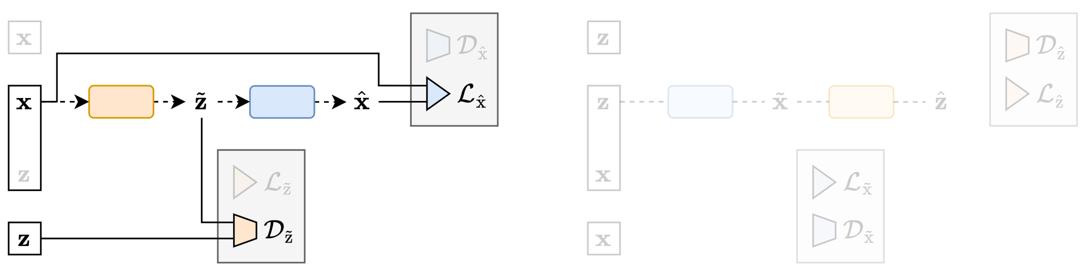

Adversarial Auto-Encoder

- Only one-way setup
- Gaussian latent → no pairs in latent space
- Could add distribution loss in reconstruction space

GAN & WGAN

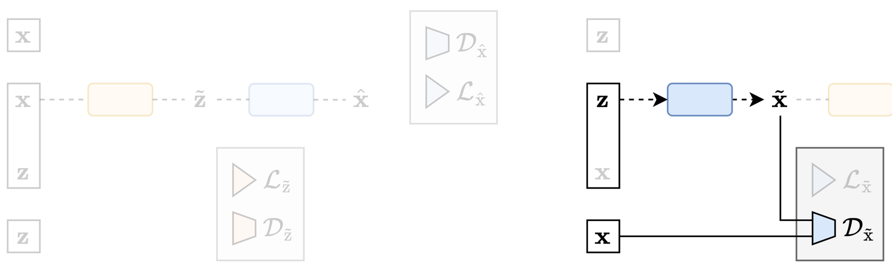

(Wasserstein) Generative Adversarial Network

- No encoder → only decoder
- No pairs

pix2pix & SRGAN

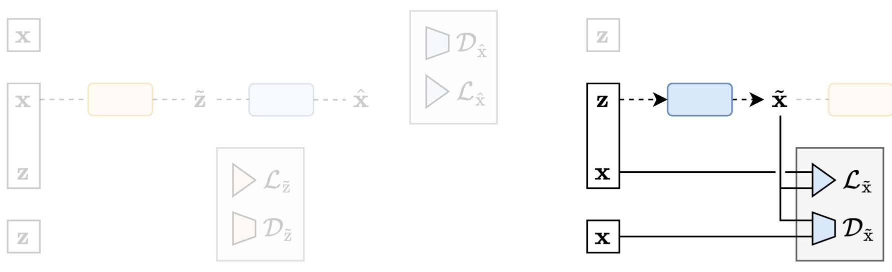

pix2pix & Super-Resolution GAN

- No encoder → only decoder
- But pairs!

CycleGAN

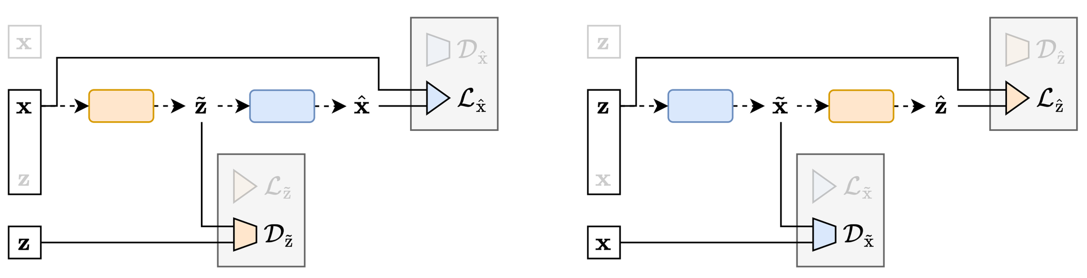

Cycle-Consistent GAN

- Two GANs with cycle consistency
  → Or a two-way AAE!
- X and Z from two sets of data (e.g. horse vs zebra)

Flow

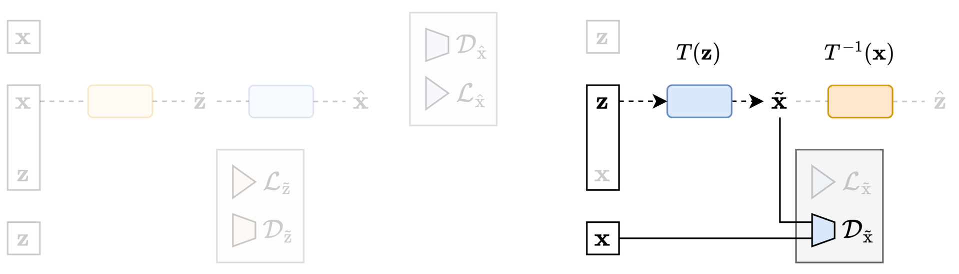

Normalizing Flow

- Encoder = Decoder^{-1}
- It's like GAN!
  → But other parametrisations and KLD computation

ALAE

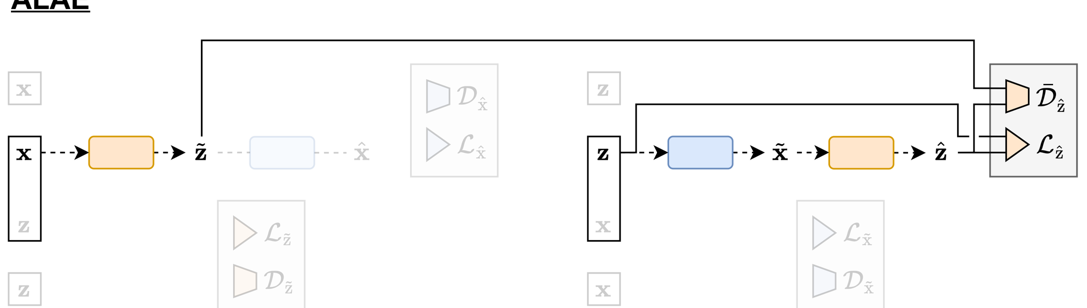

Adversarial Latent Auto-Encoder

- (Slightly) modified TURBO term
- Still maximising mutual information

Physical Latent Space

Auto-encoders with a *virtual* latent space or a *physical* latent space.

- In the virtual setting, the latent variable does not have any physical meaning (e.g. random noise, category, etc.).
- In the physical setting, the latent variable represents a part of the physical observation or measurement chain.

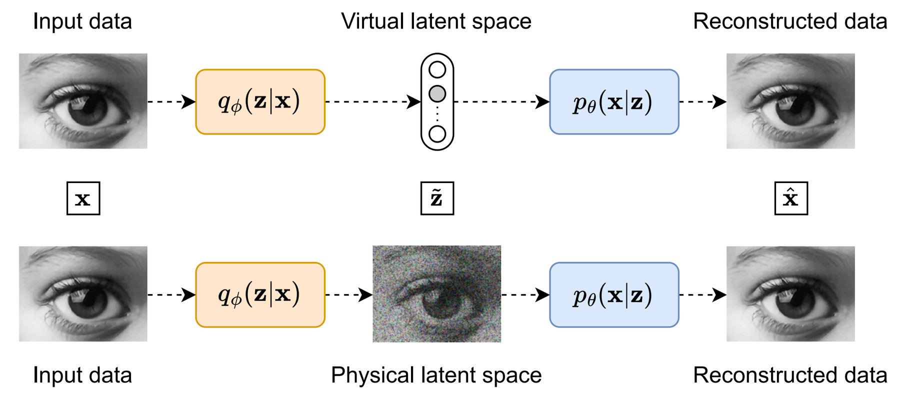

Composition of different latent space settings and several auto-encoder-like networks as internal components of a global auto-encoder architecture.

- The global latent variable can contain both a virtual and a physical parts.
- The global encoder and decoder can be nested auto-encoders with internal latent variables of any kind.

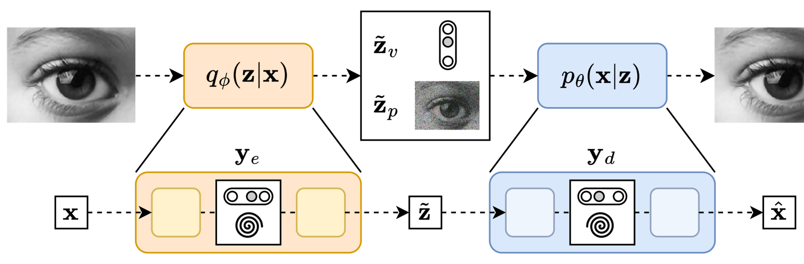

Applications

Highlights of three practical applications where TURBO shows interoperability, superior performance with respect to the models compared, more stable and efficient training.

- The top row shows a high-energy physics example, where particles with given four-momenta are created in a collider experiment and detected by a detector.
- The middle row shows a galaxy imaging example, where two pictures of the same portion of the sky are taken by two different telescopes.
- The bottom row shows a counterfeiting detection example, where a digital template is acquired by a phone camera.

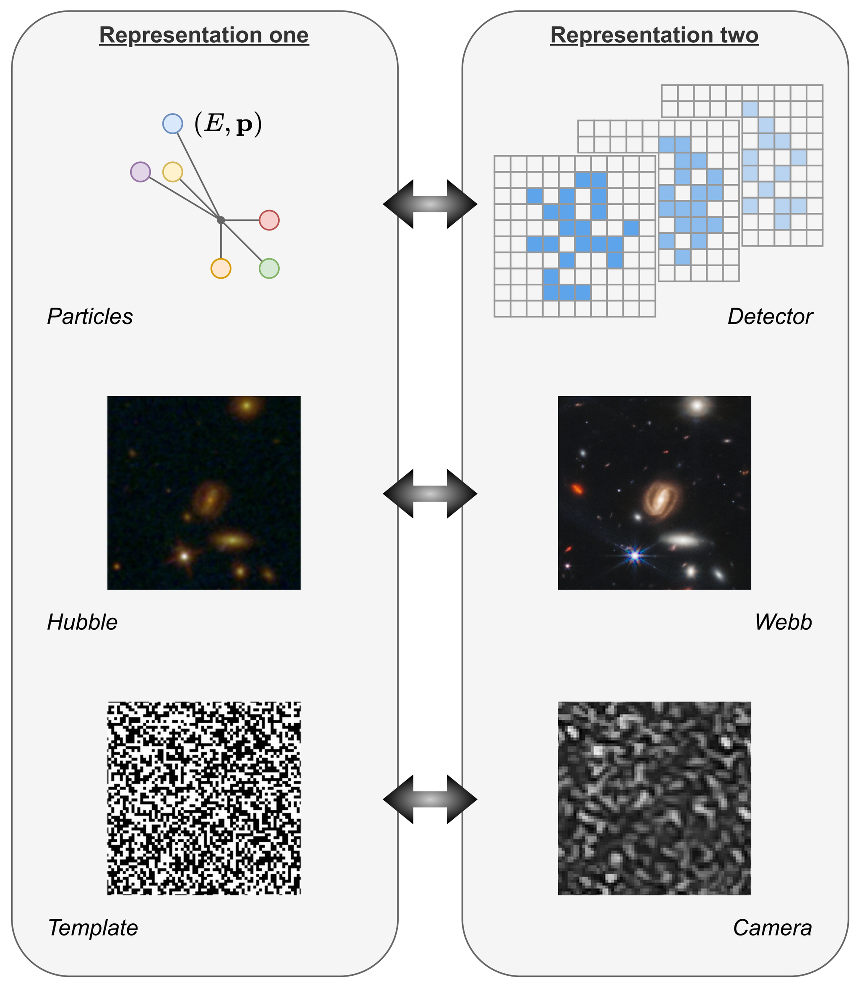

TURBO in High-Energy Physics: Turbo-Sim

Transform the real four-momenta of a set of particles created by the collision of two protons in a collider experiment into the observed four-momenta of the particles captured by detectors, and vice versa. A clever interpretation of the problem is to think of the real and the observed spaces as two different representations of the same physical system.

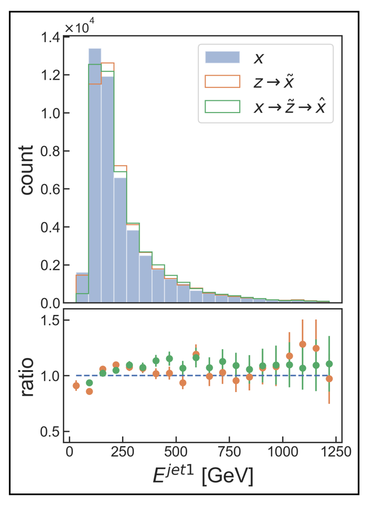

The Kolmogorov–Smirnov distance between the original data simulation and samples generated by the model. The energy of a real particle, a b-quark and the energy of the leading jet are shown, as well as the invariant mass of the top-quarks pair, which are unstable particles decaying into the real ones before flying through the detectors to be observed.

	Z space	X space	Rec. space
Model	E^b	E^{jet1}	m_{tt}
OTUS	**2.76**	5.75	15.8
Turbo-Sim	3.96	**4.43**	**2.97**

TURBO in Astronomy: Hubble-to-Webb

Using TURBO as an image-to-image translation framework to generate simulated images of the James Webb Space Telescope from observed images of the Hubble Space Telescope and vice versa.

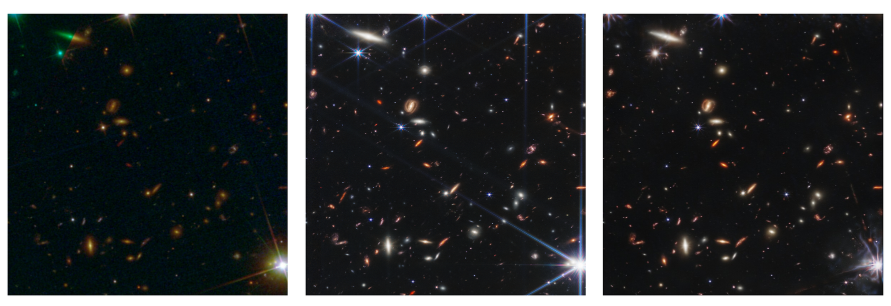

Hubble Webb TURBO

Model	MSE ↓	SSIM ↑	PSNR ↑	LPIPS ↓	FID ↓
CycleGAN	0.0097	0.83	20.11	0.48	128.1
pix2pix	**0.0021**	**0.93**	**26.78**	0.44	54.58
TURBO	0.0026	0.92	25.88	**0.41**	**43.36**

TURBO in Anti-Counterfeiting: Digital Twin

Accurately estimate the complex stochastic process of printing and generate predictions of how a digital template would appear once printed, as well as reversing the process and predict the original digital template from the printed one.

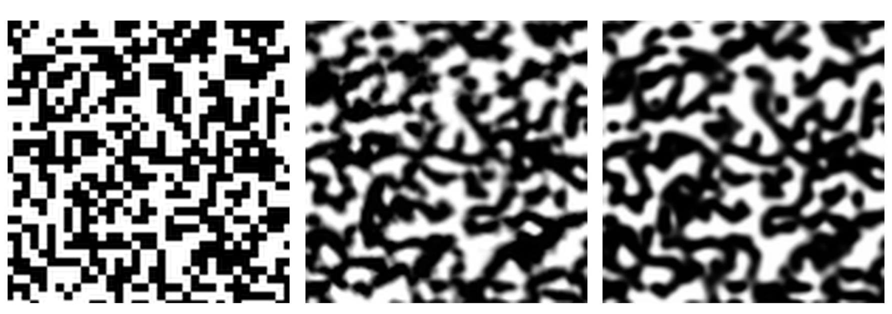

Digital Printed TURBO

Model	$FID_{x\to z}$ ↓	$FID_{z\to x}$ ↓	Hamming ↓	MSE ↓	SSIM ↑
w/o processing	*304*	*304*	*0.24*	*0.18*	*0.48*
CycleGAN	3.87	**4.45**	0.15	0.05	0.73
pix2pix	3.37	8.57	0.11	0.05	0.76
TURBO	**3.16**	6.60	**0.09**	**0.04**	**0.78**

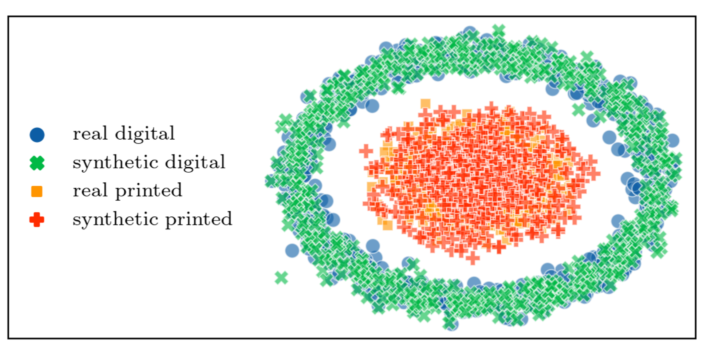

UMAP visualisation of synthetically generated templates superimposed on the corresponding real counterparts

References

1. Quétant G., Belousov Y., Kinakh V., Voloshynovskiy S.
TURBO: The Swiss Knife of Auto-Encoders.
Entropy, 2023, 25, 1471. https://doi.org/10.3390/e25101471

2. Quétant G., Drozdova M., Kinakh V., Golling T., Voloshynovskiy S.
Turbo-Sim: A generalised generative model with a physical latent space.
In Proceedings of the Workshop on Machine Learning and the Physical Sciences, NeurIPS, 2021. https://arxiv.org/abs/2112.10629

3. Belousov Y., Pulfer B., Chaban R., Tutt J., Taran O., Holotyak T., Voloshynovskiy S.
Digital twins of physical printing-imaging channel.
In Proceedings of the IEEE International Workshop on Information Forensics and Security, 2022. https://doi.org/10.1109/WIFS55849.2022.9975439