Génération d'images

GAN, MAE, DAE, Diffusion

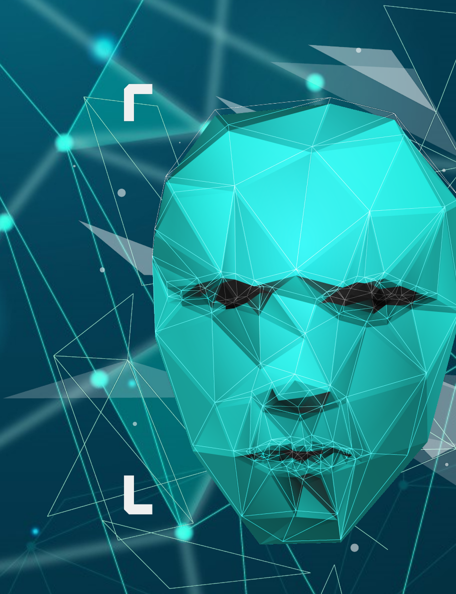

Introduction

- Génération d'image (DallE, Stable Diffusion, Midjourney

- Restauration d'image (colorisation, débruitage…)

- Modification d'image (inpainting, outpainting…)

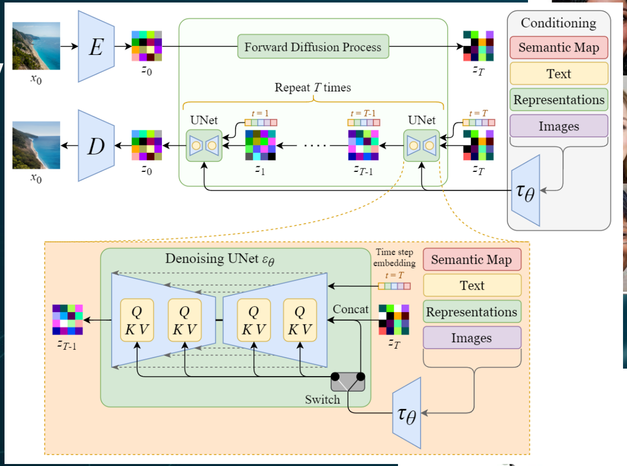

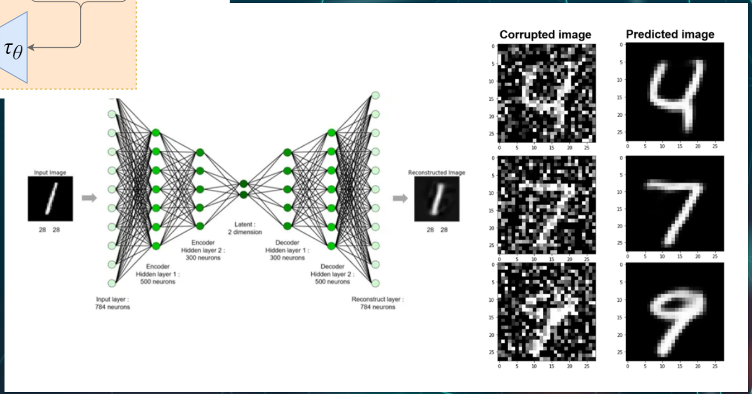

DAE: Denoising Autoencoders (1987)

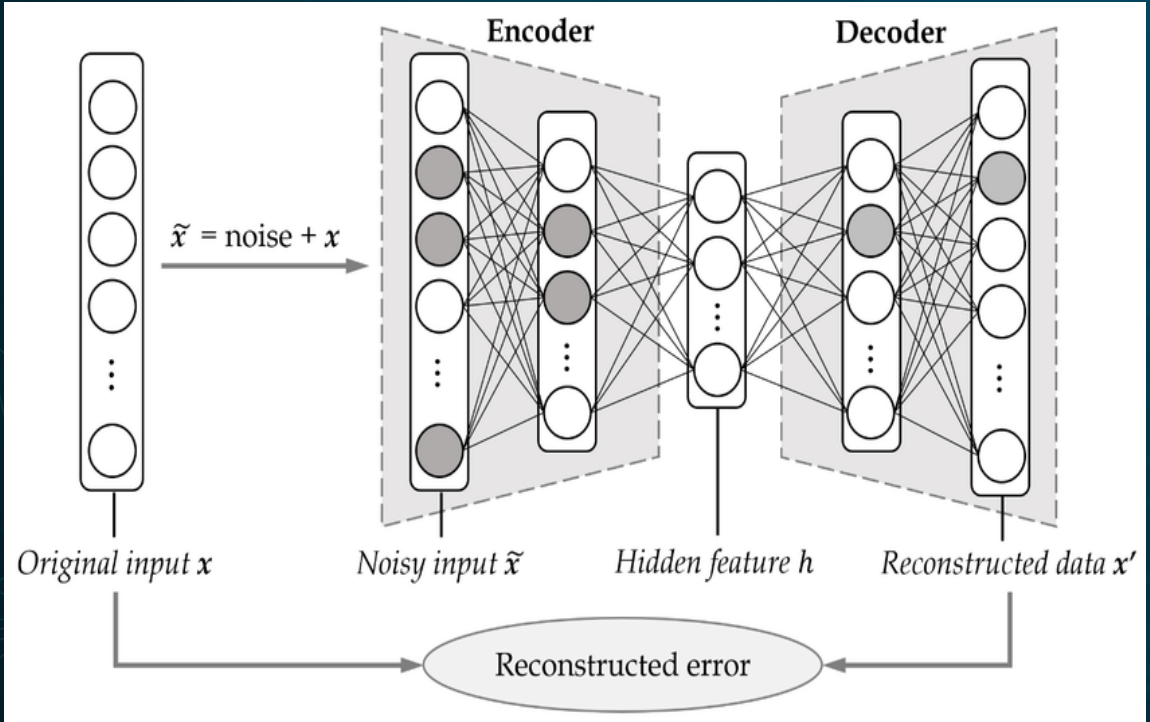

DAE: Denoising Autoencoders (1987)

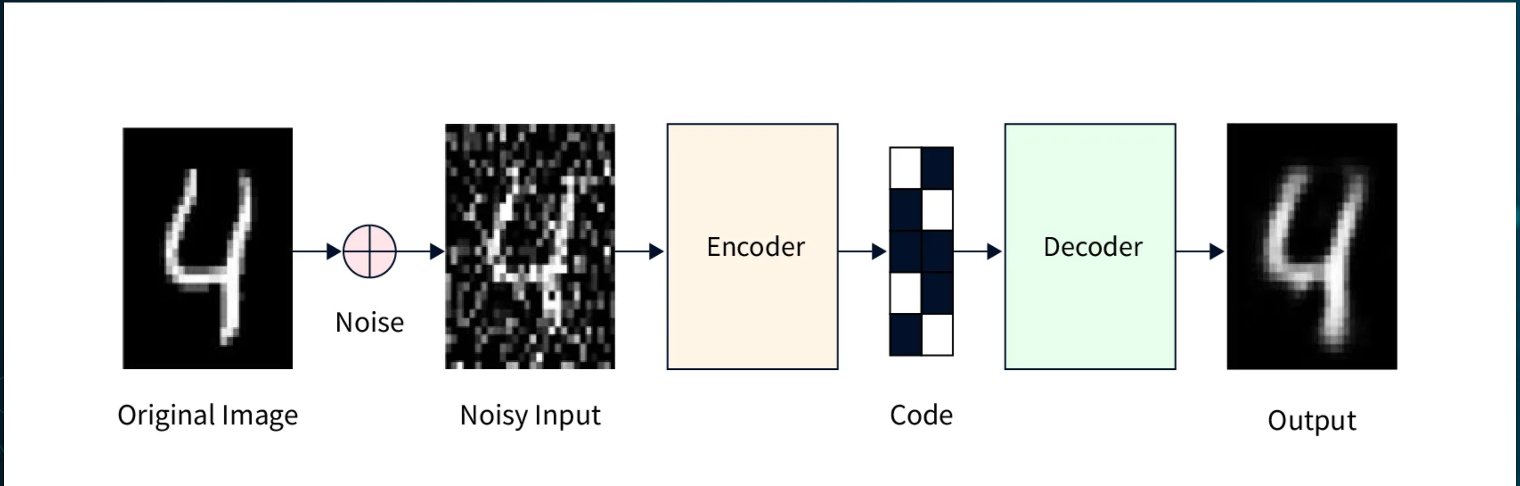

GAN: Generative Adversarial Networks (2014)

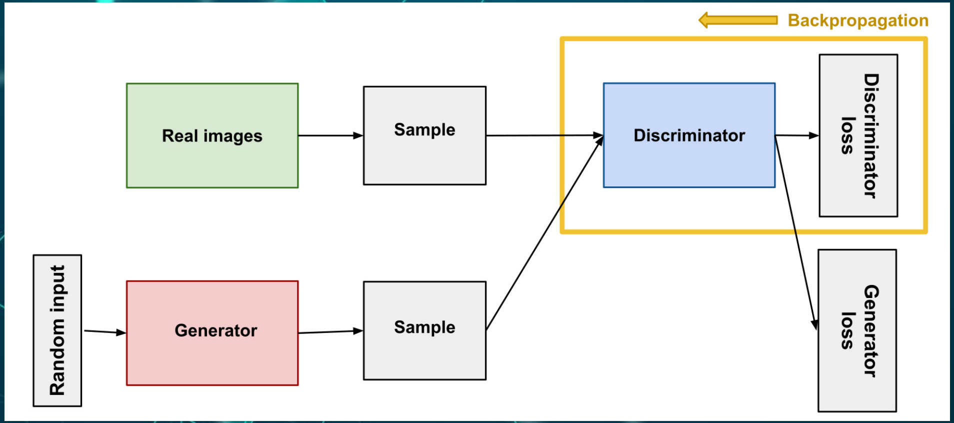

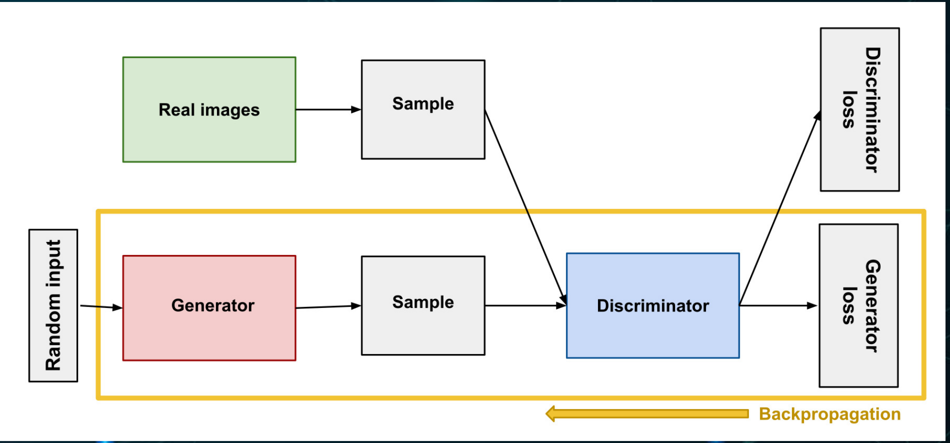

GAN: Generative Adversarial Networks (2014)

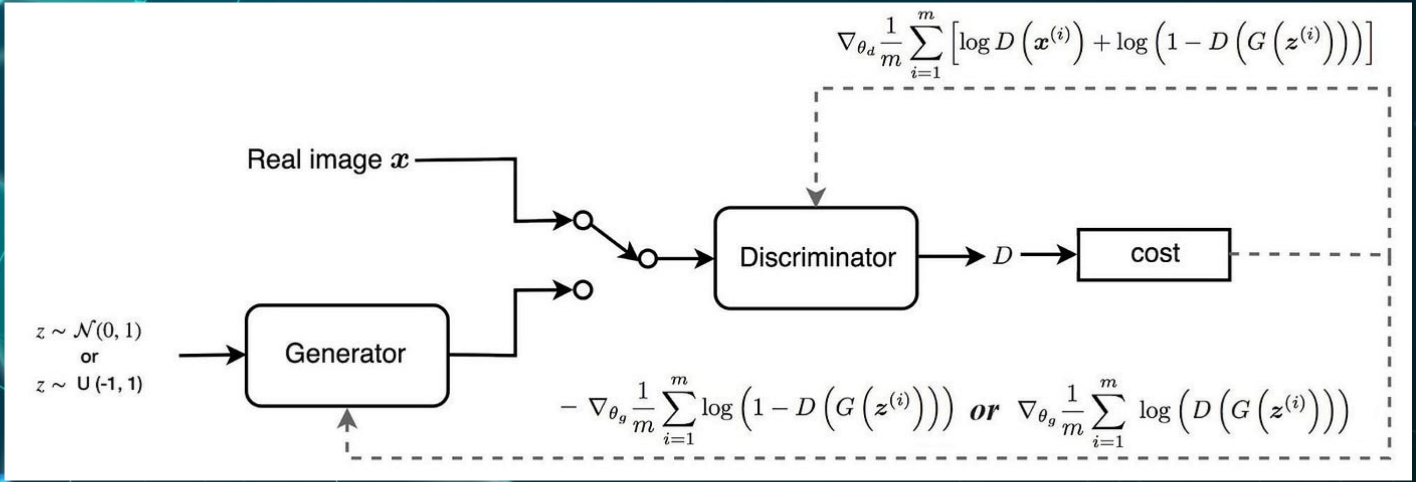

MAE: Masked Autoencoders (2021)

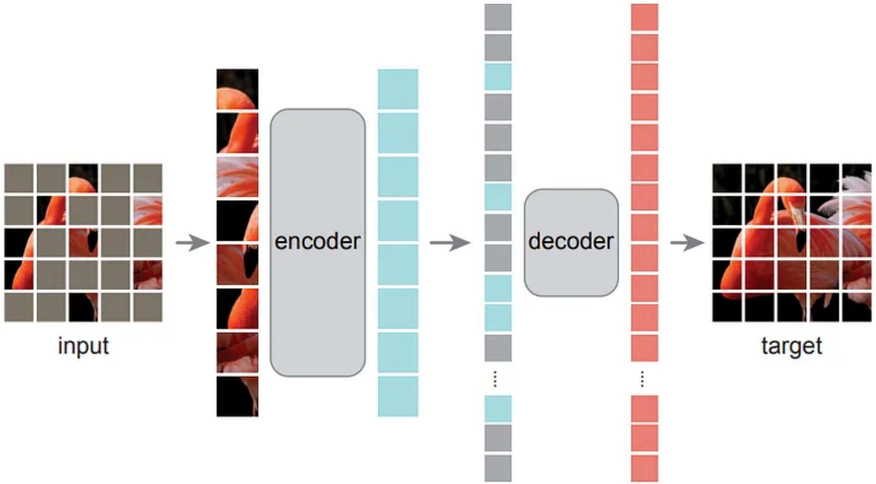

Figure 1. **Our MAE architecture**. During pre-training, a large random subset of image patches (*e.g*., 75%) is masked out. The encoder is applied to the small subset of *visible patches*. Mask tokens are introduced *after* the encoder, and the full set of encoded patches and mask tokens is processed by a small decoder that reconstructs the original image in pixels. After pre-training, the decoder is discarded and the encoder is applied to uncorrupted images (full sets of patches) for recognition tasks.

Masked Autoencoders Are Scalable Vision Learners

Kaiming He[*,†] Xinlei Chen[*] Saining Xie Yanghao Li Piotr Dollár Ross Girshick

[*]equal technical contribution [†]project lead

Facebook AI Research (FAIR)

MAE: Masked Autoencoders (2021)

Figure 3. Example results on COCO validation images, using an MAE trained on ImageNet (the same model weights as in Figure 2). Observe the reconstructions on the two right-most examples, which, although different from the ground truth, are semantically plausible.

Avantages:

- Entrainement de décodeurs large rapide.

- Grande robustesse des représentations, choix des masques.

- Self supervised learning + Transfer learning

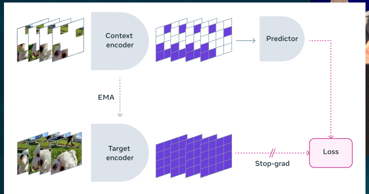

Diffusion Models

Débruitage successif analogue de la diffusion thermique.

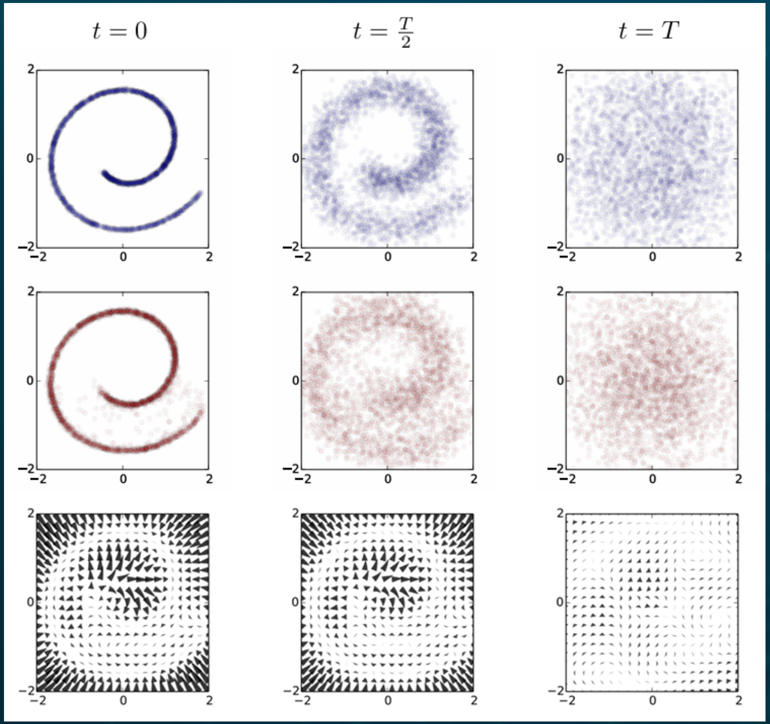

Exemple original de diffusion (Sohl-Dickstein, Jascha, et al. 2015)

Diffusion Models

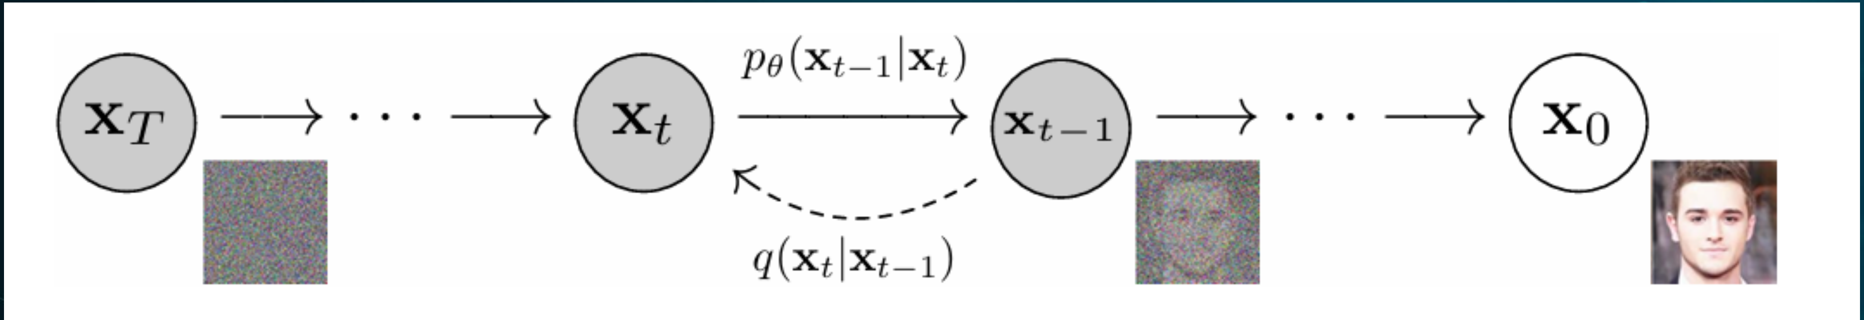

Chaîne de Markov d’un modèle Denoising Diffusion Probabilistic (Ho, Jonathan, et al. 2020)

Fonction de perte :

$$\mathbb{E}_q\left[-\log p(\mathbf{x}_T) - \sum_{t\geq 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)}{q(\mathbf{x}_t|\mathbf{x}_{t-1})}\right]$$

Diffusion Models

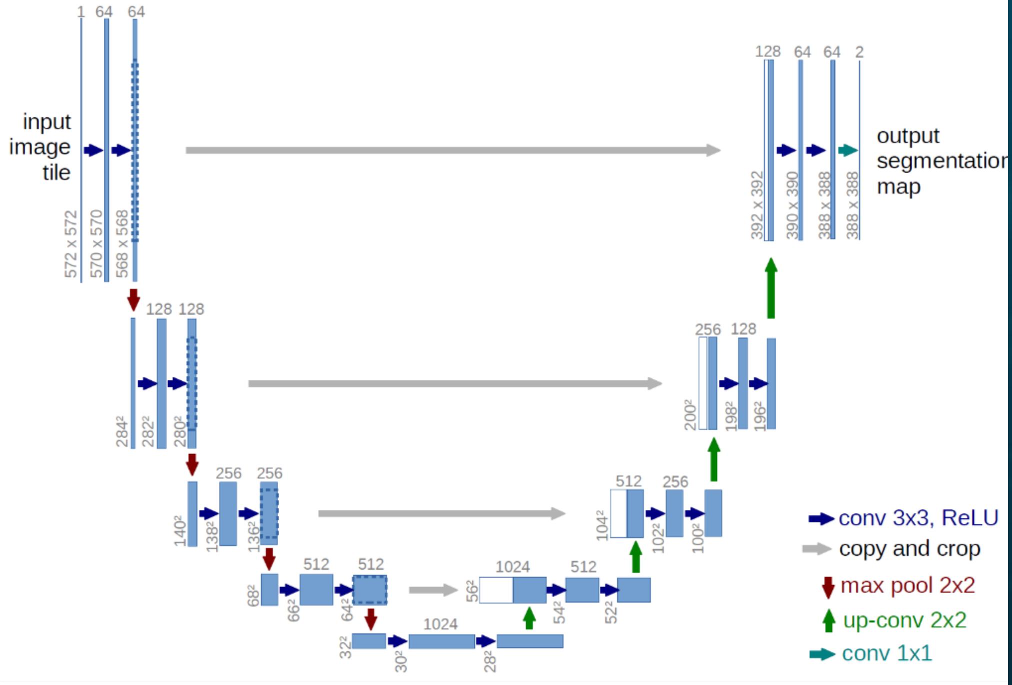

Architecture U-Net (Ho, Jonathan, et al. 2020)

Diffusion Models: Applications

Image originale

Variation
denoise 0.6

Style Picasso
denoise 0.6

Inpainting

Sources

https://developers.google.com/machine-learning/gan/discriminator?hl=fr

https://arxiv.org/abs/2111.06377

Generative Adversarial Networks,Ian J. Goodfellow and Jean Pouget-Abadie and Mehdi Mirza and Bing Xu and David Warde-Farley and Sherjil Ozair and Aaron Courville and Yoshua Bengio,year={2014}

Masked Autoencoders Are Scalable Vision Learners. Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, Ross Girshick

L'algorithme derrière Midjourney : Comprendre les modèles de diffusion. www.youtube.com, https://www.youtube.com/watch?v=IvMGTteb3EI. Consulté le 12 mars 2024.

« Introduction to Diffusion Models for Machine Learning ». News, Tutorials, AI Research, 12 mai 2022, https://www.assemblyai.com/blog/diffusion-models-for-machine-learning-introduction/.

Sohl-Dickstein, Jascha, et al. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. arXiv:1503.03585, arXiv, 18 novembre 2015. arXiv.org, https://doi.org/10.48550/arXiv.1503.03585.

Ho, Jonathan, et al. Denoising Diffusion Probabilistic Models. arXiv:2006.11239, arXiv, 16 décembre 2020. arXiv.org, https://doi.org/10.48550/arXiv.2006.11239.

THANKS!

CREDITS: This presentation template was created by **Slidesgo**, including icons by **Flaticon**, and infographics & images by **Freepik**.