MICAD2022
Medical Imaging and Computer-Aided Diagnosis

MICCAI
MICCAI Society Endorsed Event

The University of Edinburgh

Royal Academy of Engineering

Canon
CANON MEDICAL

ellis

The Alan Turing Institute

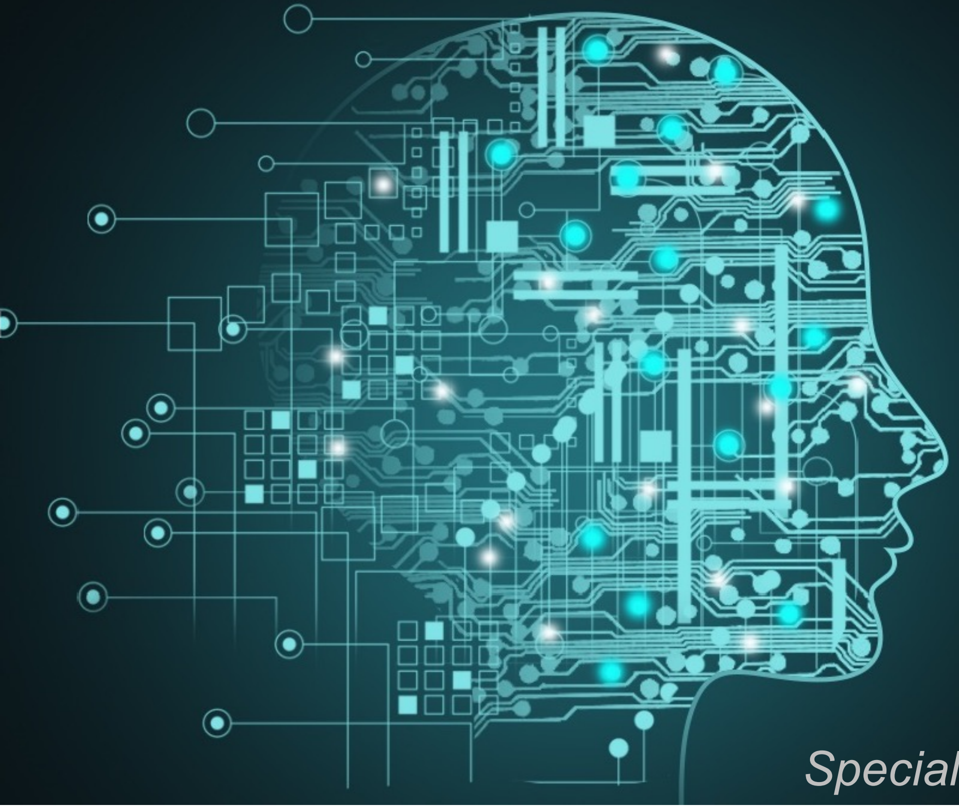

Diffusion Models in Medical Imaging and Analysis. Hype or Hope?

Canon Medical / Royal Academy of Engineering Research Chair in Healthcare AI
Professor in Machine Learning and Computer Vision
Turing Fellow & ELLIS Fellow

https://vios.science/tutorials/diffusion

Special thanks to Pedro Sanchez and all colleagues who shared slides

Organised by: University of Leicester

Supported by: Spectrum Academy, Springer

At VIOS we do interdisciplinary AI

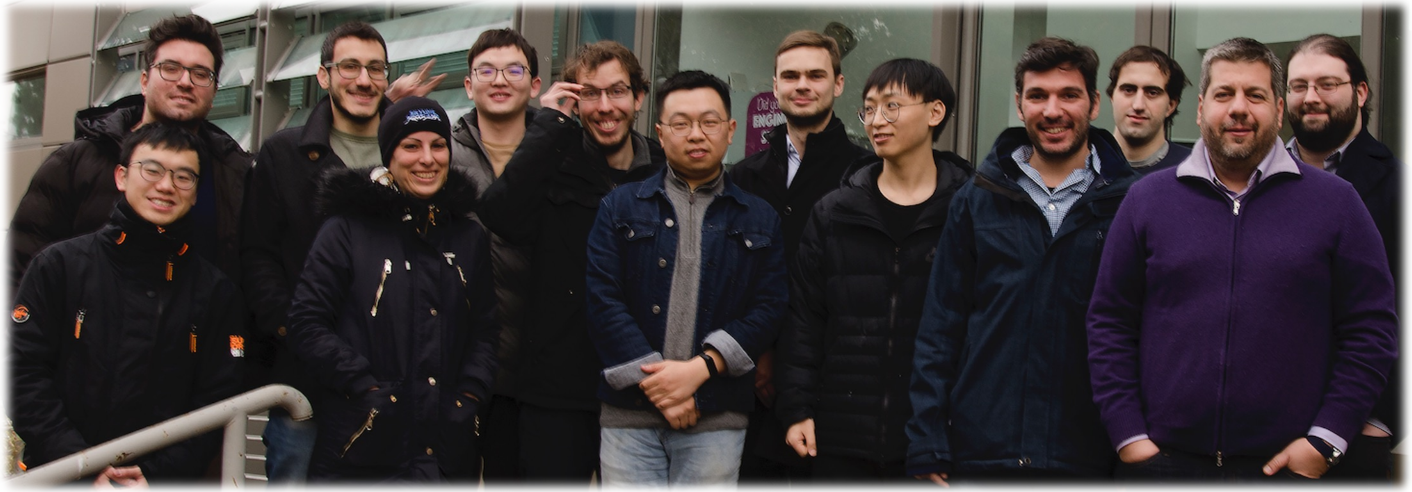






Text2Image

Stable Diffusion (Stability AI)

"Holy festival of colors with dancing people in India"

Imagen (Google Brain)

"A photo of a Persian cat wearing a cowboy hat and black leather jacket riding a bike on a beach"

Dall-E 2 (OpenAI)

"An astronaut riding a horse in a photorealistic style"

1. Saharia et al (2022). Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. arXiv:2205.11487
2. Ramesh et al (2022). Hierarchical text-conditional image generation with clip latents. arXiv:2204.06125
3. Rombach et al (2022). High-Resolution Image Synthesis with Latent Diffusion Models. CVPR

Popularity

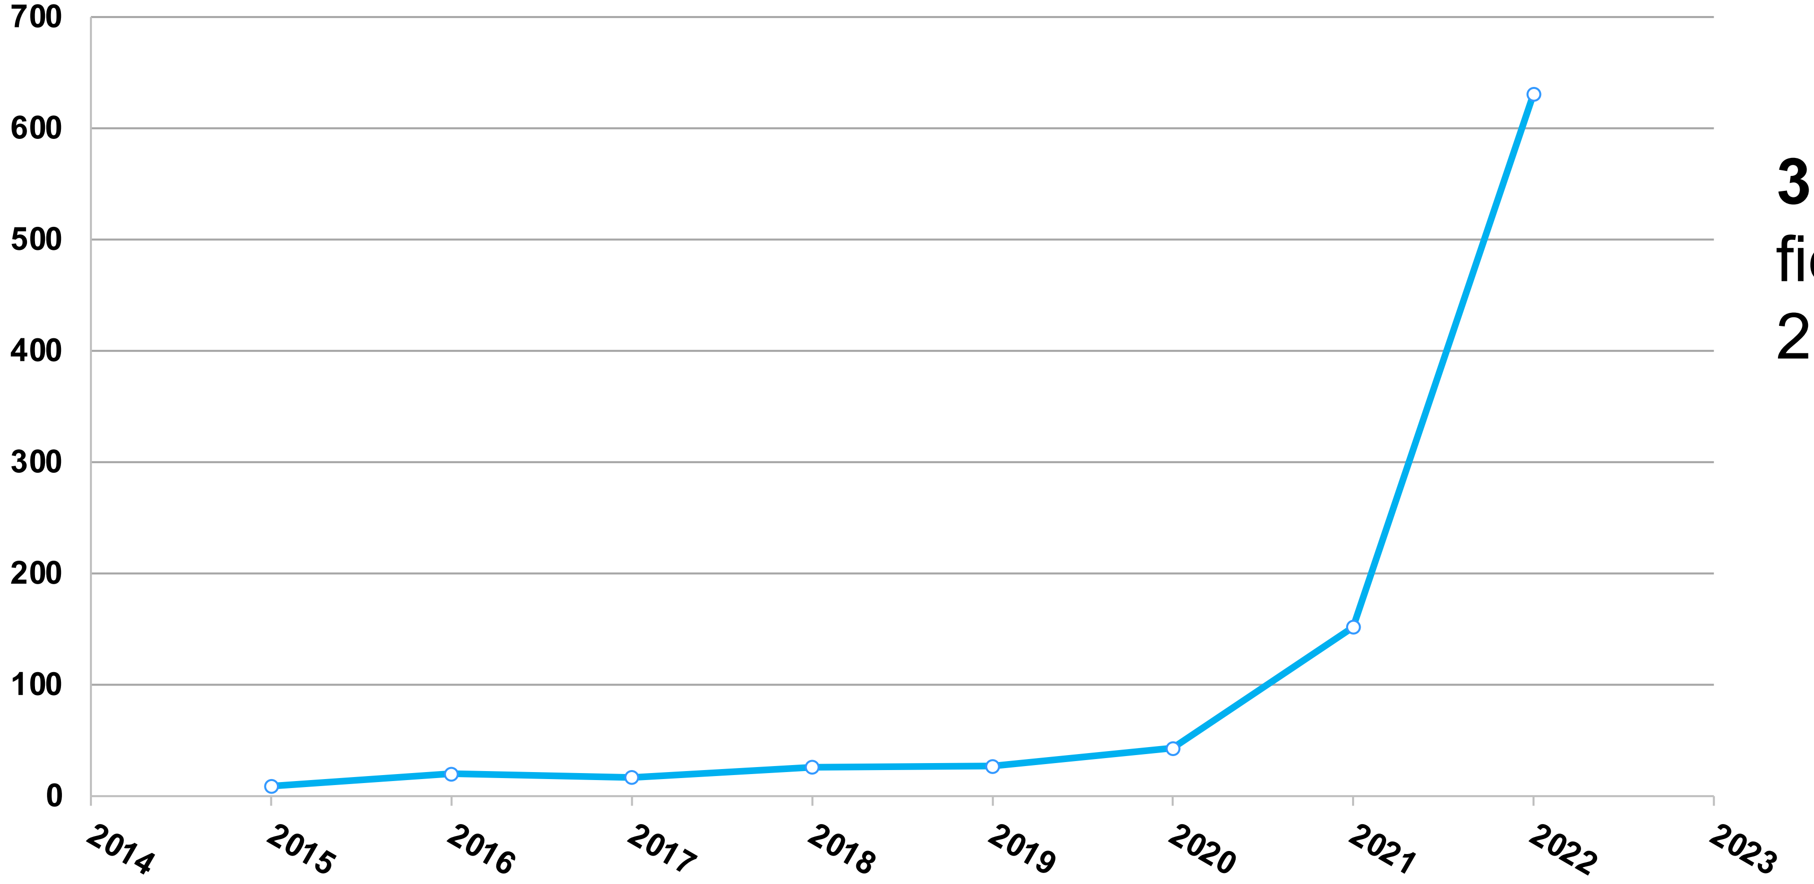

36 in our field since 2021

Source: Google Scholar query: machine learning "diffusion models" source:arxiv

The rest of the talk

- Diffusion models (a brief illustration)
- Diffusion models in medical imaging analysis
- Conclusion

Diffusion models

Generative Models

$x_i \sim p_{data}$

$i \in \{1,2,\dots,|D|\}$

$p_\theta \in \mathcal{M}$

Model family

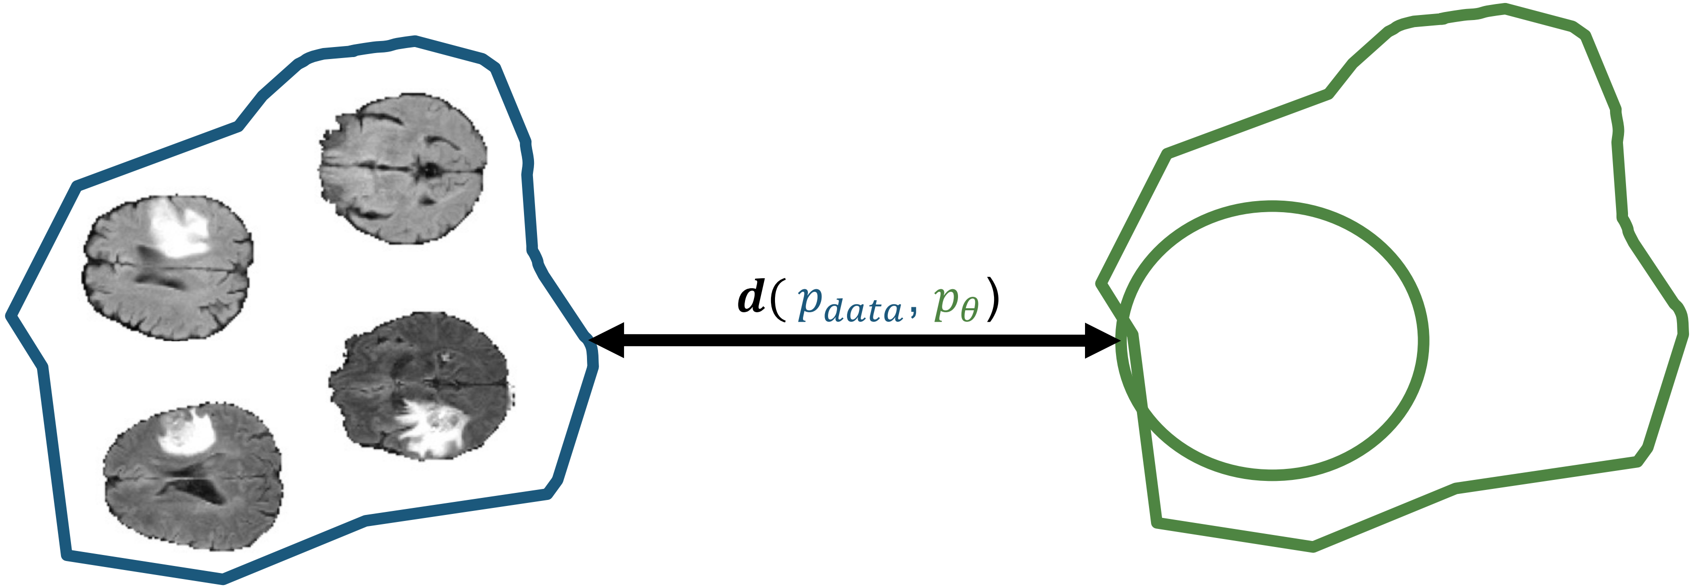

$\mathcal{M}$ = {VAE, GAN, NF, **Diffusion Models**}

https://deepgenerativemodels.github.io/notes/introduction/

Generative models

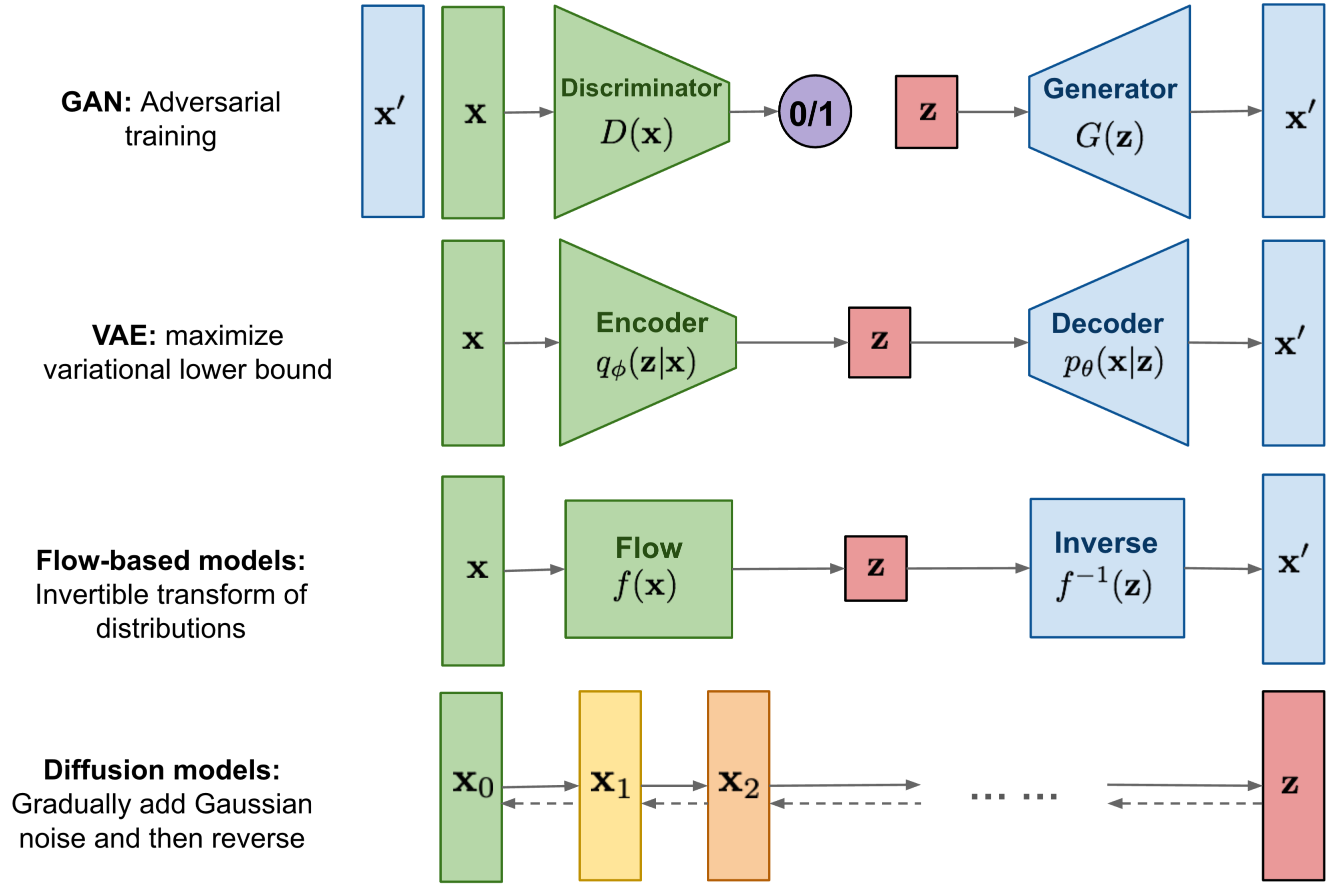

likelihood-based models

Require
- <u>inductive bias</u> to ensure a tractable normalizing constant for likelihood computation; or
- <u>surrogate objectives</u> to approximate ML training.

implicit generative models

- Require <u>adversarial</u> training:
  - notoriously unstable; leading to
  - mode collapse

diffusion models bypass both with neat tricks

Figure by Lilian Weng
https://lilianweng.github.io/posts/2021-07-11-diffusion-models/

The score

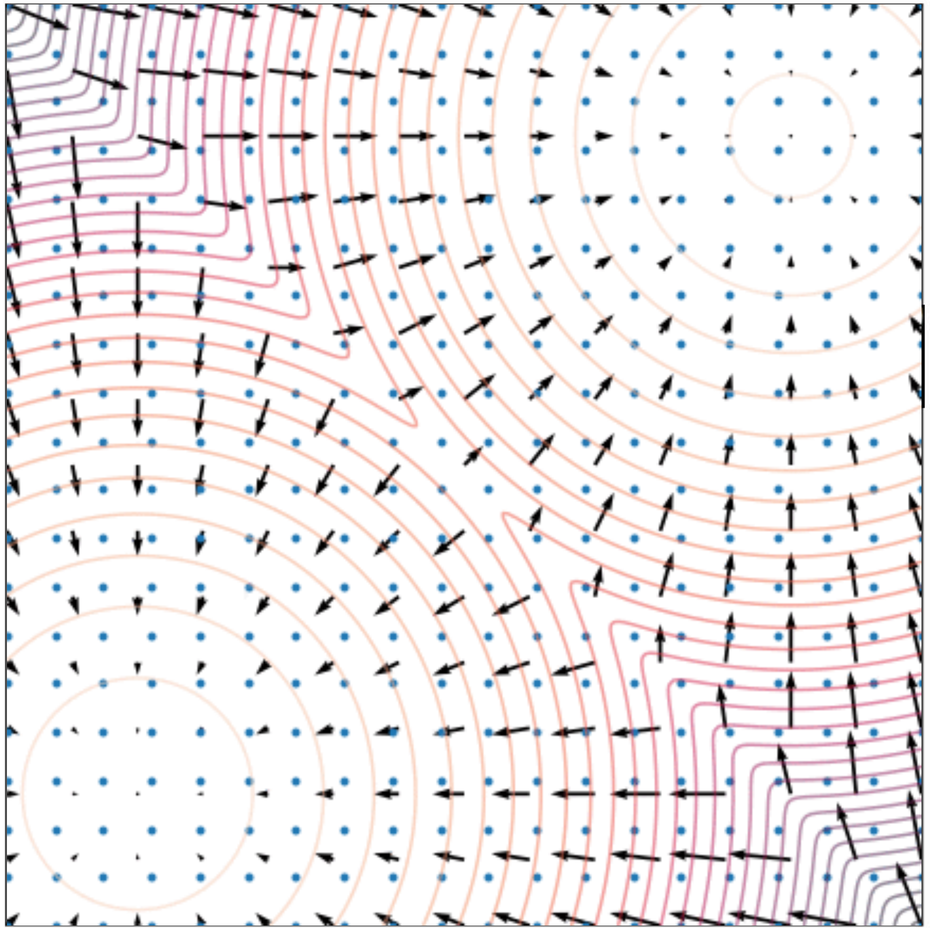

The main idea:
learn how to iteratively modify noise to move towards the data distribution

$$\nabla_x \log p(x)$$

Figure by Yang Song
https://yang-song.net/blog/2021/score/

How? Training by Denoising

$\epsilon \sim \mathcal{N}(0, \mathbb{I})$
$t \sim \mathcal{U}(0, T)$

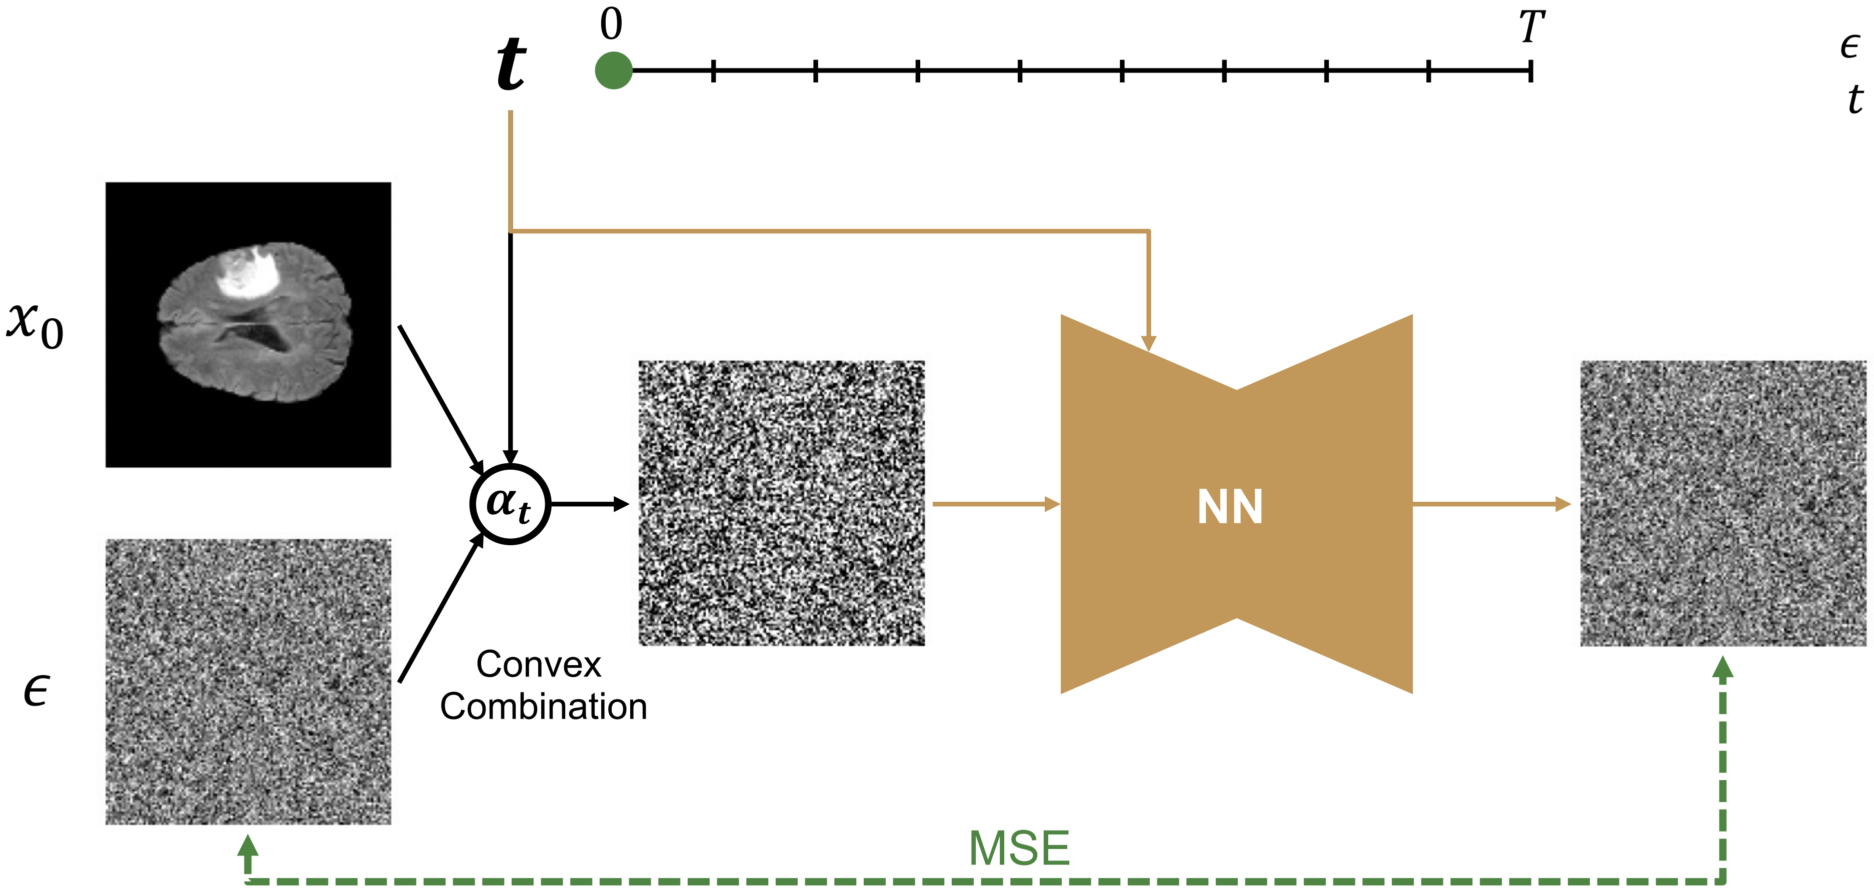

How? Inference

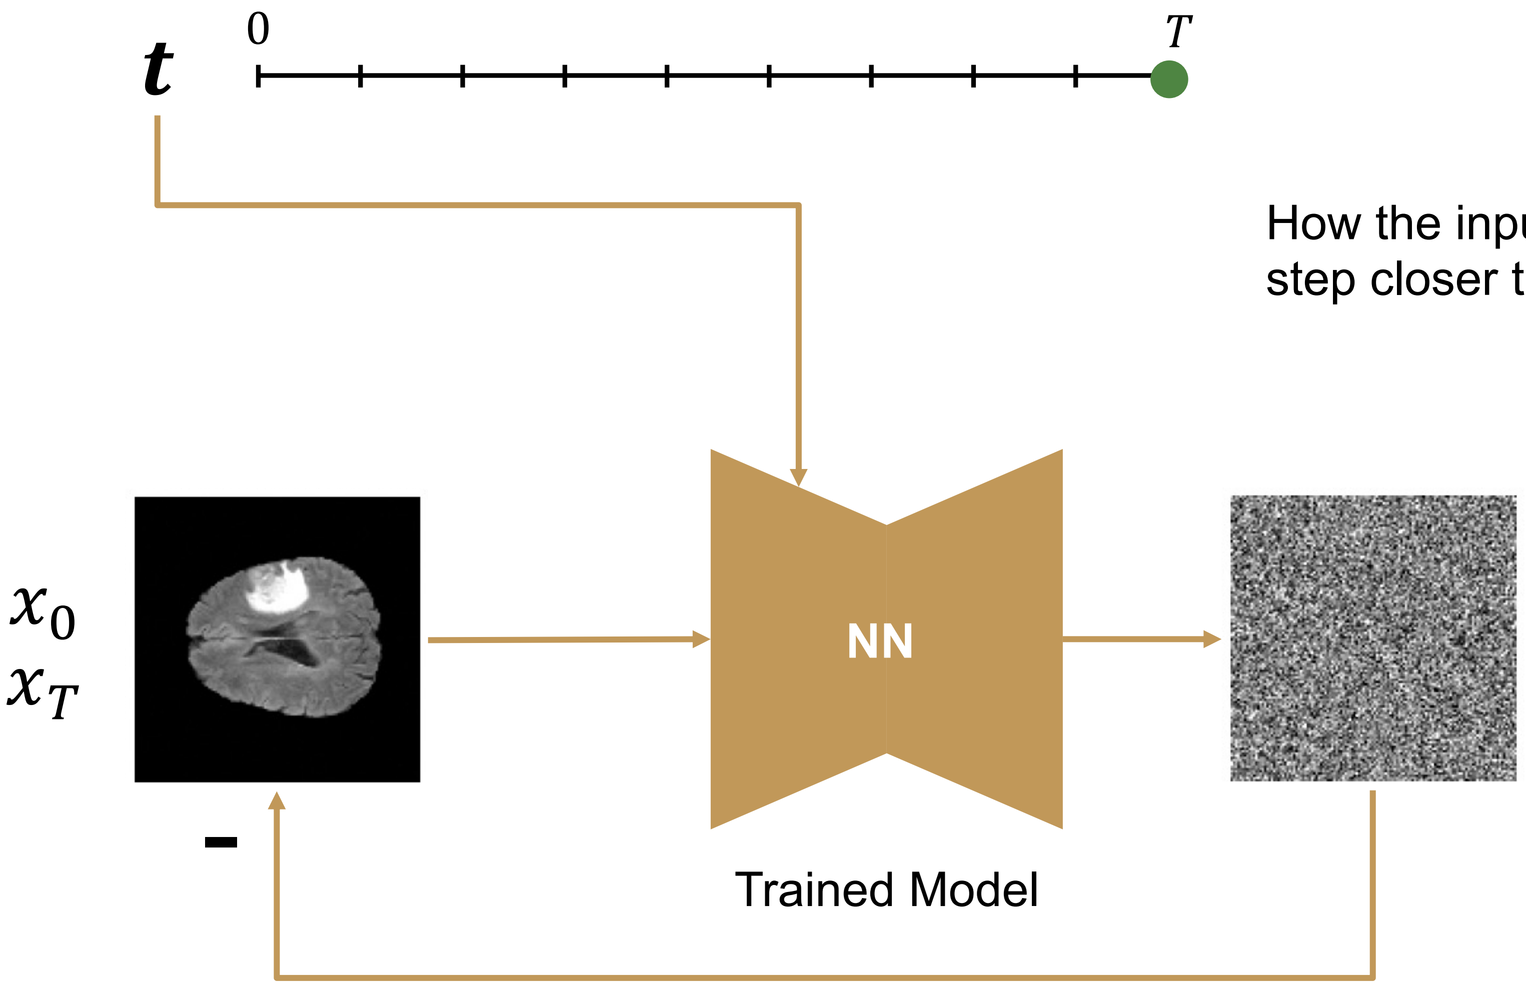

How the input can be a step closer to an image?

Architecture – Reusing the *classics*, and the *SoTA*, and *conditioning*

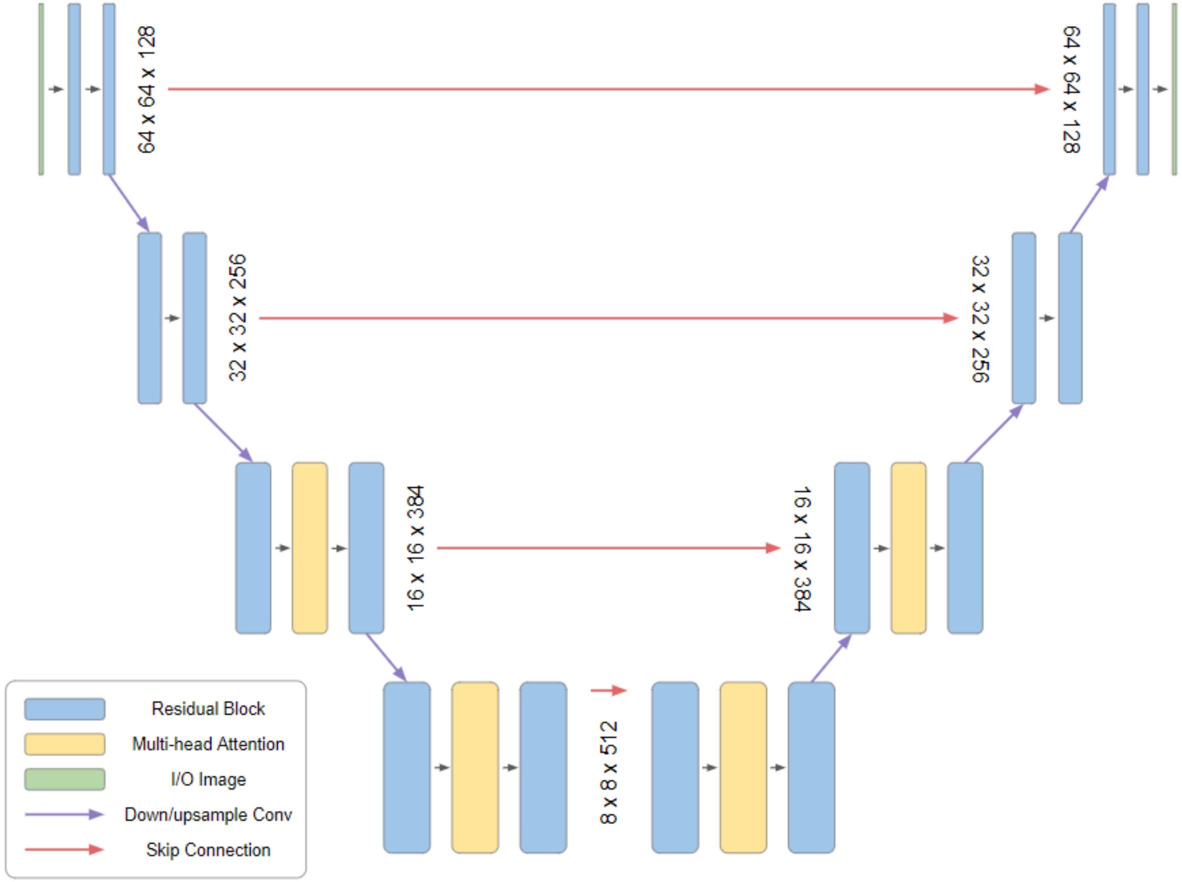

Jonathan Ho, Ajay Jain, Pieter Abbeel (2020) Denoising Diffusion Probabilistic Models. NeurIPS

Unet!

Or transformers
Or VQ-VAEs
Or…

Image reconstruction

Examples from the community

What is the task

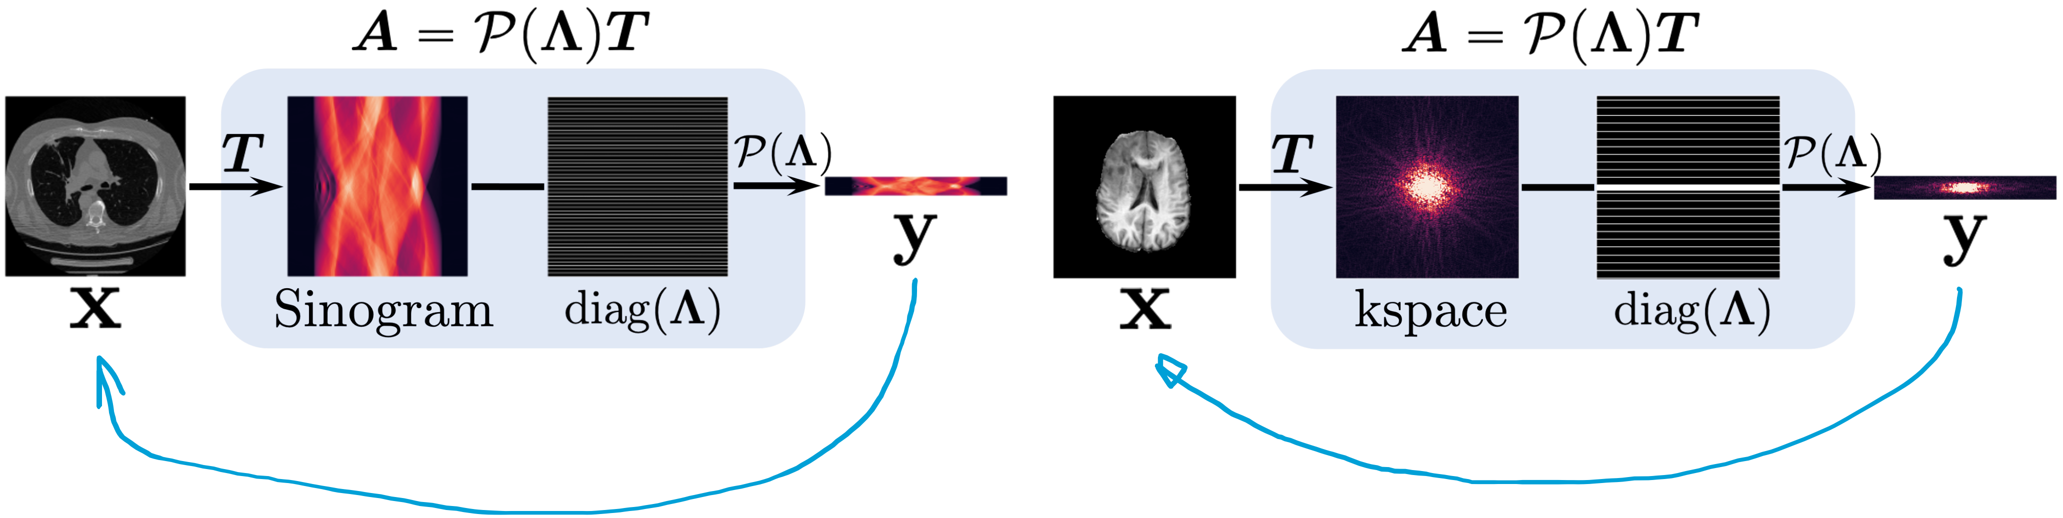

Song et al (2022) Solving Inverse Problems in Medical Imaging with Score-Based Generative Models. ICLR
Chung et al. (2022) Come-Closer-Diffuse-Faster: Accelerating Conditional Diffusion Models for Inverse Problems through Stochastic Contraction. CVPR
Luo et al. (2022) MRI Reconstruction via Data-Driven Markov Chains with Joint Uncertainty Estimation arxiv:2202.01479
Xie et al. (2022) Measurement-Conditioned Denoising Diffusion Probabilistic Model for Under-Sampled Medical Image Reconstruction. MICCAI
Peng et al. (2022) Towards Performant and Reliable Undersampled MR Reconstruction via Diffusion Model Sampling. MICCAI
Gungor et al. (2022) Adaptive Diffusion Priors for Accelerated MRI Reconstruction. arxiv:2207.05876
Cui et al. (2022) Self-Score: Self-Supervised Learning on Score-Based Models for MRI Reconstruction. Arxiv:2209.00835
Cao et al. (2022) High-Frequency Space Diffusion Models for Accelerated MRI. arxiv:2208.05481
Chung et al.(2022) Improving Diffusion Models for Inverse Problems using Manifold Constraints. arxiv:2206.00941
Chung et al. (2022) MR Image Denoising and Super-Resolution Using Regularized Reverse Diffusion. arxiv:2203.12621
Chung et al. (2021) Score-based diffusion models for accelerated MRI. MIA 2021
Hu et al. (2022) Unsupervised Denoising of Retinal OCT with Diffusion Probabilistic Model. arxiv:2201.11760
Gong et al (2022) PET image denoising based on denoising diffusion probabilistic models. arxiv:2209.06167

Figure by Song et al ICLR 2022.

MRI Reconstruction with Adaptive Diffusion Priors

<u>AdaDiff:</u> an adaptive diffusion model for accelerated MRI reconstruction

- An **unconditional** diffusion prior is trained on fully-sampled MR acquisitions
- The diffusion prior is **adapted** to the test subject by enforcing data consistency
- SoTA image fidelity and generalization performance under domain shifts

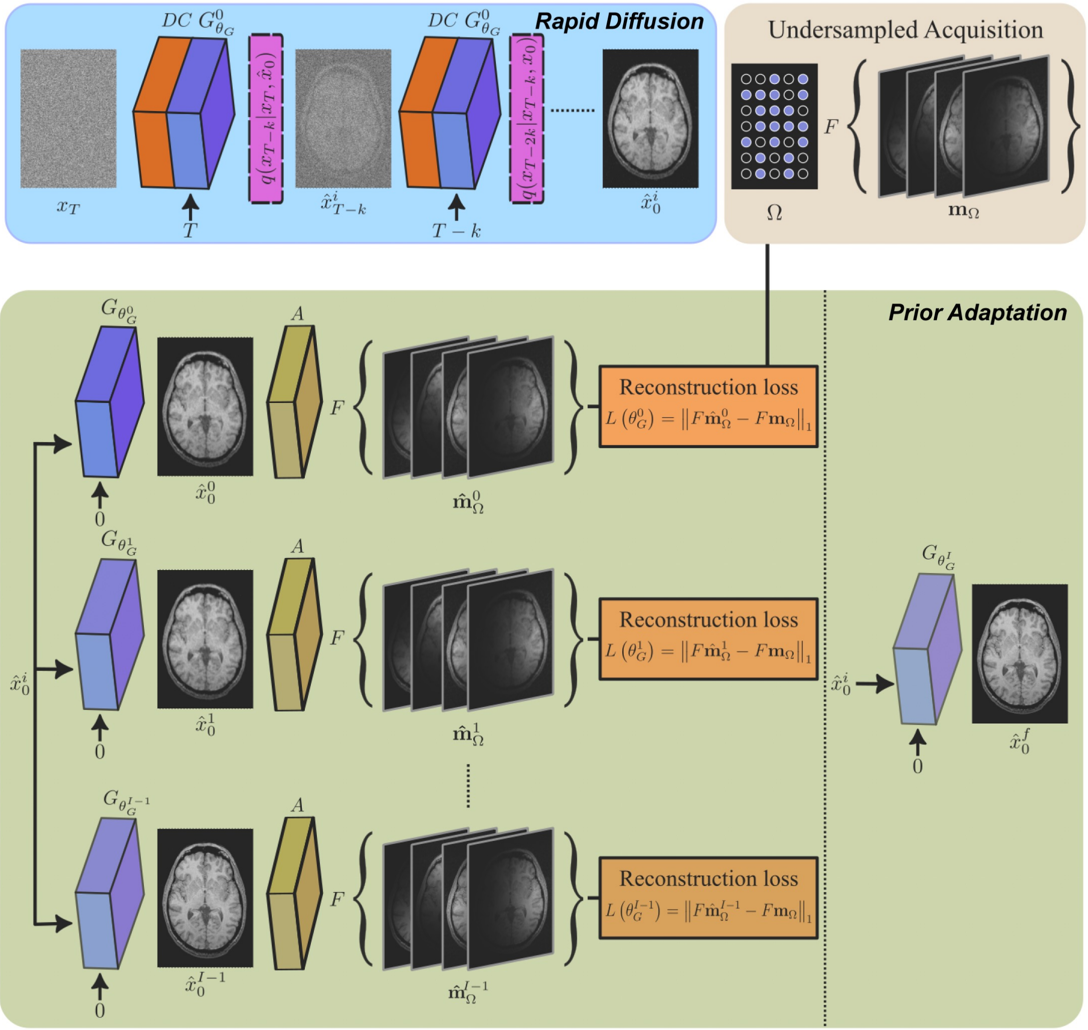

Slides courtesy of Tolga Cukur

1. Gungor et al (2022). Adaptive Diffusion Priors for Accelerated MRI Reconstruction. arXiv:2207.05876. (https://github.com/icon-lab/AdaDiff)

MRI Reconstruction with Adaptive Diffusion Priors

Models trained on the fastMRI dataset, tested on the IXI dataset at R=8x

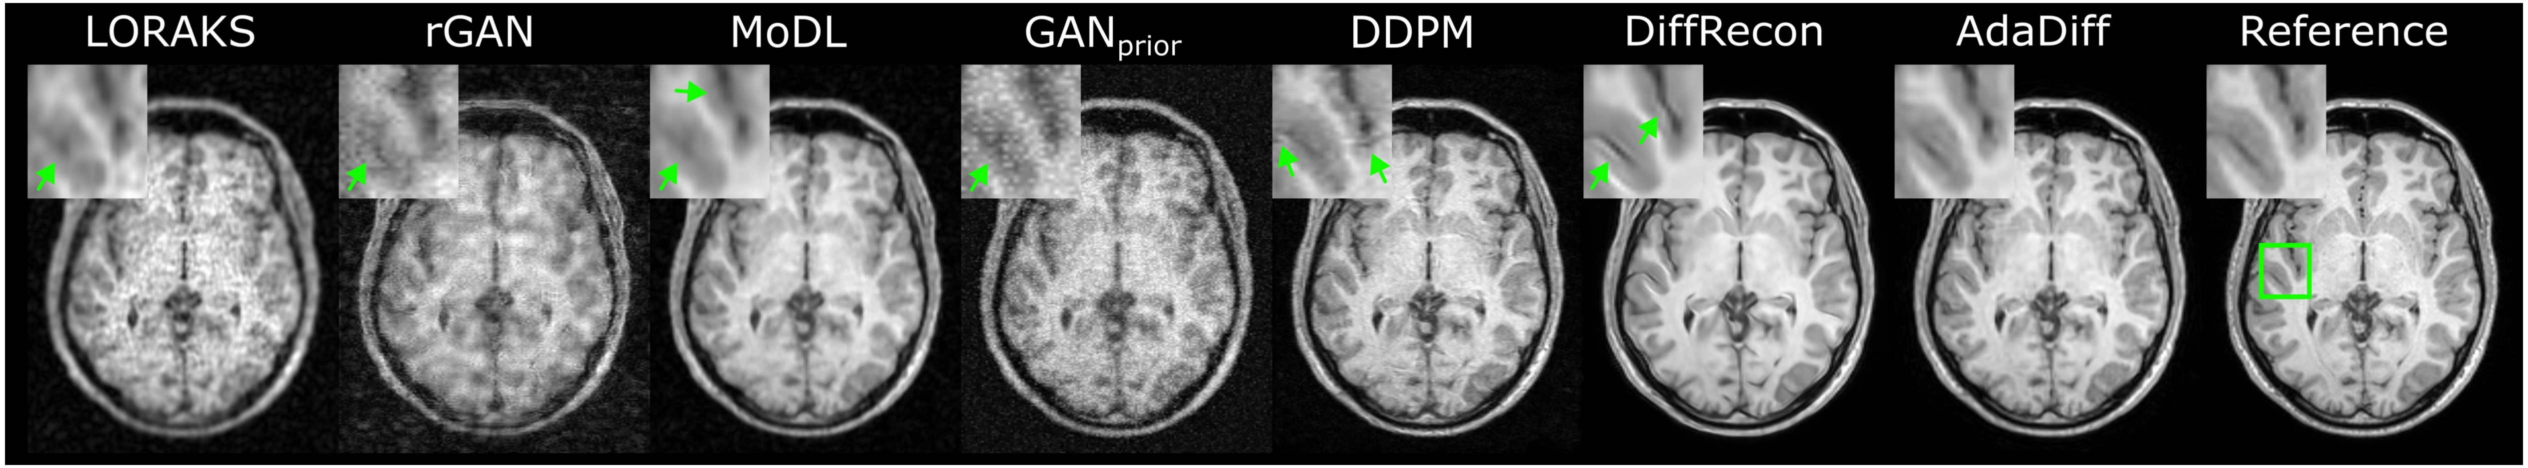

Slides courtesy of Tolga Cukur

1. Gungor et al (2022). Adaptive Diffusion Priors for Accelerated MRI Reconstruction. arXiv:2207.05876. (https://github.com/icon-lab/AdaDiff)

Image synthesis

Examples from the community

What is the task?

Real

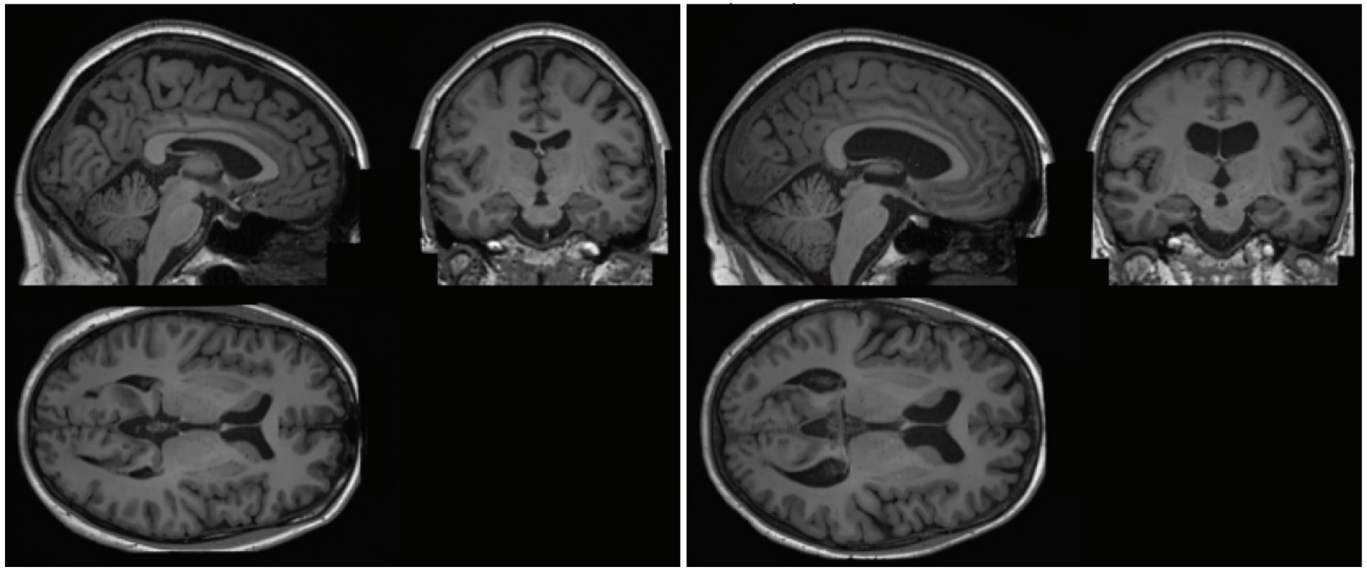

Synthetic

Pinaya et al (2022) Brain Imaging Generation with Latent Diffusion Models. MICCAI workshop
Kim et al. (2022) Diffusion Deformable Model for 4D Temporal Medical Image Generation. MICCAI
Khader et al. (2022) Medical Diffusion -- Denoising Diffusion Probabilistic Models for 3D Medical Image Generation. arXiv:2211.03364
Packhäuser et al. (2022) Generation of Anonymous Chest Radiographs Using Latent Diffusion Models for Training Thoracic Abnormality Classification Systems. arXiv:2211.01323
Ali et al. (2022) Spot the fake lungs: Generating Synthetic Medical Images using Neural Diffusion Models. arXiv:2211.00902
Rouzrokh et al. (2022) Multitask Brain Tumor Inpainting with Diffusion Models: A Methodological Report. arXiv:2210.12113
Chambon et al (2022) Adapting Pretrained Vision-Language Foundational Models to Medical Imaging Domains. arXiv:2210.04133
Lyu et al. (2022) Conversion Between CT and MRI Images Using Diffusion and Score-Matching Models. arXiv:2209.12104
Ozbey et al. (2022) Unsupervised Medical Image Translation with Adversarial Diffusion Models. arXiv:2207.08208
Meng et al. (2022) A Novel Unified Conditional Score-based Generative Framework for Multi-modal Medical Image Completion. arXiv:2207.03430

Generating high-resolution 3D brain data

- Latent Diffusion Models trained on data from UK Biobank (N = 31,740)
  - T1 MRI brain images with 1 mm^3 voxel size (160 × 224 × 160 voxels)
- Conditioned on covariates, such as:
  - Age
  - Gender
  - Ventricular and Brain volumes
- Released dataset with 100,000 synthetic brains!

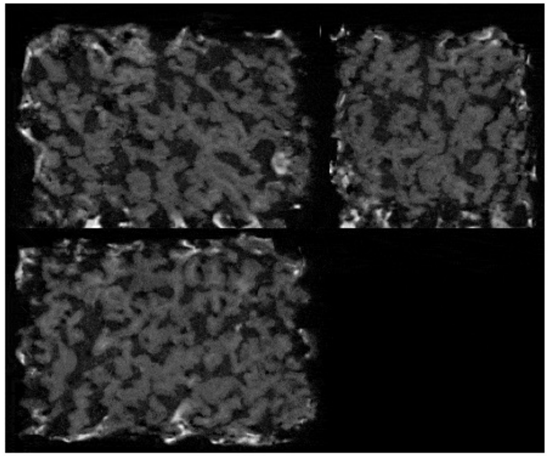

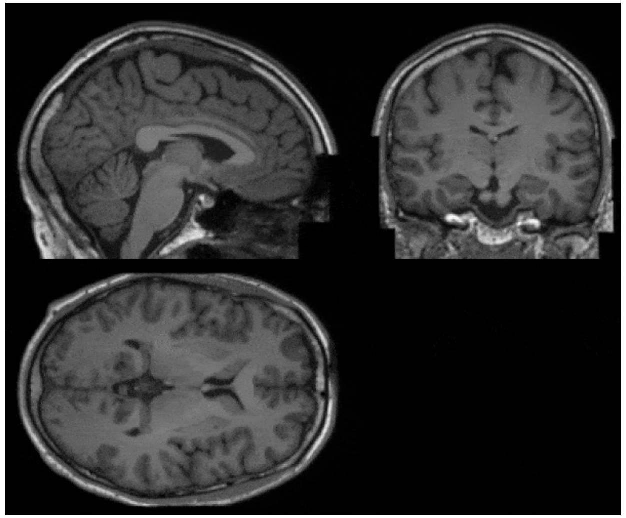

Slides courtesy of Walter H.L. Pinaya

1. Pinaya et al (2022). Brain Imaging Generation with Latent Diffusion Models. DOI:10.1007/978-3-031-18576-2_12

Generation of anonymous chest radiographs

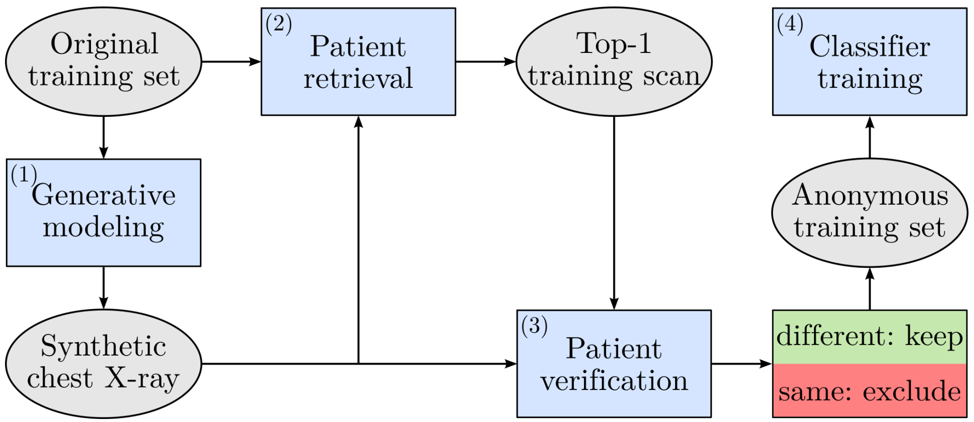

Fig. 1: Proposed privacy-enhancing image sampling strategy. Image taken from [1].

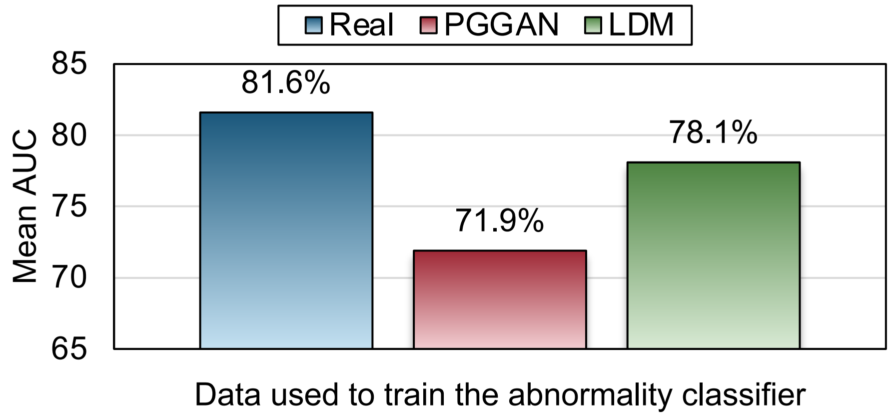

Fig. 2: Synthetic data can be used in lieu or real ones.

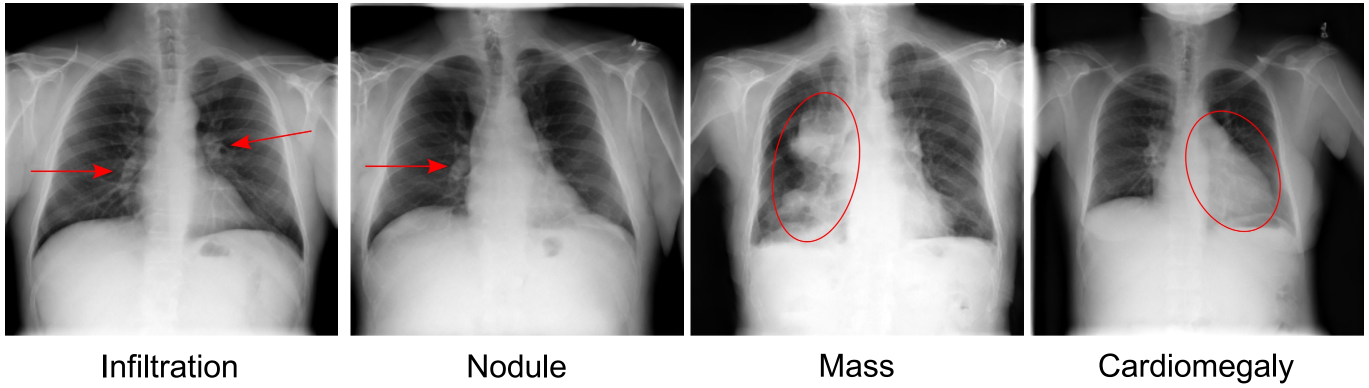

Fig. 3: Randomly selected images generated by the trained LDM. Images taken from [1].

Slides courtesy of Kai Packhäuser

1. Packhäuser et al (2022). Generation of Anonymous Chest Radiographs Using Latent Diffusion Models for Training Thoracic Abnormality Classification Systems. arXiv:2211.01323

Medical Image Translation with Adversarial Diffusion

SynDiff: an unsupervised diffusion model for medical image translation

- A non-diffusive module with cycle-consistency loss enables training on unpaired datasets
- An adversarial diffusive module maps fast source ➔ target
- Inference: 180ms with 4 steps (SynDiff) vs. 45000ms with 1000 steps (DDPM)

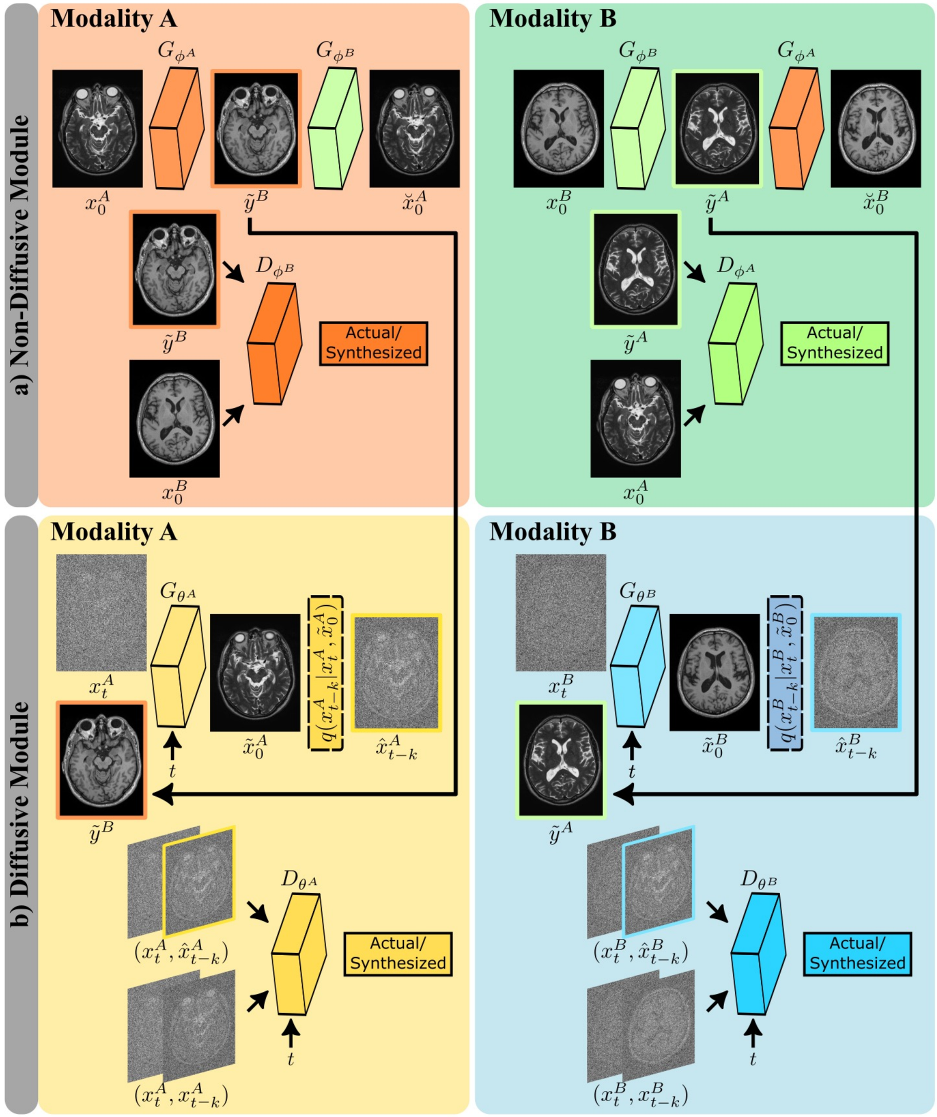

Slides courtesy of Tolga Cukur

1. Ozbey et al (2022). Unsupervised Medical Image Translation with Adversarial Diffusion Models. arXiv:2207.08208. (https://github.com/icon-lab/SynDiff)

Medical Image Translation with Adversarial Diffusion

MRI Contrast Translation

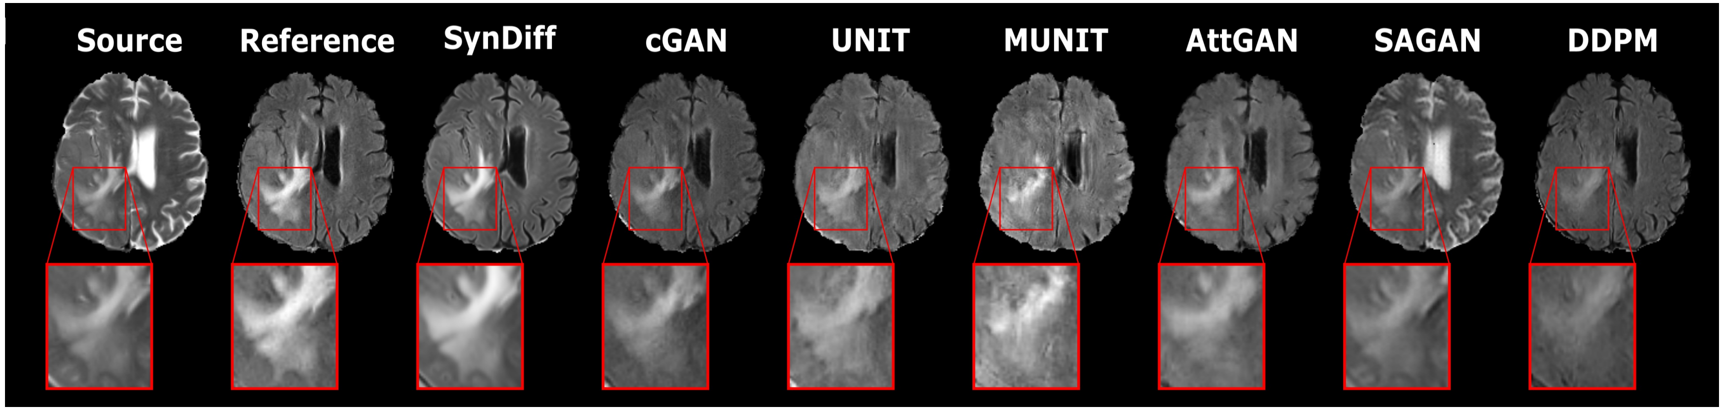

MRI to CT Translation

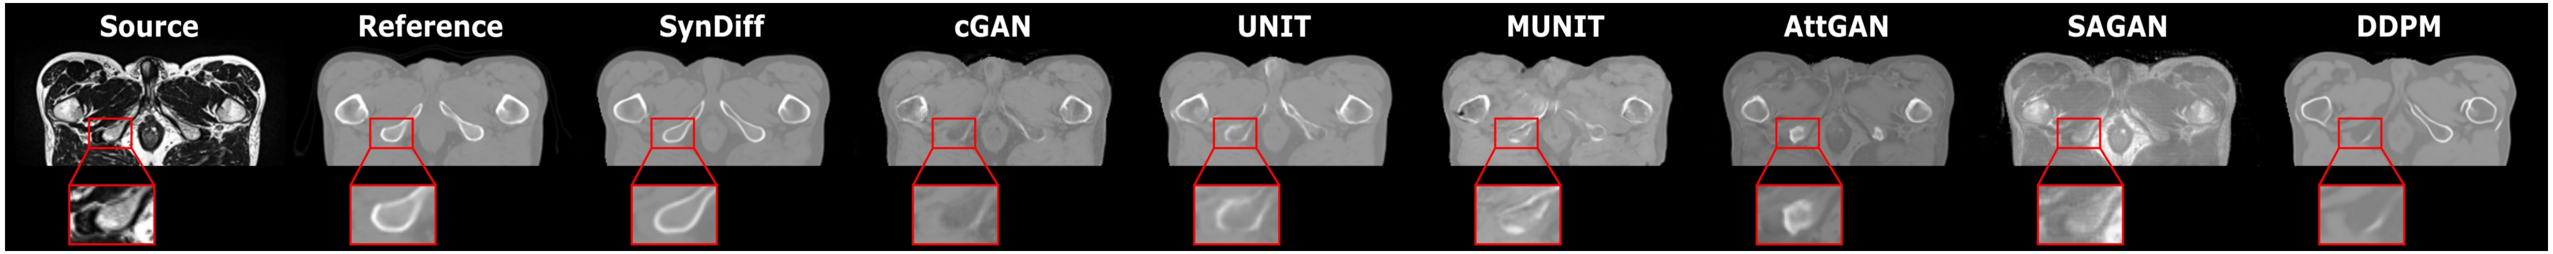

Slides courtesy of Tolga Cukur

1. Ozbey et al (2022). Unsupervised Medical Image Translation with Adversarial Diffusion Models. arXiv:2207.08208. (https://github.com/icon-lab/SynDiff)

Image segmentation

Examples from the community

What is the task?

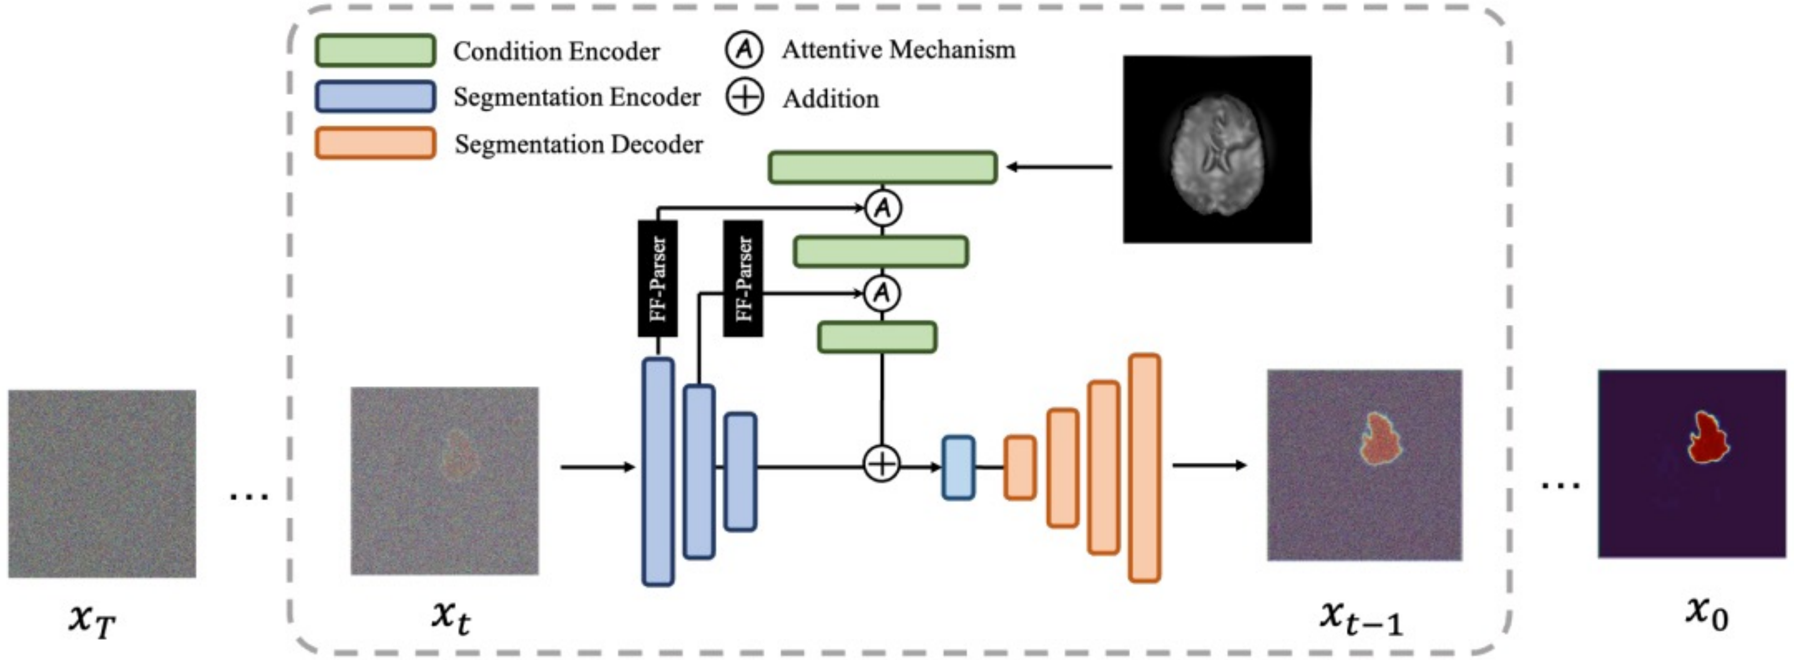

Guo et al (2022) Accelerating Diffusion Models via Pre-segmentation Diffusion Sampling for Medical Image Segmentation. arXiv:2210.17408
La Barbera et al. (2022) Anatomically constrained CT image translation for heterogeneous blood vessel segmentation. arXiv:2210.01713
Kim et al. (2022) Diffusion Adversarial Representation Learning for Self-supervised Vessel Segmentation. arXiv:2209.14566
Wu et al (2022) MedSegDiff: Medical Image Segmentation with Diffusion Probabilistic Model. arXiv:2211.00611

Figure by Wu et al arXiv 2022.

Diffusion Models for Segmentation

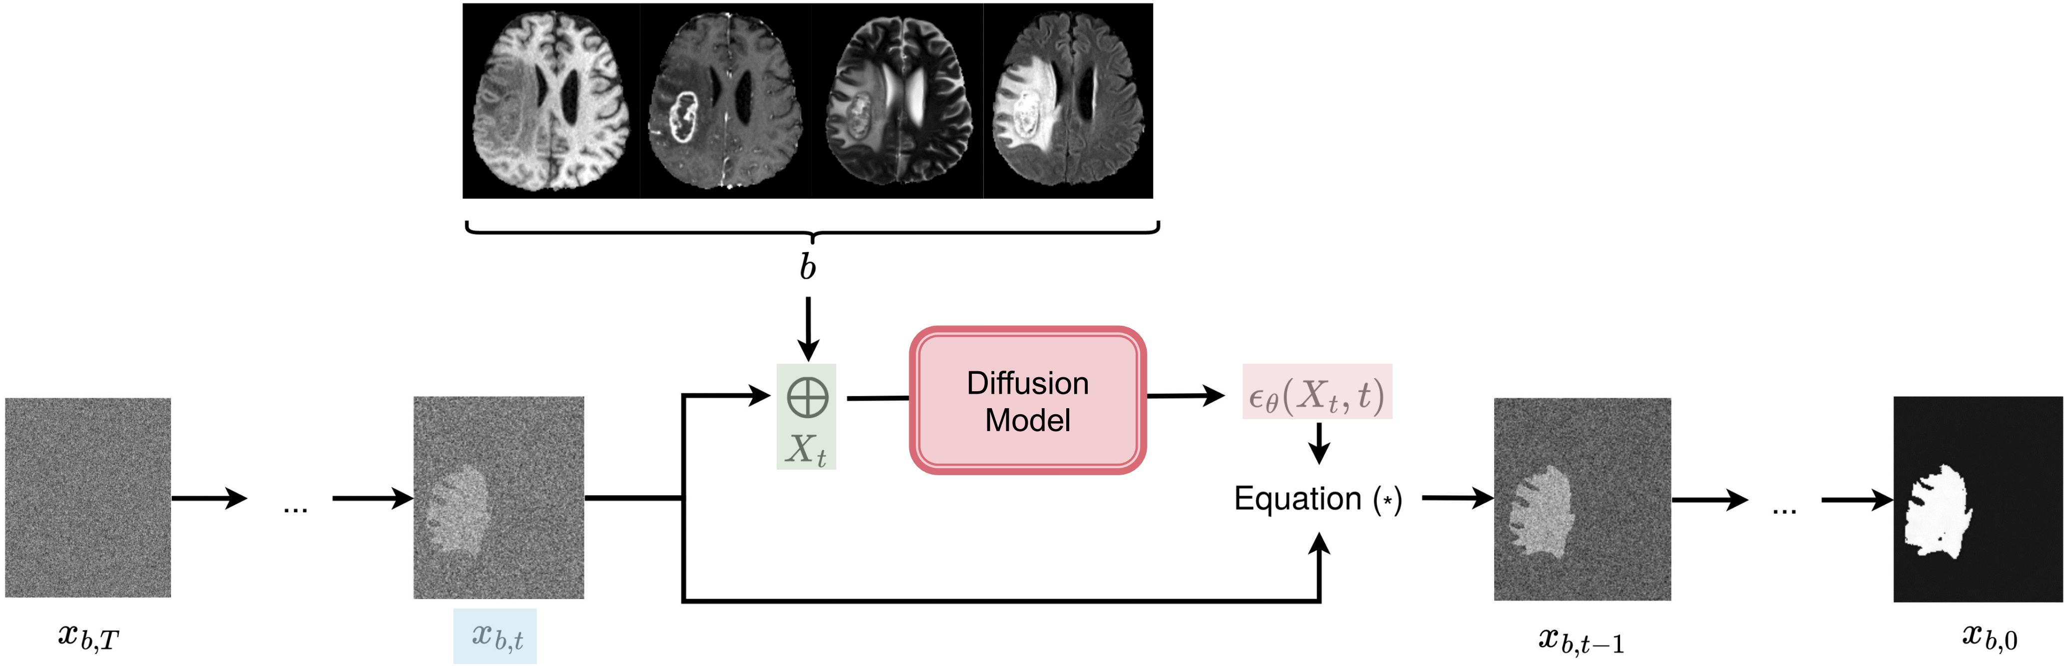

$$(*) \quad x_{b,t-1} = \frac{1}{\sqrt{\alpha_t}}\left(x_{b,t} - \frac{1-\alpha_t}{\sqrt{1-\overline{\alpha}_t}}\epsilon_\theta(X_t, t)\right) + \sigma_t \mathbf{z}, \quad \text{with } \mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$$

- Anatomical information is added by concatenating the input images *b* to the noisy segmentation mask $x_{b,t}$ *in* every step *t.*

Slides courtesy of Julia Wolleb

Wolleb et al (2022). Diffusion Models for Implicit Image Segmentation Ensembles, *MIDL 2022.* arXiv:2112.03145

Diffusion Models for Segmentation

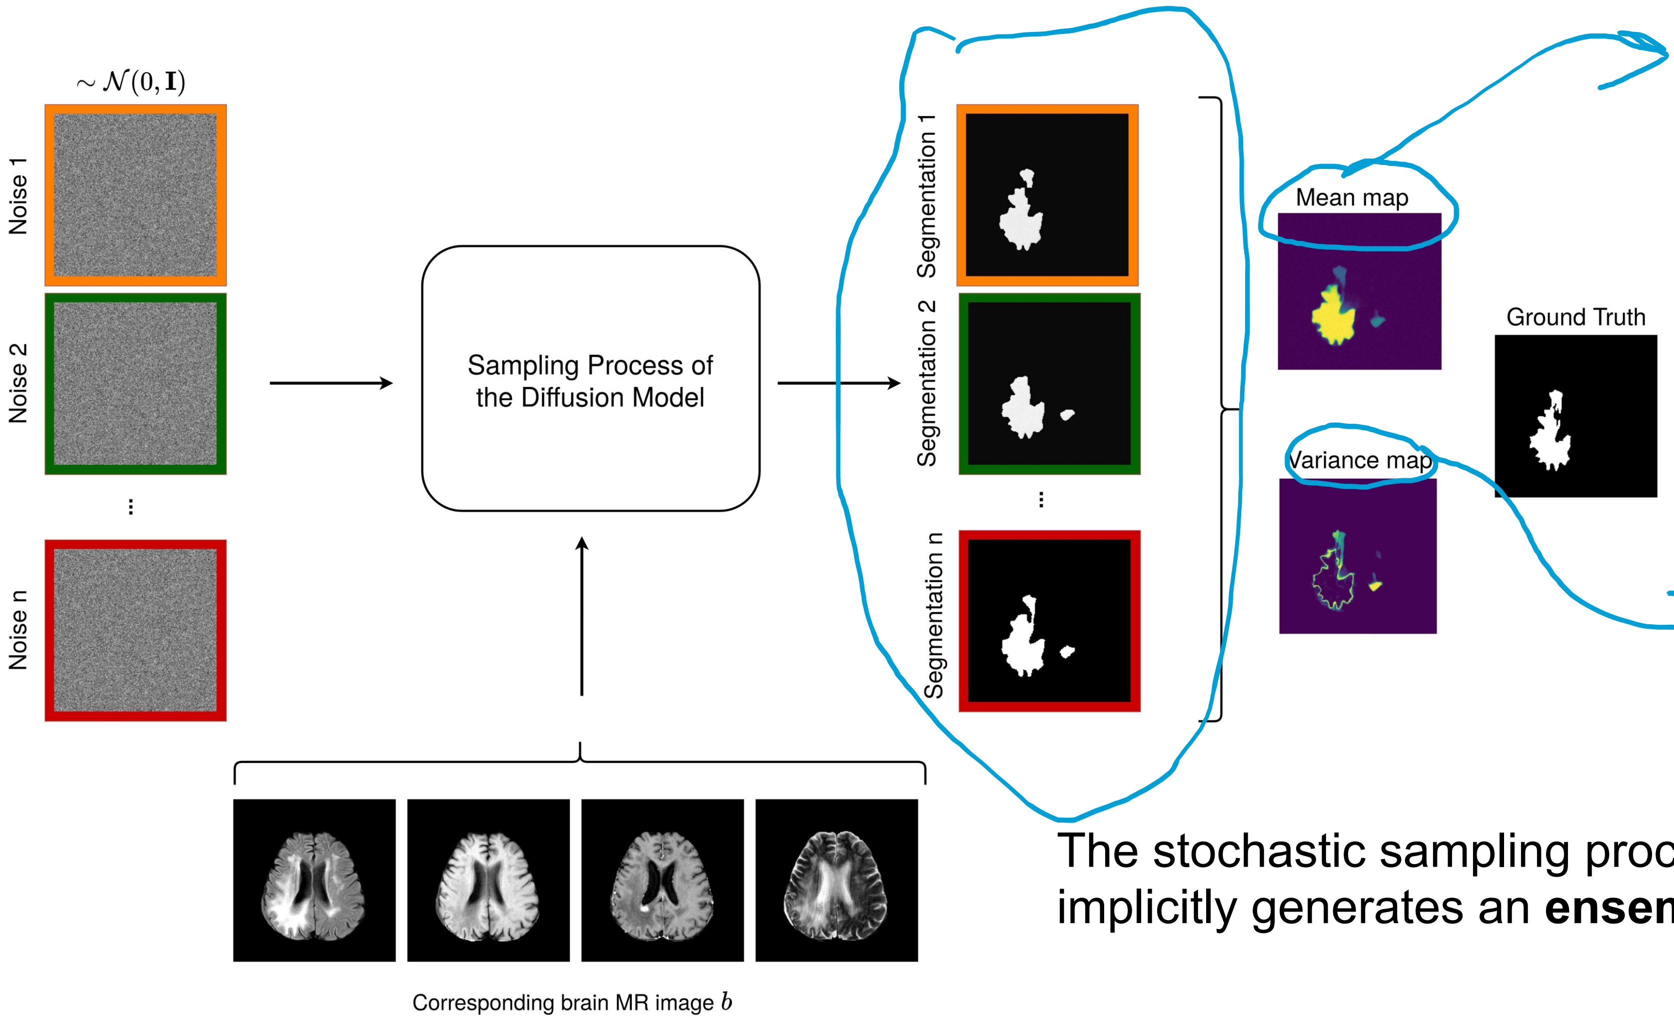

thresholded *mean* map improves **performance**

pixel-wise **uncertainty** map for model interpretability

The stochastic sampling process, implicitly generates an **ensemble** of masks

Slides courtesy of Julia Wolleb

Wolleb et al (2022). Diffusion Models for Implicit Image Segmentation Ensembles, *MIDL 2022.* arXiv:2112.03145

Diffusion adversarial representation learning

- Segment vessels of angiography images without labelled data
- **Diffusion module**: Provides latent features of background signals/vessel structures
- **Generation module**: Generates both synthetic images and vessel masks using our switchable SPADE layer

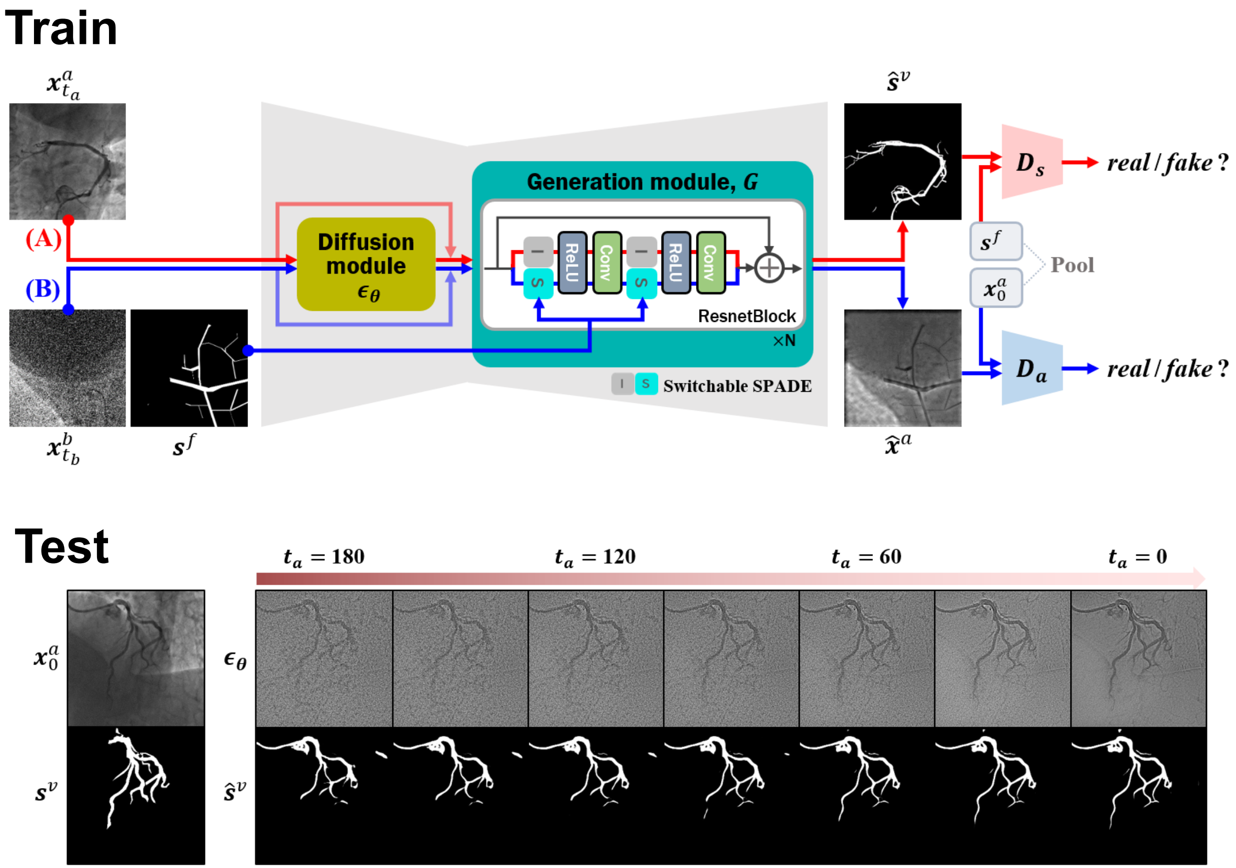

Slides courtesy of Boah Kim & Jong Chul Ye

Kim and Oh et al (2022) Diffusion Adversarial Representation Learning for Self-supervised Vessel Segmentation. In peer review

Diffusion adversarial representation learning

- Achieves state-of-the-art performance among the un-/self-supervised methods
- Robust to noisy images
- Generalisation capability on
  1) External data of X-ray angiography
  2) Cross-modal data of retinal imaging

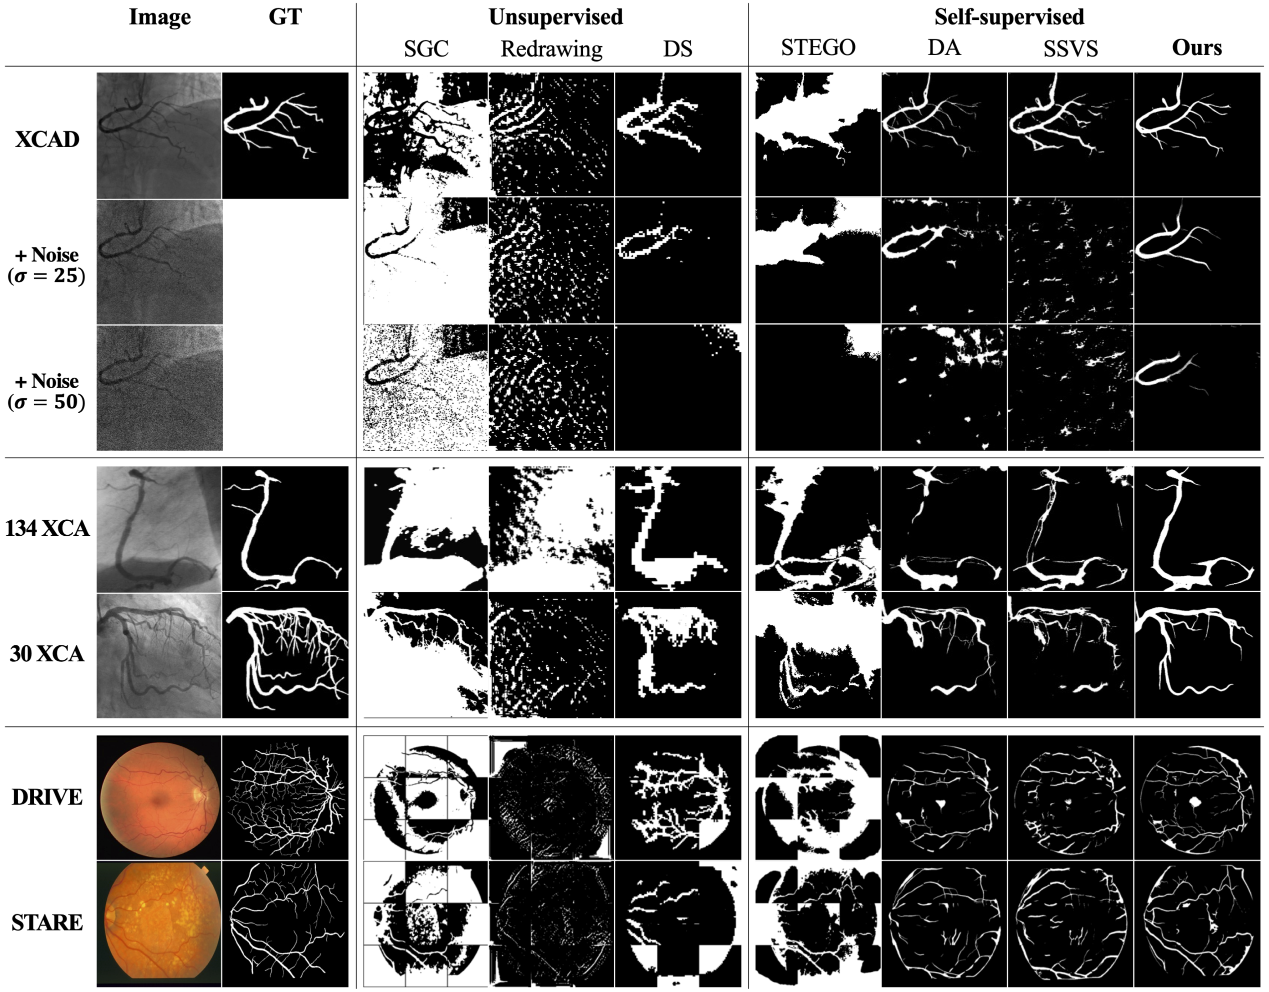

Slides courtesy of Boah Kim & Jong Chul Ye

Image registration

Examples from the community

DiffuseMorph

- Image registration along the continuous trajectory
- **Diffusion network**: estimates a conditional score function
- **Deformation network**: yields the registration fields & provide the deformed image

Loss function

$$\min_{G_\theta, M_\psi} L_{diffusion}(c, x_t, t) + \lambda L_{regist}(m, f)$$

$$L_{diffusion}(c, x_t, t) = \mathbb{E}_{\epsilon, x_t, t} \| G_\theta(c, x_t, t) - \epsilon \|_2^2$$

$$L_{regist}(m, f) = -(m(\phi) \otimes f) + \lambda_\phi \Sigma \| \nabla \phi \|^2$$

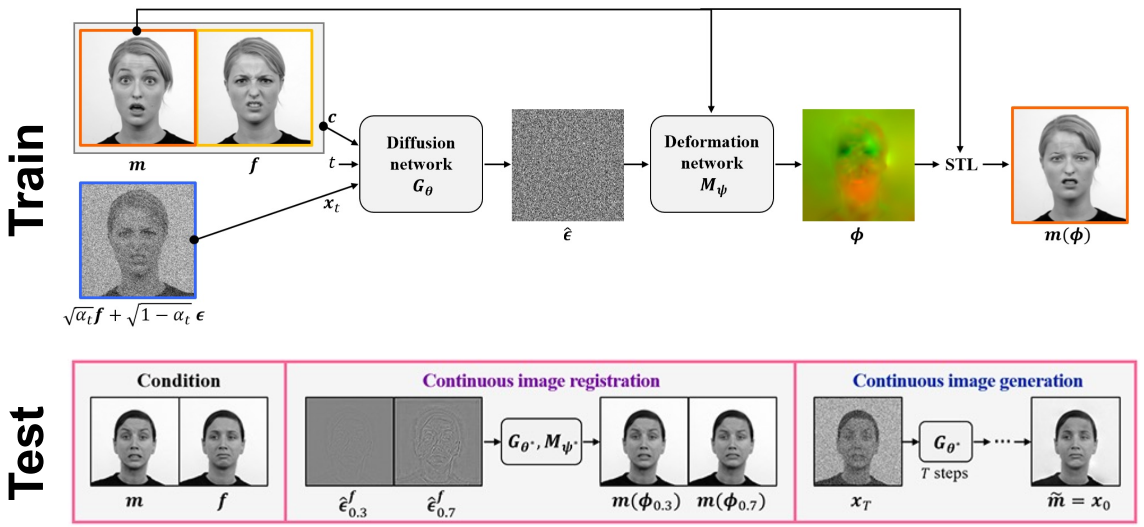

Algorithm 1 Continuous image registration

```
1: Input: Conditional images, c = (m, f)
2: Output: Deformed moving image, m(φ_η)
3: Set the latent feature ε̂^f = G_θ*(c, f, 0)
4: for η ∈ [0, 1] do
5:     ε̂^f_η ← η · ε̂^f
6:     φ_η ← M_ψ*(m, ε̂^f_η)
7: end for=
8: return m(φ_η)
```

Algorithm 2 Synthetic image generation process

```
1: Input: Conditional images, c = (m, f)
2: Output: Synthetic deformed image, x
3: Set T ∈ (0, T_train)
4: Sample x_T = √(ᾱ_T) m + √(1 − ᾱ_T) ε, where ε ∼ N(0, I)
5: for t = T, T − 1, ..., 1 do
6:     z ∼ N(0, I)
7:     x_{t−1} ← 1/√(1−β_t) (x_t − β_t/√(1−α_t) G_θ*(c, x_t, t)) + σ_t z
8: end for
9: return x_0
```

Slides courtesy of Boah Kim & Jong Chul Ye

1. Kim et al (2022). DiffuseMorph: Unsupervised Deformable Image Registration Along Continuous Trajectory Using Diffusion Models. ECCV 2022

DiffuseMorph

- Intra-subject 3D cardiac MR image registration

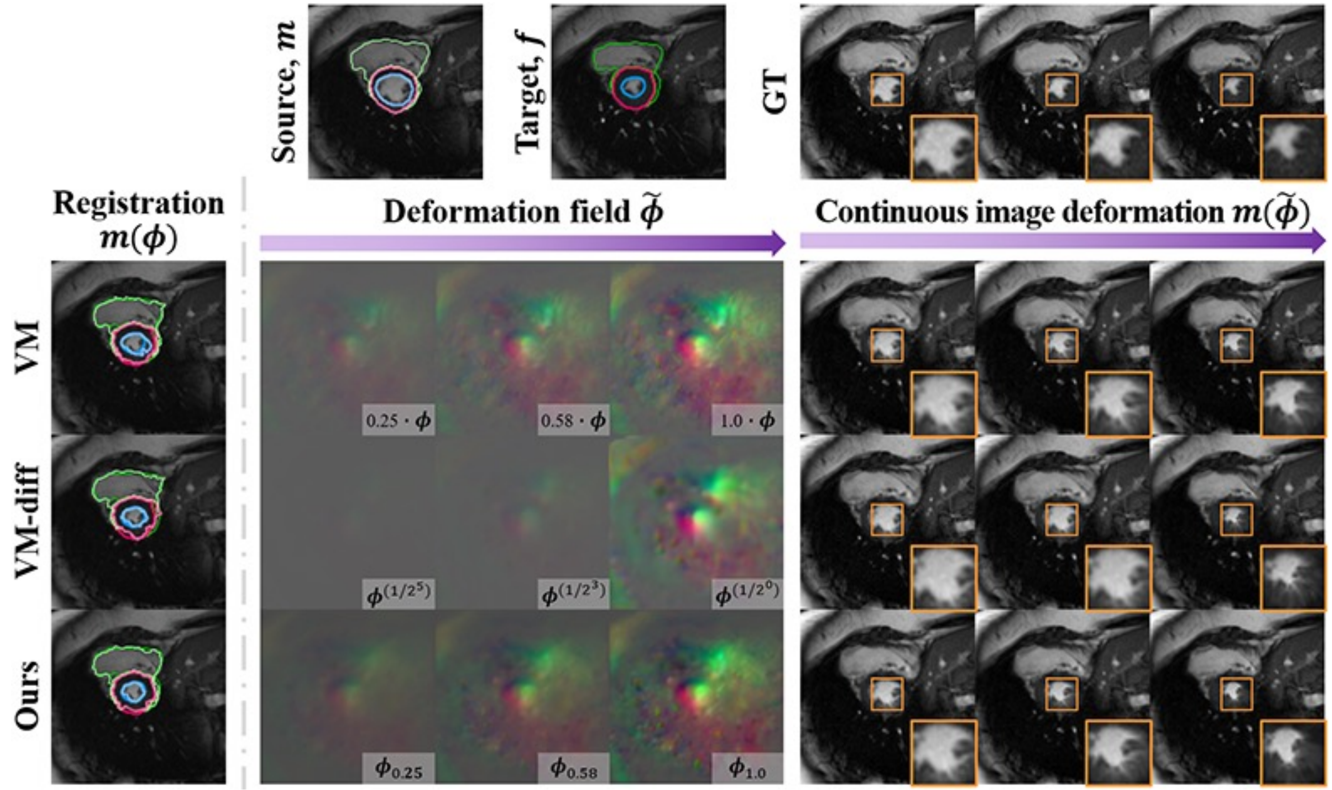

Methods	Dice	$\lvert J_\phi \rvert \leq 0$ (%)
Initial	0.642 (0.188)	-
VM [1]	0.787 (0.113)	0.169 (0.109)
VM-diff [2]	0.794 (0.104)	0.291 (0.188)
Ours	**0.802 (0.109)**	**0.161 (0.082)**

Slides courtesy of Boah Kim & Jong Chul Ye

1. Kim et al (2022). DiffuseMorph: Unsupervised Deformable Image Registration Along Continuous Trajectory Using Diffusion Models. ECCV 2022

Anomaly detection

Examples from the community

What is the task?

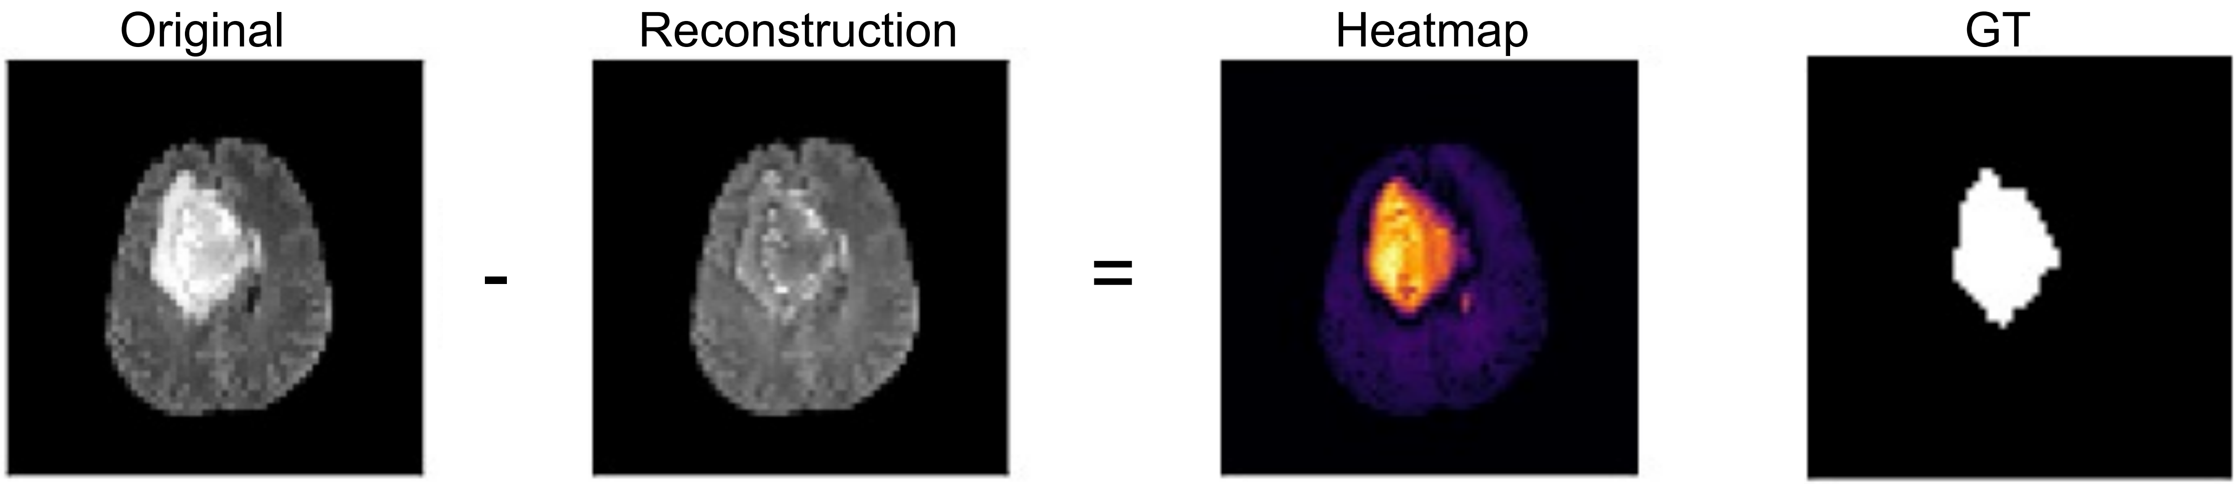

Sanchez et al. (2022) What is Healthy? Generative Counterfactual Diffusion for Lesion Localization. MICCAI workshop
Pinaya et al (2022) Fast Unsupervised Brain Anomaly Detection and Segmentation with Diffusion Models. MICCAI
Wolleb et al (2022) Diffusion Models for Medical Anomaly Detection. MICCAI
Wyatt et al (2022) AnoDDPM: Anomaly Detection with Denoising Diffusion Probabilistic Models using Simplex Noise. CVPR workshop

Figure by Sanchez et al

Unsupervised Anomaly Segmentation

- Latent Diffusion Model (LDM) learns the distribution of healthy brain data
- Compression (Vector-Quantised VAE) scales for high-resolution images

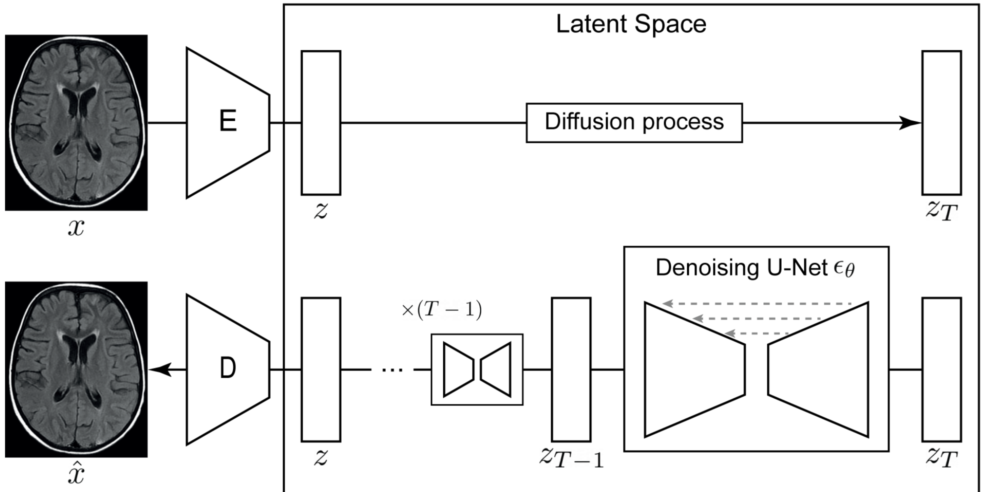

LDM identify regions with a **low** likelihood of belonging to a healthy dataset

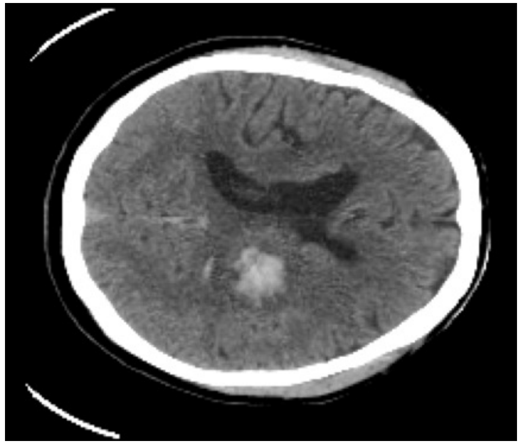

Reverse/denoising process is used to **inpaint** these regions and “**heal**” the possible anomalies

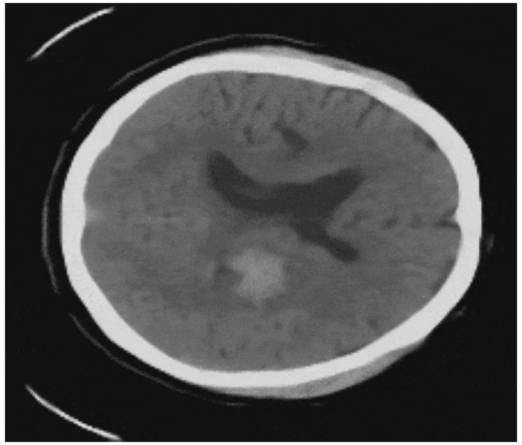

Slides courtesy of Walter H.L. Pinaya

1. Pinaya et al (2022). Fast Unsupervised Brain Anomaly Detection and Segmentation with Diffusion Models. DOI:10.1007/978-3-031-16452-1_67

Weakly-supervised Lesion Segmentation

- *Image-wise* labels to *pixel-wise* prediction
- **Manipulate** images between two classes

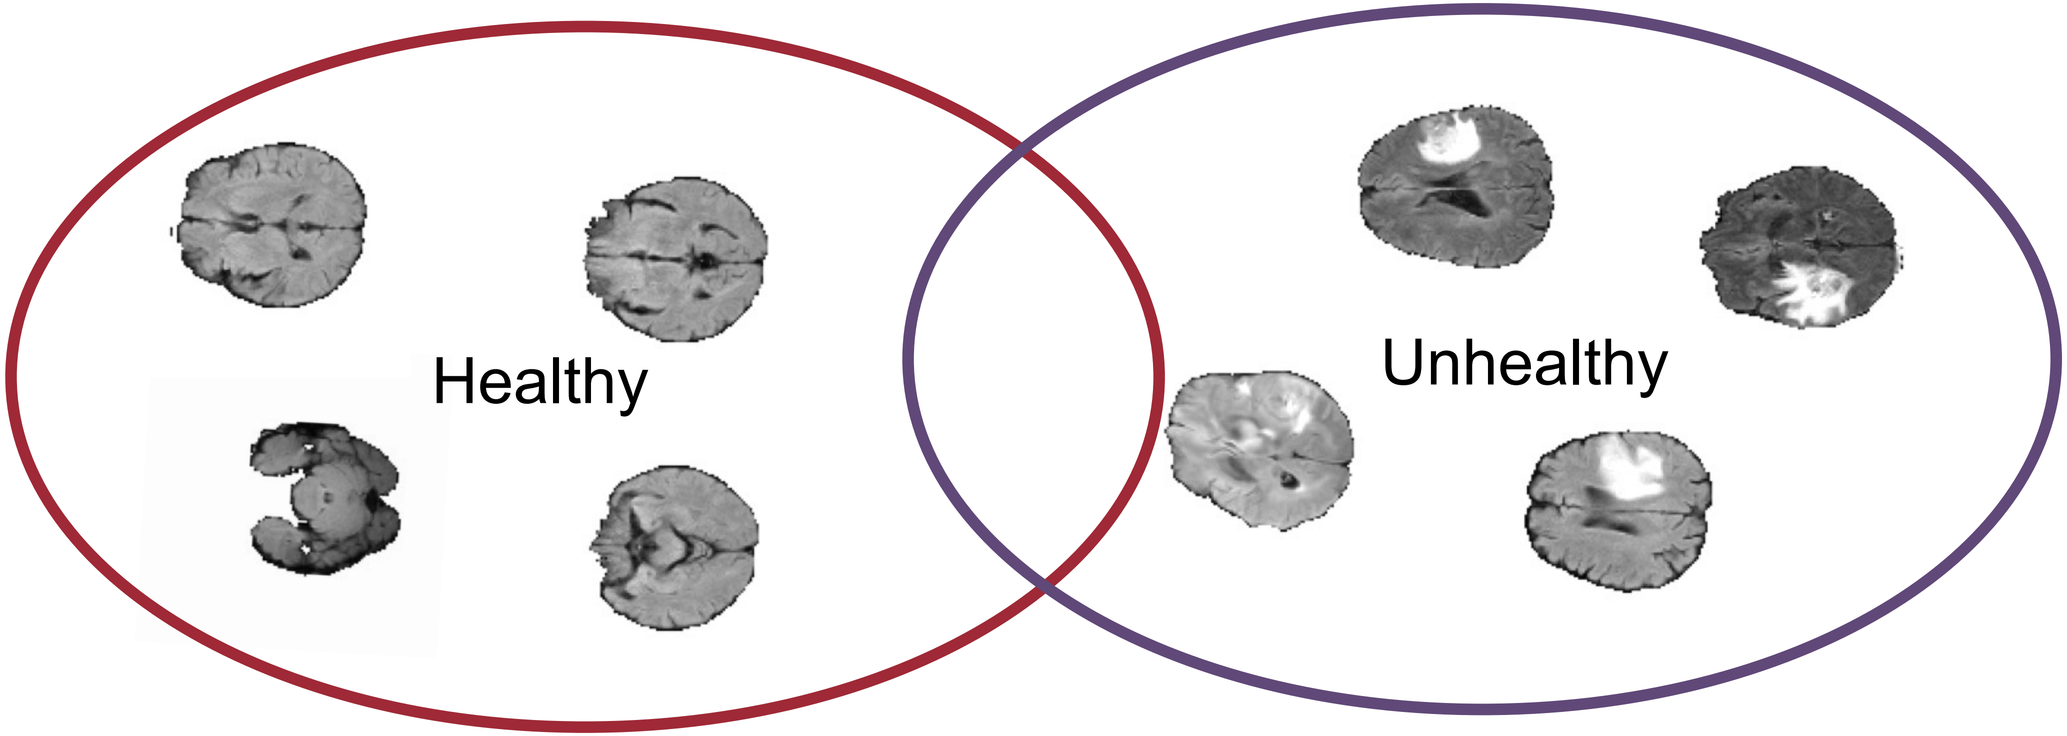

Diffusion Models for Medical Anomaly Detection

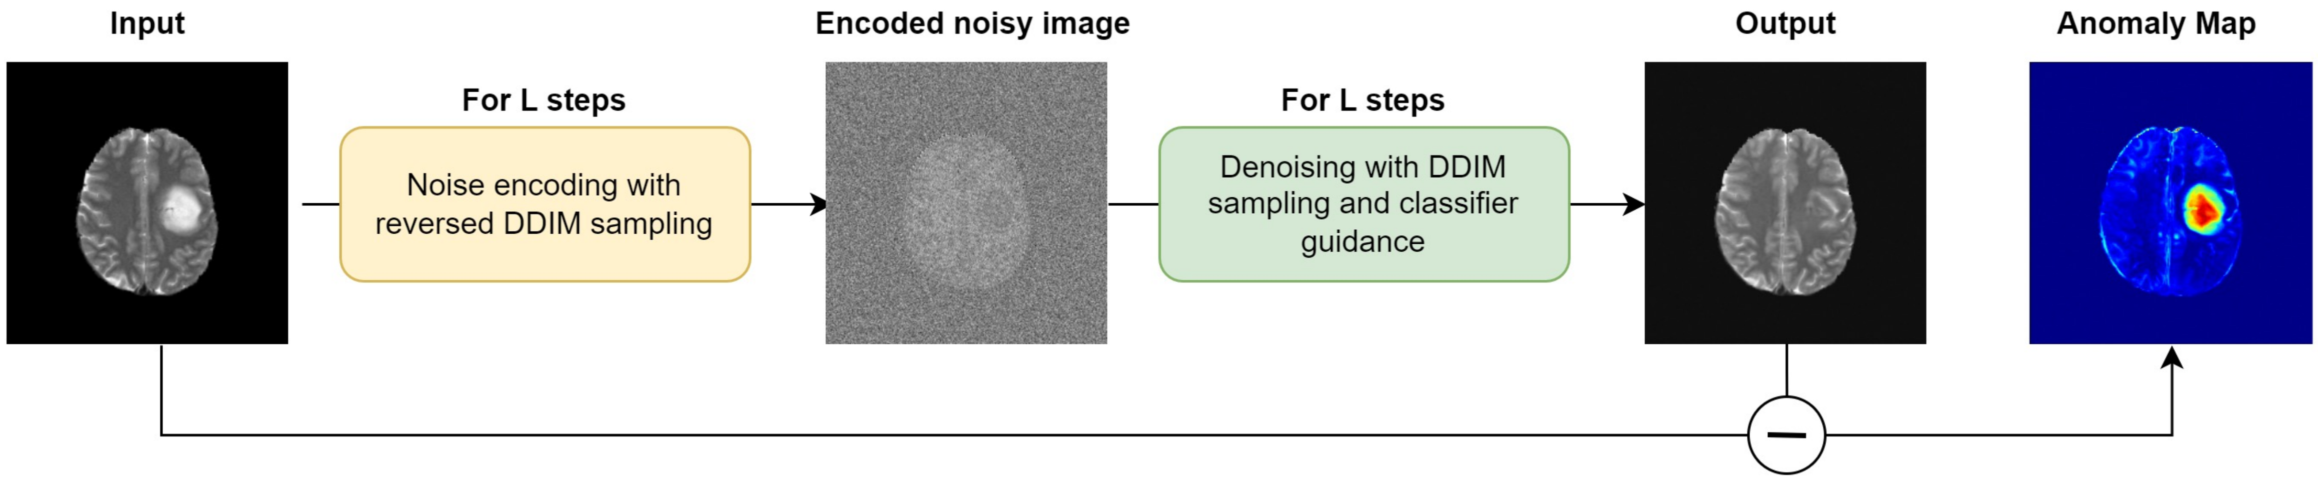

- Uses the iterative deterministic encoding and decoding scheme of **denoising diffusion implicit models**
- Translation to a healthy subject is based on **gradient guidance** of an external binary classifier

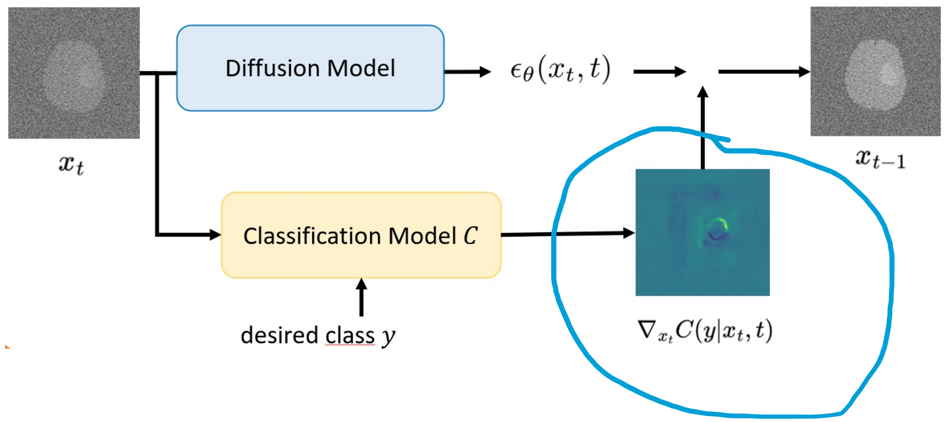

Slides courtesy of Julia Wolleb

1. Wolleb et al (2022). Diffusion Models for Medical Anomaly Detection, *MICCAI 2022.* arXiv:2203.04306

Lesion Localization with Diffusion

Classifier-free guidance

1. Encoding - Empty condition
2. Decoding - Target class

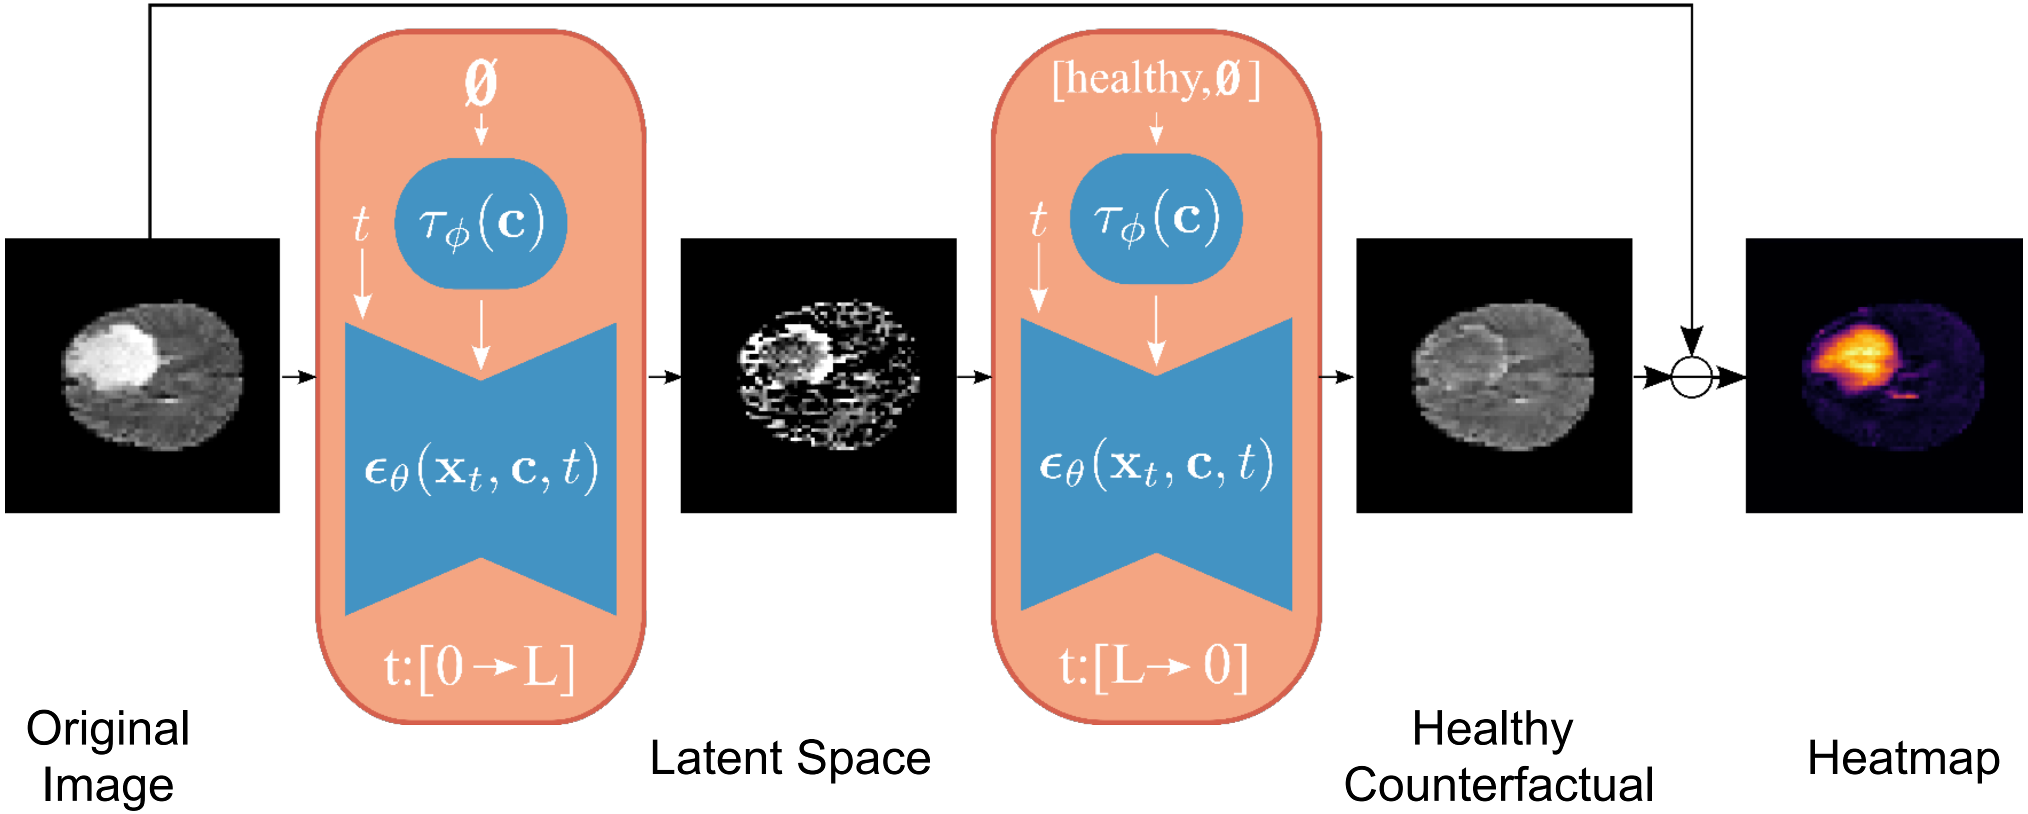

1. Sanchez et al (2022). What is Healthy? Generative Counterfactual Diffusion for Lesion Localization. DGM4*MICCAI 2022.* arXiv:2207.12268

Causality

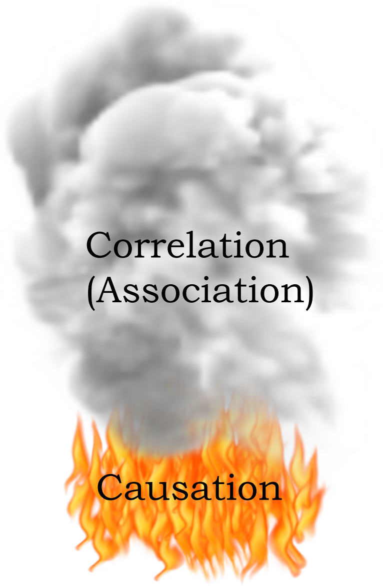

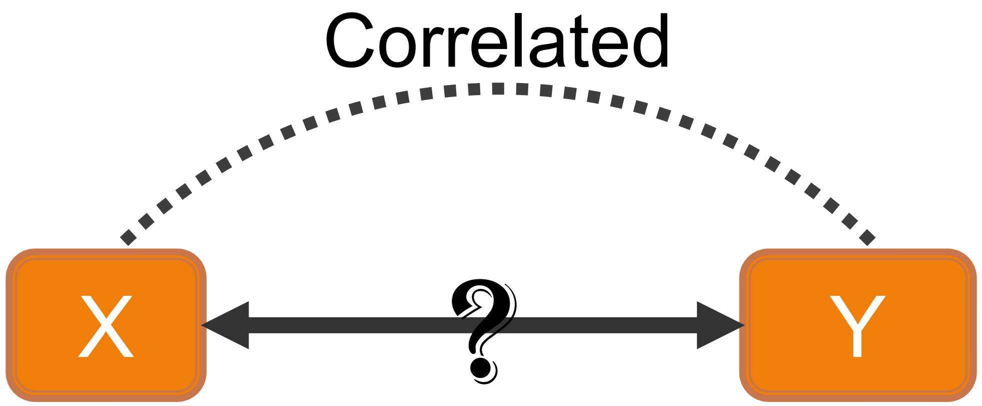

Diffusion for causal counterfactuals

"how an image should change to be classified as another class?"

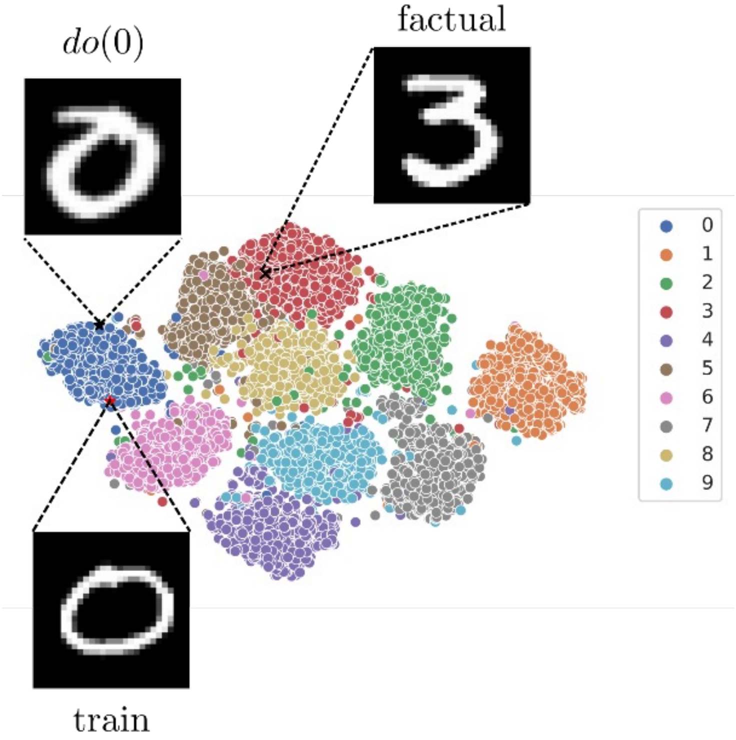

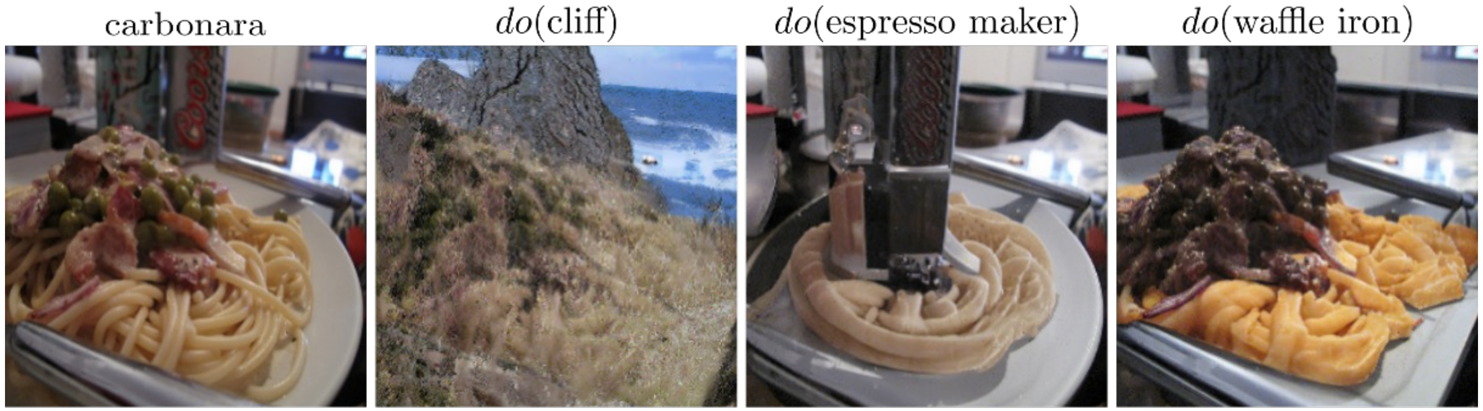

Sanchez et al (2022). Diffusion Causal Models for Counterfactual Estimation. CLeaR *2022.* arXiv:2202.10166

Diffusion for causal counterfactuals

We use diffusion ODE for estimating a **“latent space”** at t = T

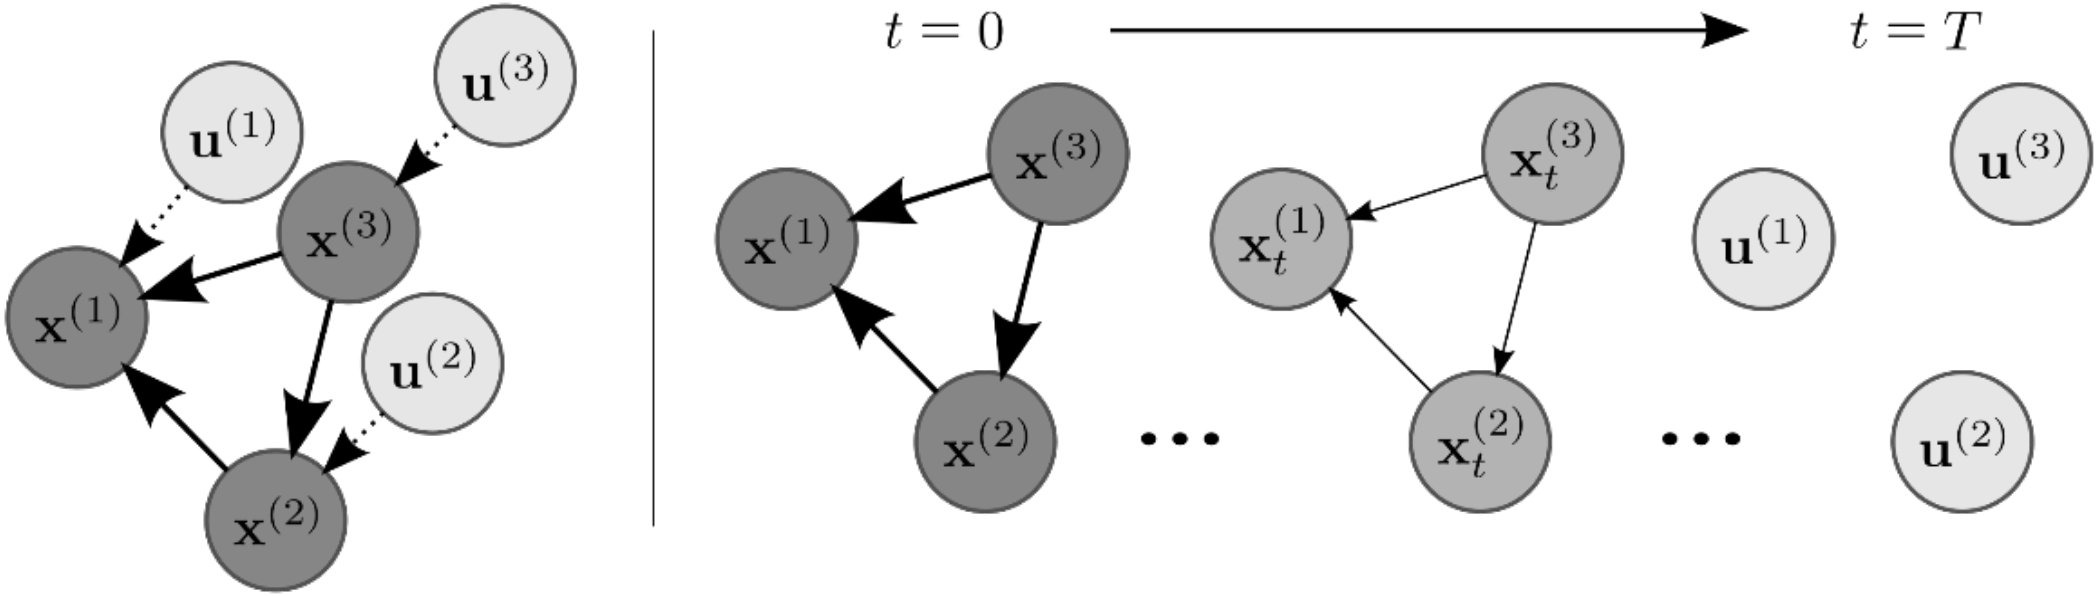

Sanchez et al (2022). Diffusion Causal Models for Counterfactual Estimation. CLeaR *2022.* arXiv:2202.10166

Diffusion for causal counterfactuals

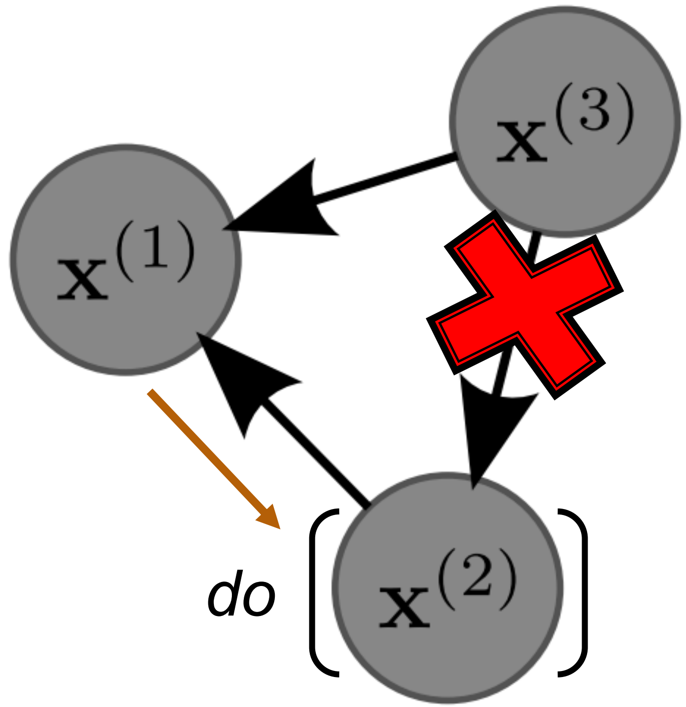

Given a **known** causal structure, learn to **estimate** causal effect of an **intervention**

$$\nabla_{x^{(1)}} \log p(\mathbf{x}^{(1)} \mid \mathbf{x}^{(2)})$$
$$\propto$$
$$\nabla_{x^{(1)}} \boxed{\log p(\mathbf{x}^{(2)} \mid \mathbf{x}^{(1)})} + \nabla_{x^{(1)}} \log p(\mathbf{x}^{(1)})$$

Sanchez et al (2022). Diffusion Causal Models for Counterfactual Estimation. CLeaR *2022.* arXiv:2202.10166

Diffusion for causal discovery

- Discovering causal (graph) structure from data
- Diffusion models help **identifying causal relationships**

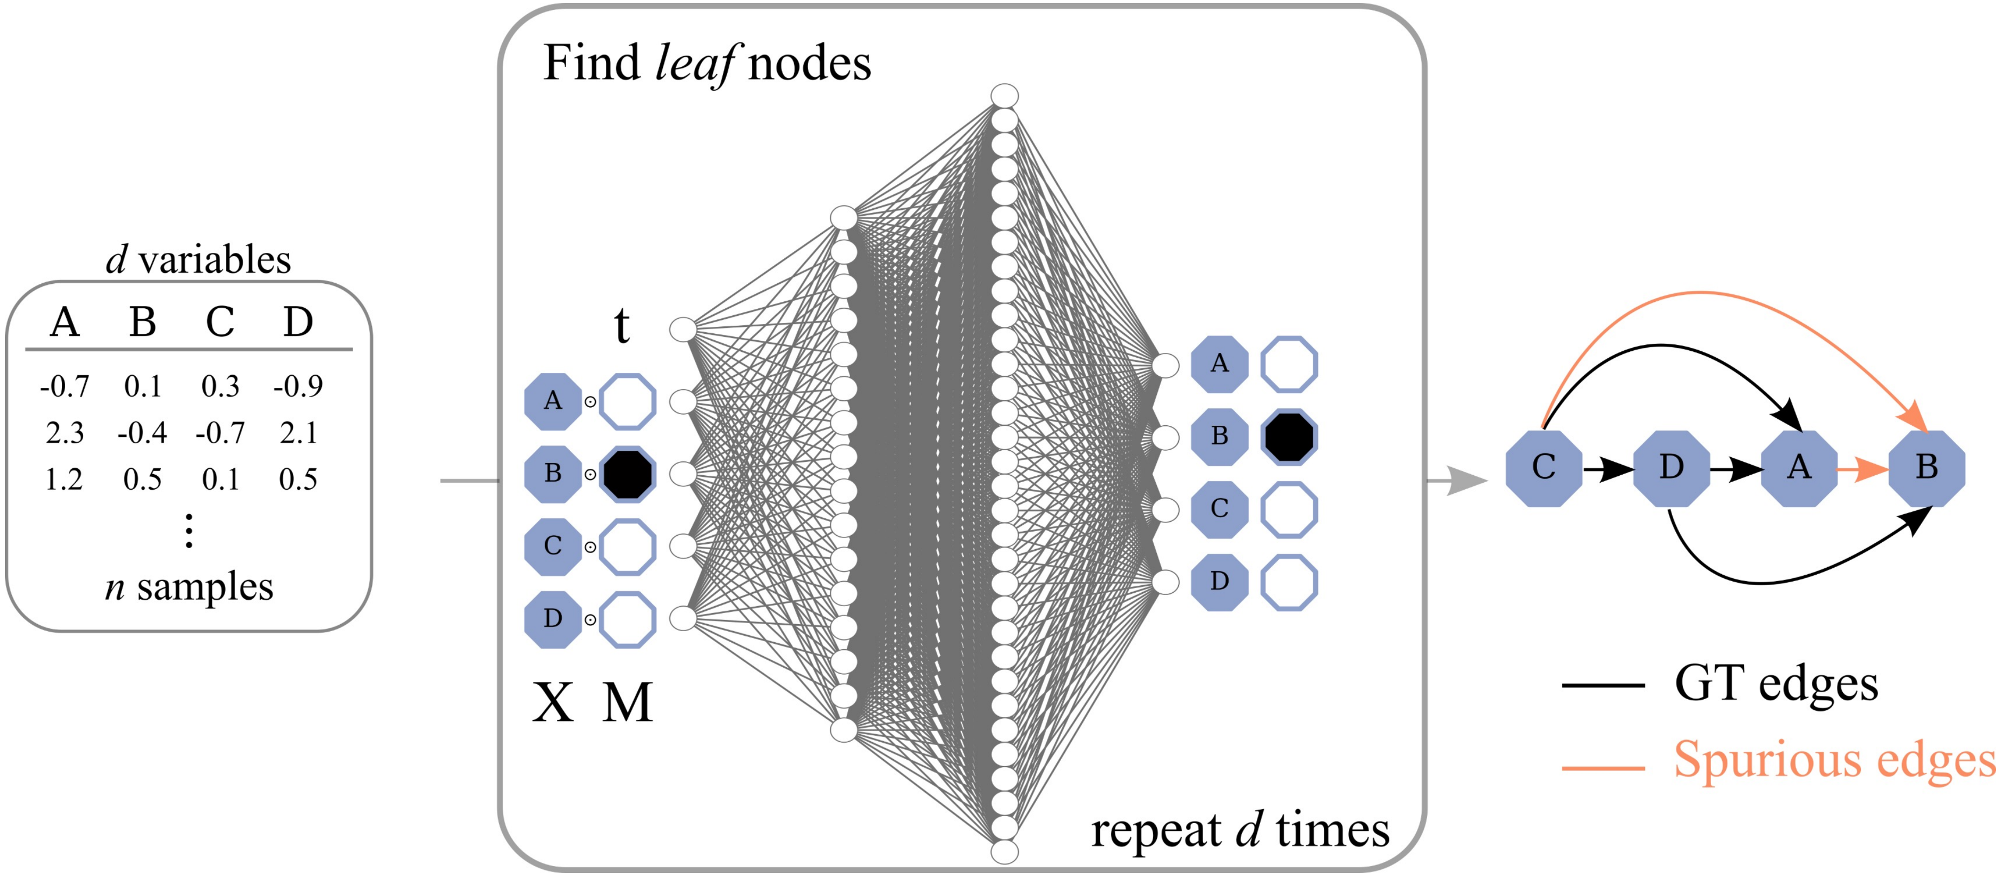

Sanchez et al (2022). Diffusion Models for Causal Discovery via Topological Ordering. arXiv:2210.06201

Conclusion

Hype or Hope?
- Tremendous growth points to hope

Opportunities
- Data manipulation
- Multimodal integration

Challenges
- Sampling speed
  - ✓ Deterministic sampling helps (ODE reparameterization)[1]
- Latent space
  - ➢ Are the Unet deepest representations the latent space?[2]
  - ➢ Is the noise the latent space?[3]
  - ➢ Are the conditions the latent space?[4]

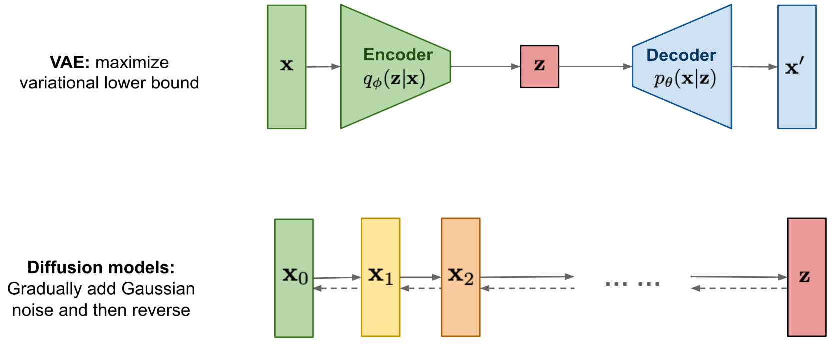

1. Song et al. (2021) Denoising Diffusion Implicit Models. ICLR
2. Kwon et al. (2022) Diffusion Models already have a Semantic Latent Space. arxiv:2210.10960
3. Ho et al. (2020) Denoising Diffusion Probabilistic Models. NeurIPS
4. Abstreiter et al (2021) Diffusion-Based Representation Learning. arXiv:2105.14257

Useful key references, gits to watch etc

https://vios.science/tutorials/diffusion

- Surveys:
  - https://arxiv.org/abs/2209.02646 (general)
  - https://arxiv.org/abs/2209.00796 (general)
  - https://arxiv.org/abs/2209.04747 (vision)
  - https://arxiv.org/abs/2211.07804 (in medical imaging and analysis)
- Github collections:
  - https://github.com/amirhossein-kz/Awesome-Diffusion-Models-in-Medical-Imaging
  - https://github.com/heejkoo/Awesome-Diffusion-Models
- Tutorials:
  - https://lilianweng.github.io/posts/2021-07-11-diffusion-models/ (A nice introductory blog)
  - https://yang-song.github.io/blog/2021/score/ (An amazing blog from one the pioneers)
  - https://arxiv.org/abs/2208.11970 (a VAE perspective)
  - https://cvpr2022-tutorial-diffusion-models.github.io
  - https://huggingface.co/blog/annotated-diffusion
  - https://huggingface.co/docs/diffusers

Thanks to my team

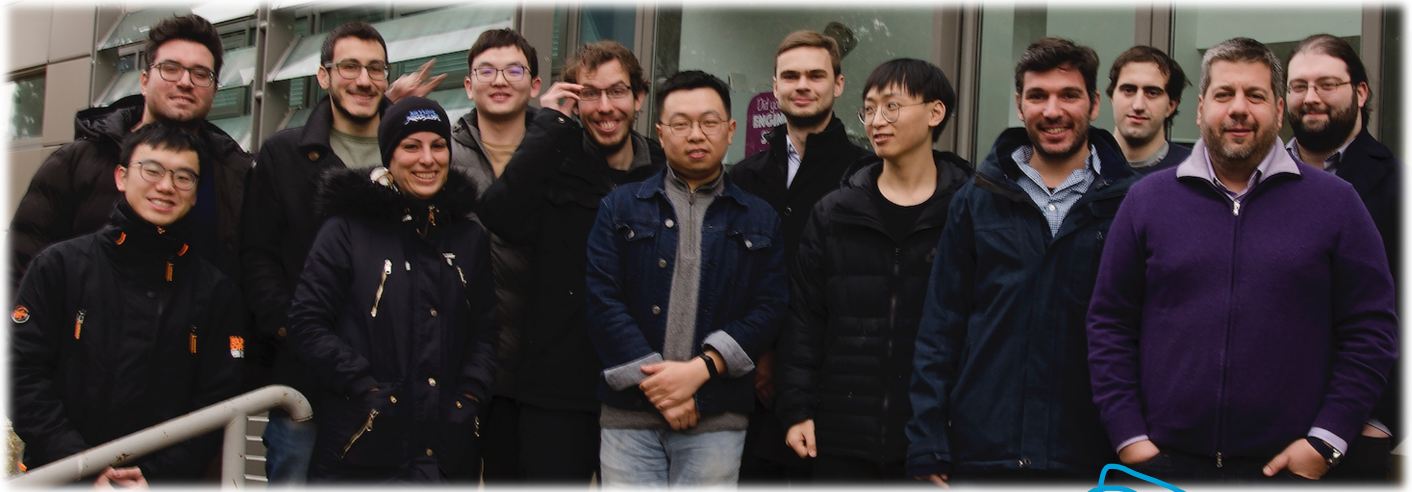






We have several PhD/RA openings if you want to join us!

vios.science

…my collaborators… and those who shared slides

UK

- S. Weir
- A. Smout
- A. O'Neil
- A. Frangi
- D. Newby
- S. Semple
- G. Papanastasiou
- M. Williams

World

- R. Dharmakumar
- T. Arbel
- L. Maier-Hein
- S. Bakas
- X. Papademetris
- N. Merchant
- H. Scharr
- P. Perata

Shared slides:

- K. Packhäuser, A. Maier, FAU
- T. Çukur, Bilkent
- J. Cardoso, W. Diaz Sanz, KCL
- J. Wolleb, P. Cattin, U Basel
- J Chul Ye, KAIST

…funders and you!

EPSRC Engineering and Physical Sciences Research Council

Royal Academy of Engineering

Canon CANON MEDICAL

dstl

Innovate UK

NIH National Institutes of Health Turning Discovery Into Health

BBSRC bioscience for the future

MRC Medical Research Council

The Alan Turing Institute

NVIDIA

iCAIRD

Funding:

NIH 4R01HL091989
NIH 1R01HL136578

EP/N510129/1

EP/P022928/1
BB/N02334X/1
BB/P023487/1

Sotirios A. Tsaftaris, PhD

Email: S.Tsaftaris@ed.ac.uk
URL: http://vios.science