



THE UNIVERSITY
of EDINBURGH



DiMEDIA: Diffusion Models in Medical Imaging and Analysis

<https://vios.science/tutorials/DiMEDIA-2024>

Nefeli Gkouti
Archimedes/Athena RC,
University of Athens



Prof. Sotirios A. Tsaftaris
Archimedes/Athena RC,
University of Edinburgh



Tutorial Schedule



First Half



- ❑ Part 1: Introduction (30 mins)
 - ❖ What? Why? How?
 - ❖ Denoising Diffusion Models
 - ❖ Understanding and Intuition
- ❑ Part 2: Advanced Topics (30 mins)
 - ❖ Sampling Strategies
 - ❖ Conditioning

Second Half

- ❑ Q&A
- ❑ Medical Imaging Applications (30 mins)
 - ❖ Synthesis
 - ❖ Segmentation
 - ❖ Anomaly Detection
 - ❖ Reconstruction
 - ❖ Registration


This tutorial is an abridged version of the successful diffusion tutorials in MICCAI 2023 in Vancouver, and ISBI 2024 in Athens

**Tutorial on
Diffusion Models for Medical Imaging**

Pedro P. Sanchez	Dr. Julia Wolleb	Prof. Jorge Cardoso
Dr. Walter Pinaya	Prof. Dorit Merhof	Prof. Sotirios A. Tsaftaris








**DiMEDIA:
Diffusion Models in Medical Imaging and Analysis**

Yuyang Xue University of Edinburgh	Nefeli Gkouti Archimedes/Athena RC, University of Athens	Dr. Julia Wolleb University of Basel	Pedro Sanchez University of Edinburgh	Prof. Sotirios A. Tsaftaris Archimedes/Athena RC, University of Edinburgh
---------------------------------------	--	---	--	---




Acknowledgement



**Tutorial on
Diffusion Models for Medical Imaging**

Pedro P. Sanchez Dr. Julia Wolleb Prof. Jorge Cardoso
Dr. Walter Pinaya Prof. Dorit Merhof Prof. Sotirios A. Tsaftaris








MICCAI 2023 - Tutorial

Some slides are adopted/inspired/borrowed from these very nice tutorials!

**Denoising Diffusion Models:
A Generative Learning Big Bang**

Jiaming Song Chenlin Meng Arash Vahdat



CVPR 2023 - Tutorial



**DiMEDIA:
Diffusion Models in Medical Imaging and Analysis**

Yuyang Xue Nefeli Gkouti Dr. Julia Wolleb Pedro Sanchez Prof. Sotirios A. Tsaftaris
University of Edinburgh Archimedes/Athena RC, University of Athens University of Basel University of Edinburgh Archimedes/Athena RC, University of Edinburgh



ISBI 2024 - Tutorial

We went to learn a few things about you

We will use wooclap:

- Are you familiar with diffusion models
- Are you familiar with generative model theory?
- What are your primary objectives for today?

More will come...

How to participate?



1 Go to wooclap.com

2 Enter the event code in the top banner

Event code
FROCKY


 Enable answers by SMS

 [Copy participation link](#)

Diffusion Models

What? Why? How?

What? Generative Models



What? Generative Models

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

Density Estimation

$$p_\theta(x)$$

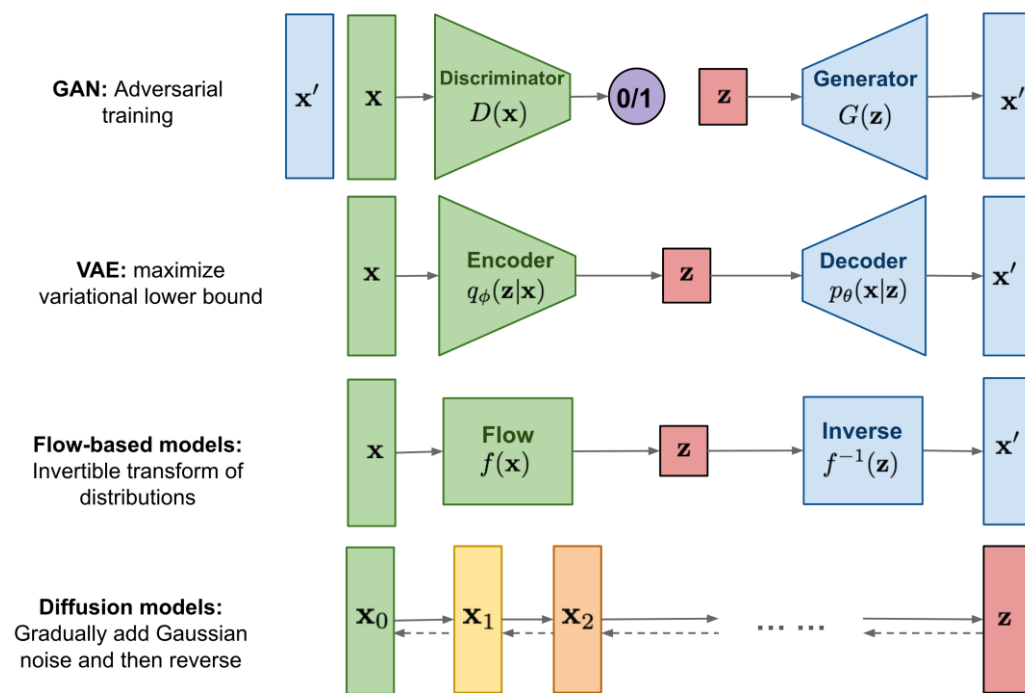
Sampling

$$x_{new} \sim p_\theta$$

Unsupervised Representation Learning

$$z \leftarrow p_\theta(x)$$

What? Generative models



likelihood-based models

Require

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

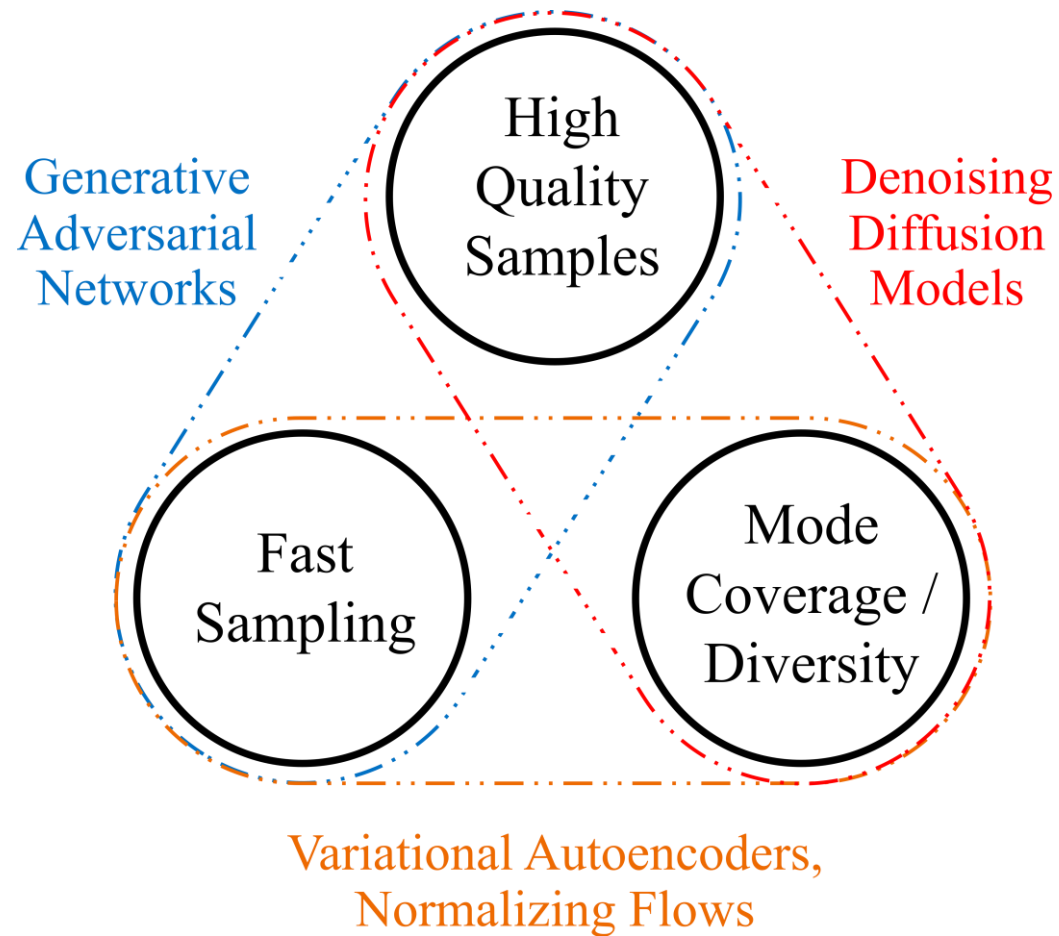
implicit generative models

Require adversarial training:

- notoriously unstable; leading to
- mode collapse


diffusion models bypass both with neat tricks

Sampling Trilemma



Why? Unprecedented Quality

“realistic photo of a cybernetic Eagle”




1. Realism
2. Control
3. Prior



“...
Tribe taking a
selfie ...”

“A dystopian male face made
of volcanic lava, mysterious,
image containing secret
codes”



Why? Community Push

Companies

Big models and data



stability.ai



IDEOGRAM
AI

Open-Source

Ease of Use



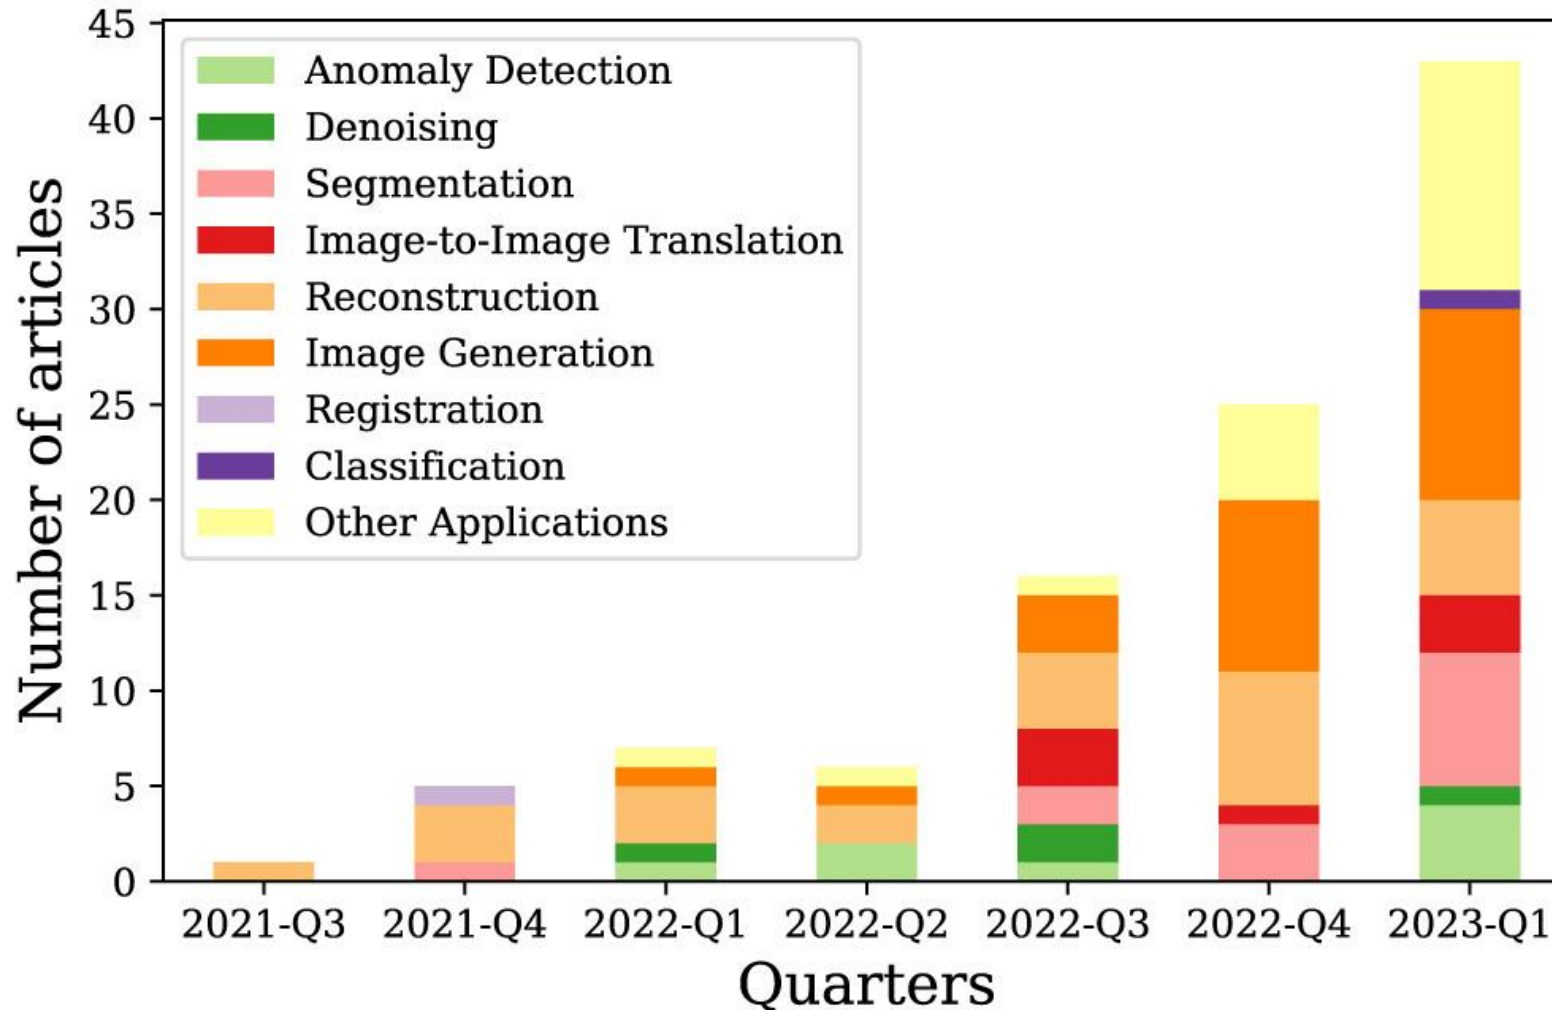
Generative Models

Why? Research Community Push


	ISBI	MICCAI	ICLR	CVPR	ICML	NeurIPS
2024	-	-	7%	12%	-	-
2023	7%	7%	3%	4%	4%	6%
2022	4%	2%	1%	>1%	1%	1%
2021	1%	>1%	>1%	≈0%	>1%	>1%

Percentage of accepted papers in great conferences that include diffusion in their title is growing

Why? Medical Imaging Popularity



Why? Medical Imaging Applications



Realism

Control

Prior

Diffusion Models

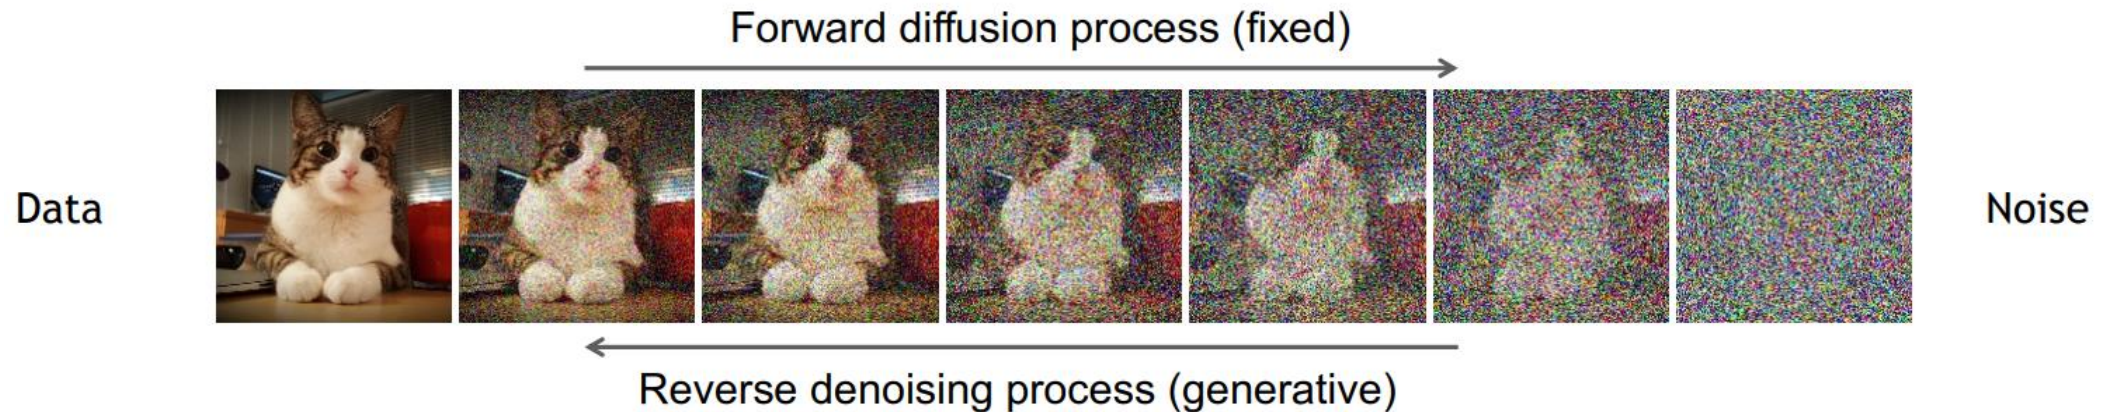
How?



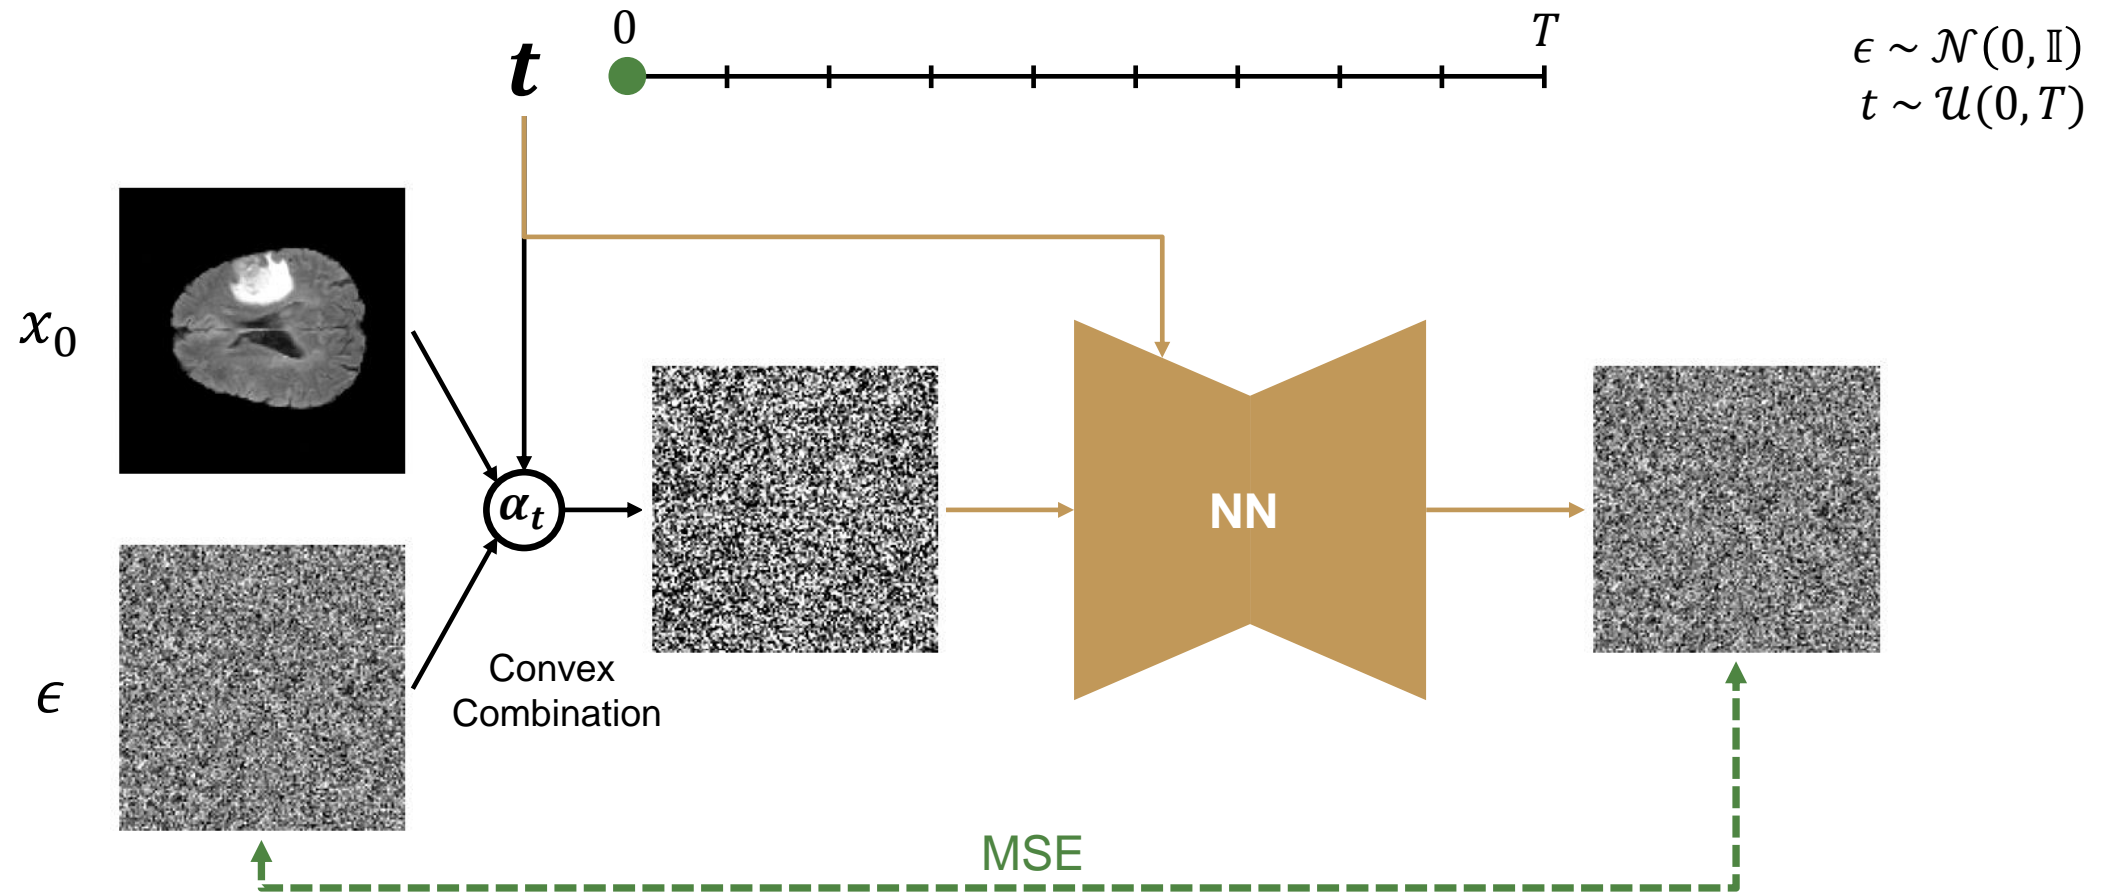
Training by Denoising

How? Training by Denoising

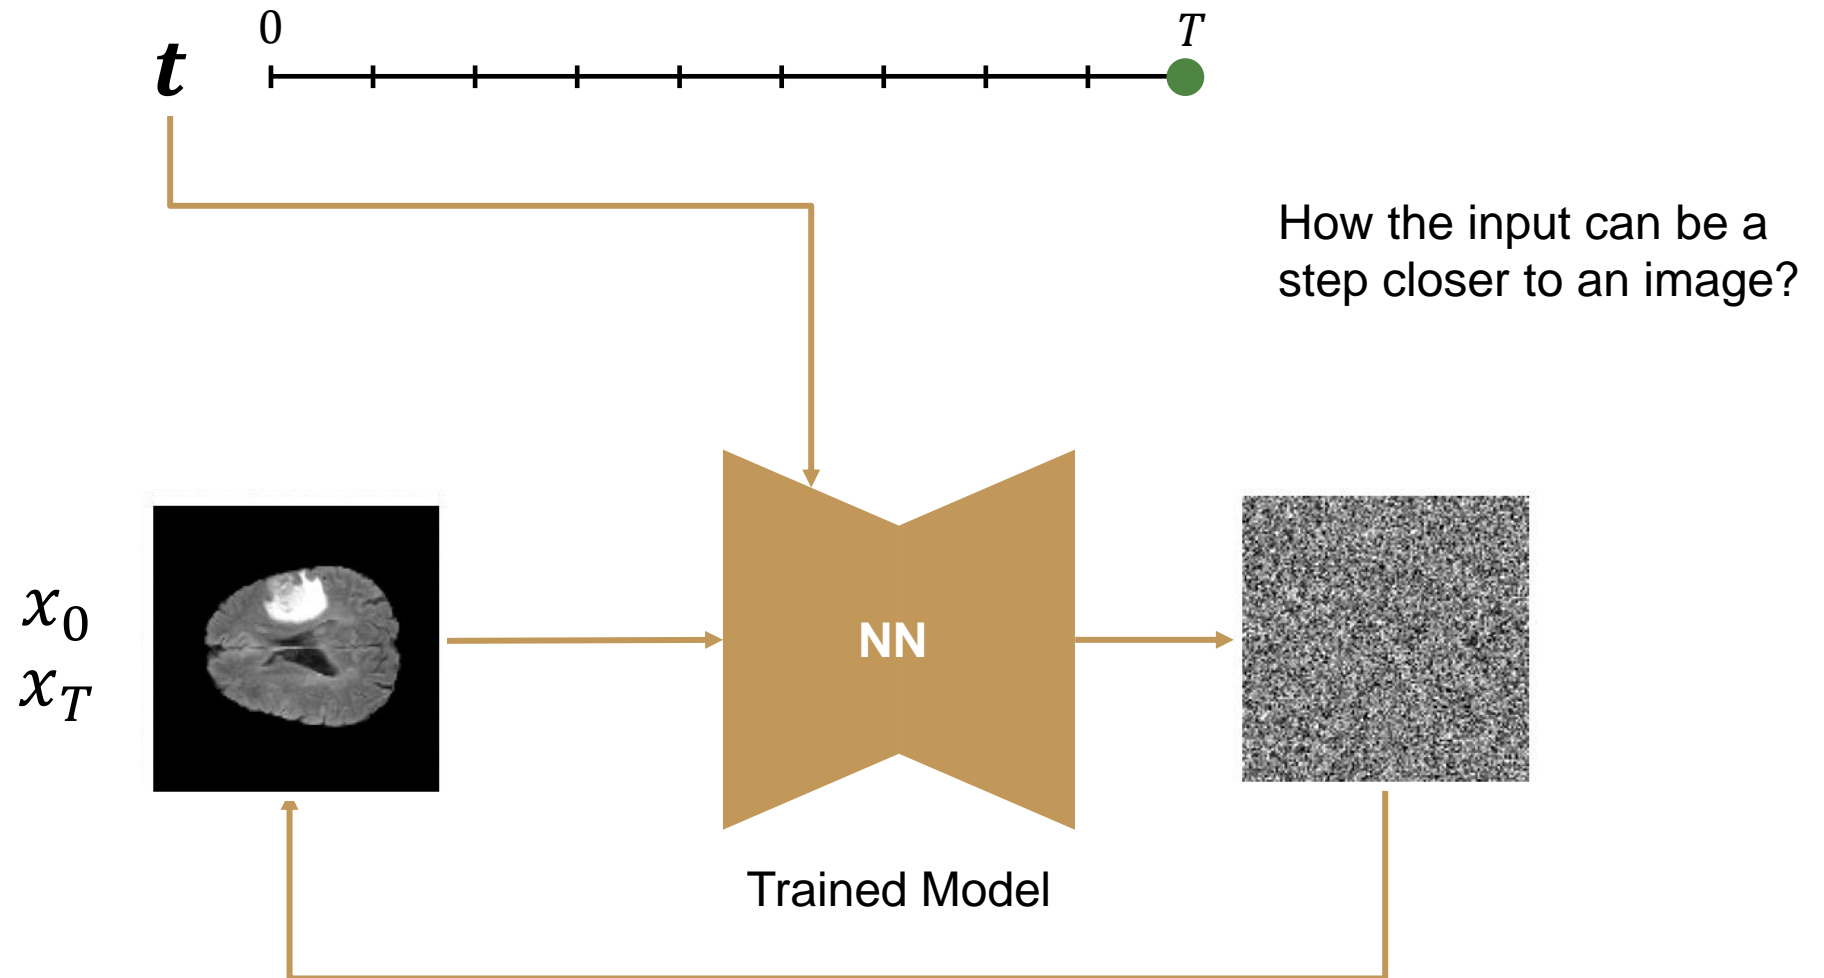
- Add gradually noise
- **Training by denoising**



How? Training by Denoising



How? Inference



Tutorial Schedule

First Half

- ❑ Part 1: Introduction (30 mins)

- ❖ What? Why? How?

- ❖ Denoising Diffusion Models

- ❖ Understanding and Intuition

- ❑ Part 2: Advanced Topics (30 mins)

- ❖ Sampling Strategies

- ❖ Conditioning

We are here!



Second Half

- ❑ Q&A

- ❑ Medical Imaging Applications (30 mins)

- ❖ Synthesis

- ❖ Segmentation

- ❖ Anomaly Detection

- ❖ Reconstruction

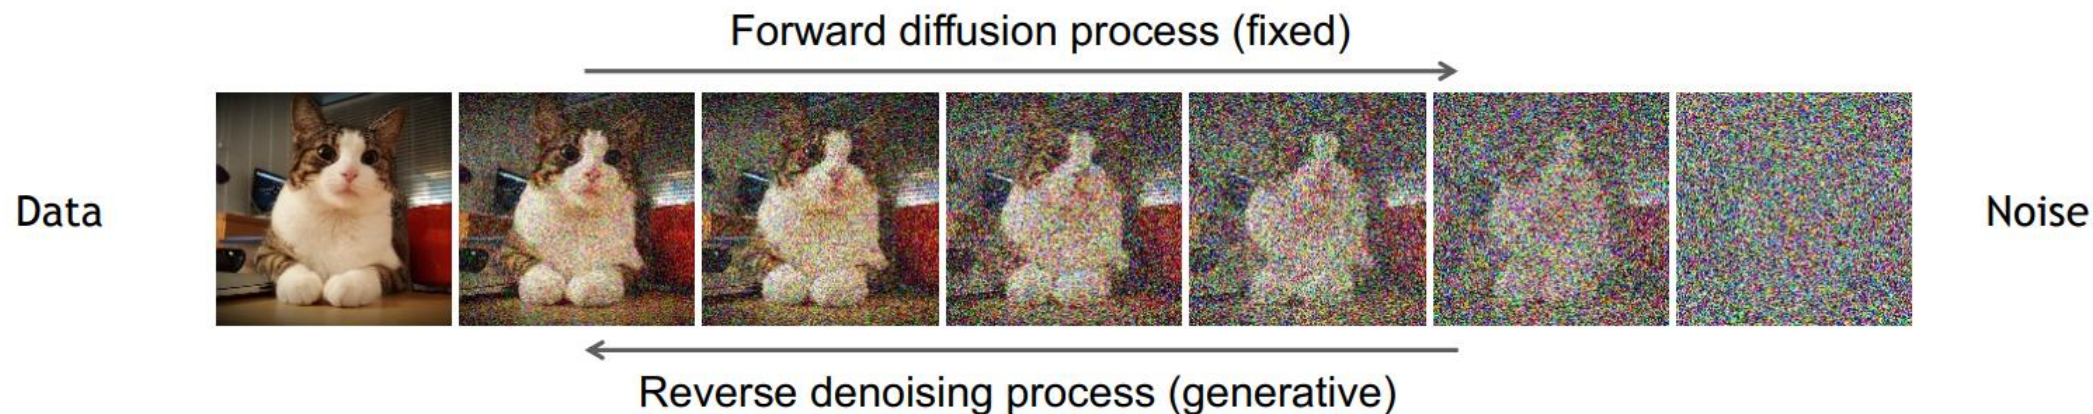
- ❖ Registration

Denoising Diffusion Model

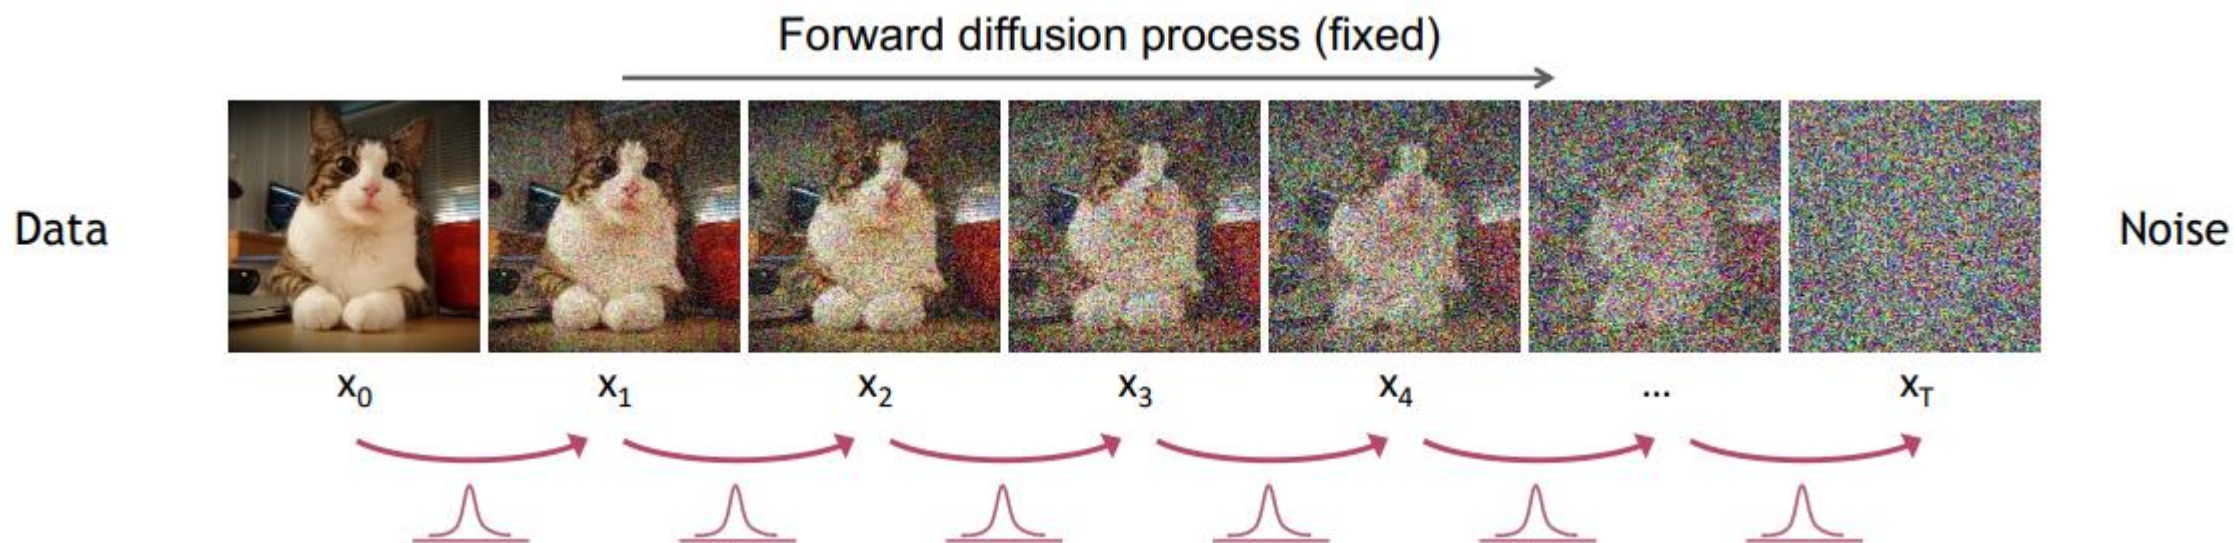
Denoising Diffusion Models

Denoising diffusion models consist of two processes:

- Forward diffusion process that gradually adds noise to input
- Reverse denoising process that learns to generate data by denoising



Forward diffusion process




We use normal distribution to generate a noisy image conditioned on previous image

$$q(x_t|x_{t-1}) = N(x_t; \sqrt{\beta_t}x_{t-1}, (1 - \beta_t)I)$$

Forward diffusion process

joint distribution $\rightarrow q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1})$



Given the step β_s , we can generate a particular step t

$$q(x_t|x_0) = N(x_t; \sqrt{\bar{a}_t}x_0, (1 - \bar{a}_t)I) \quad \text{where } \bar{a}_t = \prod_{s=1}^t (1 - \beta_s)$$

↑
diffusion kernel



For sampling at timestep t: $x_t = \sqrt{\bar{a}_t}x_0 + (1 - \bar{a}_t)\epsilon$ where $\epsilon \sim N(0, I)$

β_t is the noise schedule such that $\bar{a}_t \rightarrow 0 \Rightarrow q(x_T|x_0) \approx N(x_T; 0, I)$

The diffusion kernel at timestep T can be approximated using **standard normal distribution**

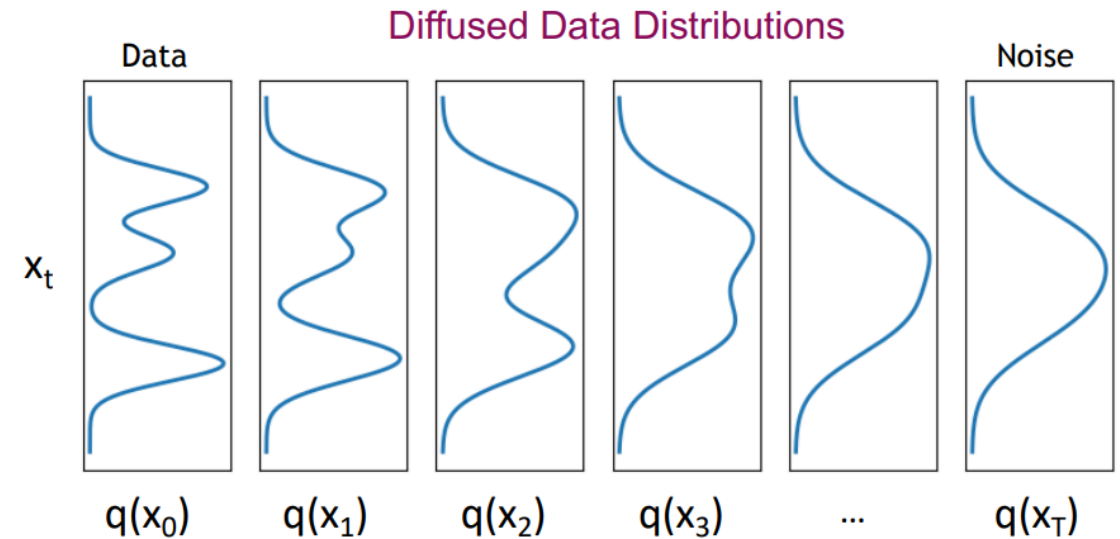


x_T

Forward diffusion process

The diffusion kernel $q(x_t|x_0)$ is Gaussian convolution

$$q(x_t) = \int q(x_0, x_t) dx_0 = \int q(x_0) \underbrace{q(x_t|x_0)}_{\text{diffusion kernel}} dx_0$$



We can sample $x_t \sim q(x_t)$
by first sampling $x_0 \sim q(x_0)$ and then sampling $x_t \sim q(x_t|x_0)$

Generation by Denoising

Recall $q(x_T) \sim N(x_T; 0, I)$

Start from $x_T \sim N(x_T; 0, I)$


Iteratively $x_{t-1} \sim q(x_{t-1}|x_t)$

But $q(x_{t-1}|x_t)$ is **intractable!**

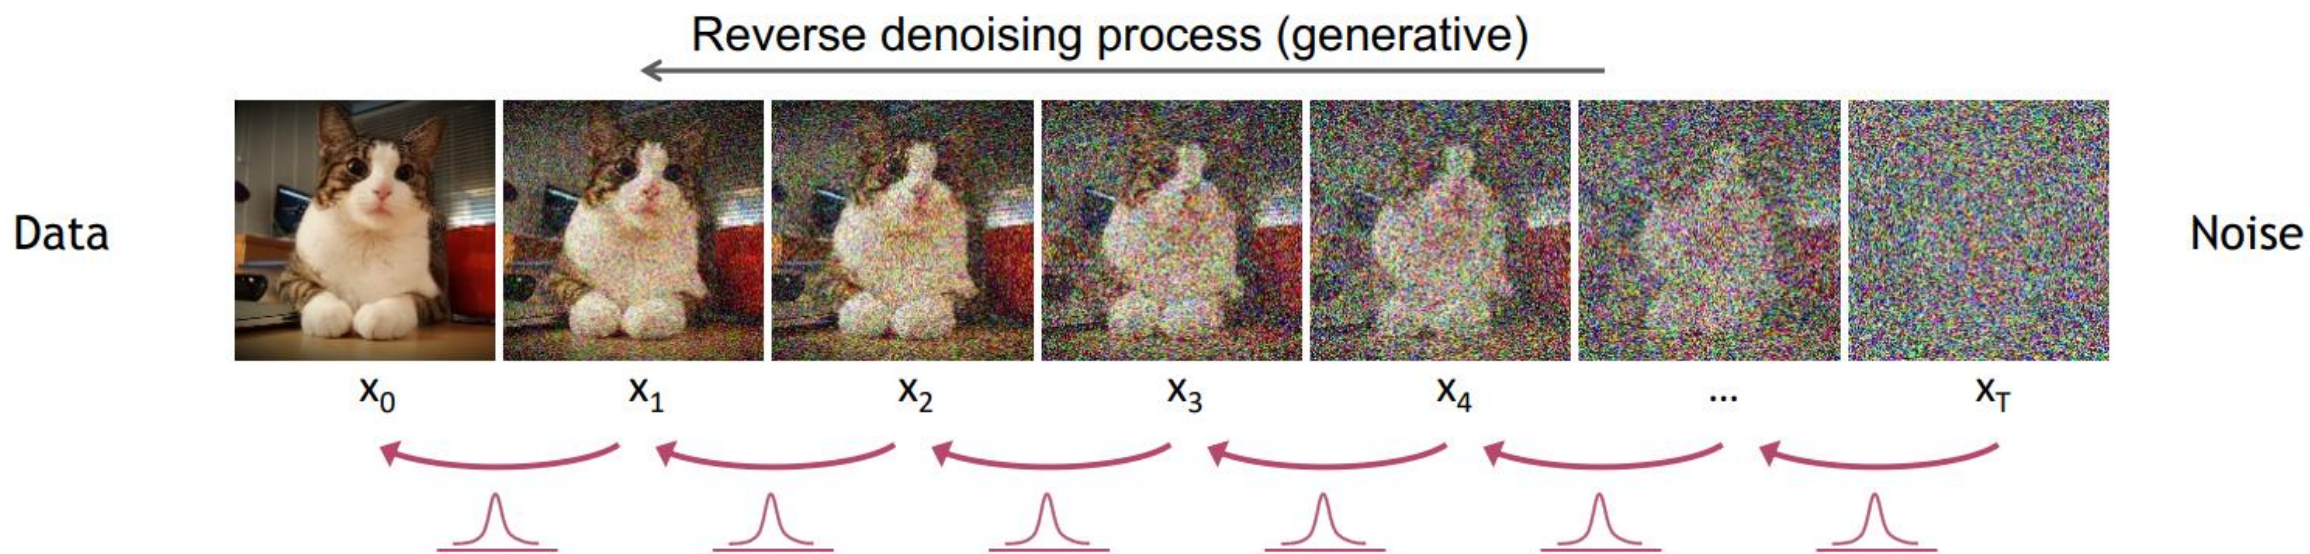
We can approximate it with a normal distribution.

if β_t is small in forward process

We need a **parametric model** to mimic $q(x_{t-1}|x_t)$



Reverse diffusion process



We define a denoising distribution $p_{\theta}(x_{t-1}|x_t)$

Reverse diffusion process

$$p(x_T) = N(x_T; 0, I)$$

$$p_\theta(x_{t-1}|x_t) = N(x_{t-1}; \underbrace{\mu_\theta(x_t, t)}_{\mu_\theta \text{ is a trainable network (like U-net)}}, \sigma_t^2 I)$$

$$\rightarrow \underbrace{p_\theta(x_{0:T})}_{\text{joint distr. of full reverse trajectory}} = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t)$$

For training, we form **variational upper bound (ELBO)** that is commonly used for training variational autoencoders:

$$E_{q(x_0)}[-\log p_\theta(x_0)] \leq E_{q(x_0)q(x_{1:T}|x_0)}[-\log \frac{p_\theta(x_{0:T})}{q(x_{1:T}|x_0)}] =: ELBO$$

Reverse diffusion process

$$E_{q(x_0)}[-\log p_\theta(x_0)] \leq E_{q(x_0)q(x_{1:T}|x_0)}[-\log \frac{p_\theta(x_{0:T})}{q(x_{1:T}|x_0)}] =: ELBO$$

recall $x_t = \sqrt{\bar{a}_t}x_0 + (1 - \bar{a}_t)\epsilon$

After some simple arithmetic operations, the variational objective is:

$$ELBO = E_{x_0 \sim q(x_0), t \sim U\{1, T\}, \epsilon \sim N(0, I)}[\lambda_t \|\epsilon - \overbrace{\epsilon_\theta(\sqrt{\bar{a}_t}x_0 + (1 - \bar{a}_t)\epsilon, t)}^{x_t}\|^2]$$

λ_t is a function of β_t

What does the diffusion model optimise?



How to participate?



1

Go to wooclap.com

2

Enter the event code in the top banner

Event code
FROCKY

Enable answers by SMS

[Copy participation link](#)

Tutorial Schedule

First Half

□ Part 1: Introduction (30 mins)

- ❖ What? Why? How?
- ❖ Denoising Diffusion Models
- ❖ Understanding and Intuition

□ Part 2: Advanced Topics (30 mins)

- ❖ Sampling Strategies
- ❖ Conditioning


Second Half

□ Q&A

□ Medical Imaging Applications (30 mins)

- ❖ Synthesis
- ❖ Segmentation
- ❖ Anomaly Detection
- ❖ Reconstruction
- ❖ Registration

 We are here!



Understanding and Intuition

Score Function

Likelihood models want to learn $p_{\theta}(x)$ directly

Diffusion models want to learn the score $\nabla_x \log p_{\theta}(x)$

(i.e. the gradient with respect to the input of the loglikelihood)

a distribution can be written as:

$$p_{\theta}(x) = \frac{e^{-f_{\theta}(x)}}{Z_{\theta}}$$

Score Function

a distribution can be written as:

$$p_{\theta}(\mathbf{x}) = \frac{e^{-f_{\theta}(\mathbf{x})}}{Z_{\theta}}$$


$$\log p_{\theta}(\mathbf{x}) = \log e^{-f_{\theta}(\mathbf{x})} - \log Z_{\theta}$$

$$\nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}) = -\nabla_{\mathbf{x}} f_{\theta}(\mathbf{x}) - \nabla_{\mathbf{x}} \log Z_{\theta}$$

\Downarrow
 ϵ_{θ}

How to learn it?

Mixture of two Gaussians
Score function (the vector field)
Density function (contours)



The score is pointing to the areas of biggest mass


Denoising Score Matching

How to **learn** the score?

$$\mathbb{E}_{q(\mathbf{x})} \left\| \underbrace{\boldsymbol{\epsilon}_{\theta}(\mathbf{x})}_{\text{Diffusion Model}} - \underbrace{\nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x})}_{\text{Score}} \right\|_2^2$$

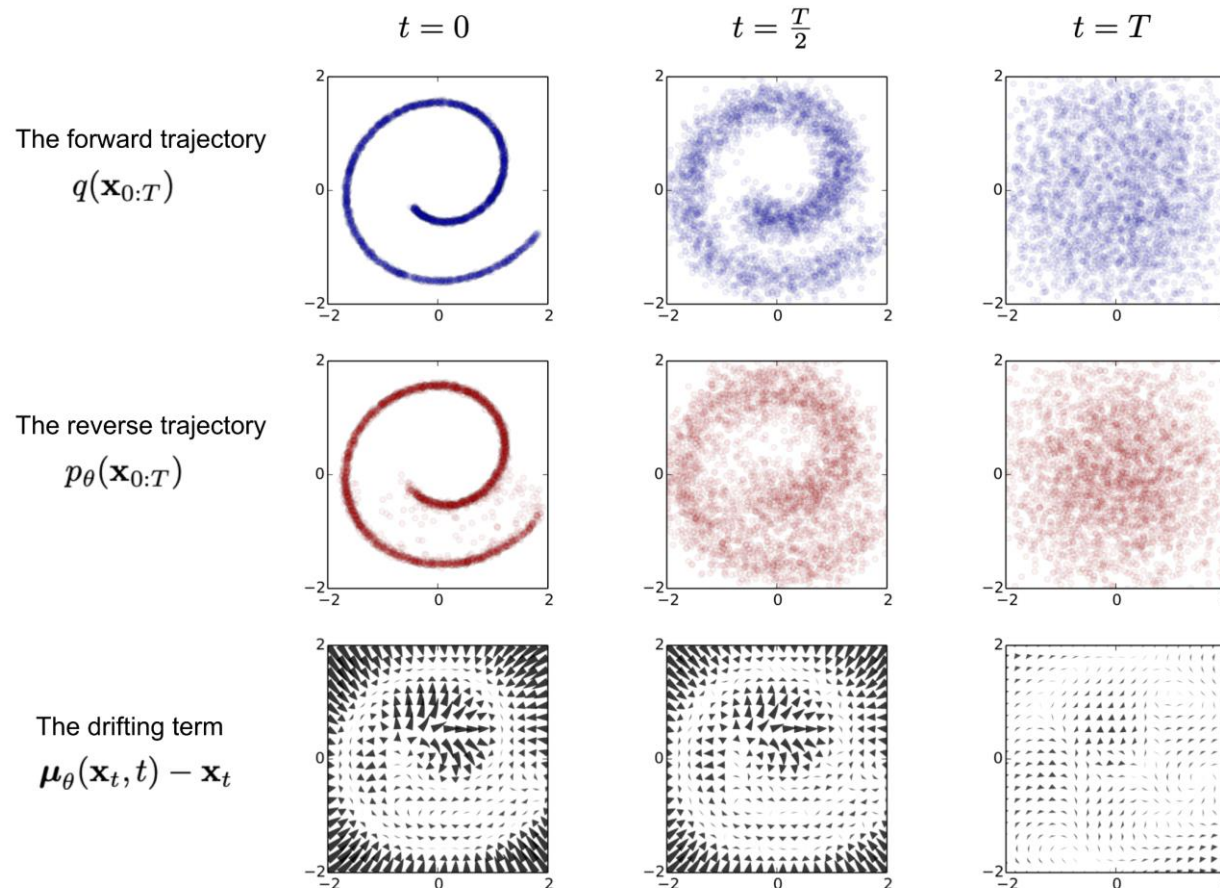
$$\mathbb{E}_{q(\mathbf{x})} \left\| \boldsymbol{\epsilon}_{\theta}(\mathbf{x}) - \nabla_{\mathbf{x}} \log q(\mathbf{x}_t | \mathbf{x}) \right\|_2^2$$

Perturbation at many scales



Learning in **low** density regions


Diffusion Models Learn the Gradient



$$\nabla_x \log p(\mathbf{x})$$

Diffusion and Differential Equations


- Perturbation process is a Stochastic Differential Equation (SDE)
 - From complex to simple
 - Allow different values for SDE modelling



$$dx = f(x, t)dt + g(t)dw$$

Reversing the Process is Generation

- Samplers are discrete solutions of the reverse-time SDE




$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t) \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt + g(t) d\mathbf{w}$$

Fourier Transform

$$\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x} + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$$


Fourier Transform

$$\mathcal{F}(\mathbf{x}_t) = \sqrt{\alpha_t} \mathcal{F}(\mathbf{x}) + \sqrt{1 - \alpha_t} \mathcal{F}(\boldsymbol{\epsilon})$$



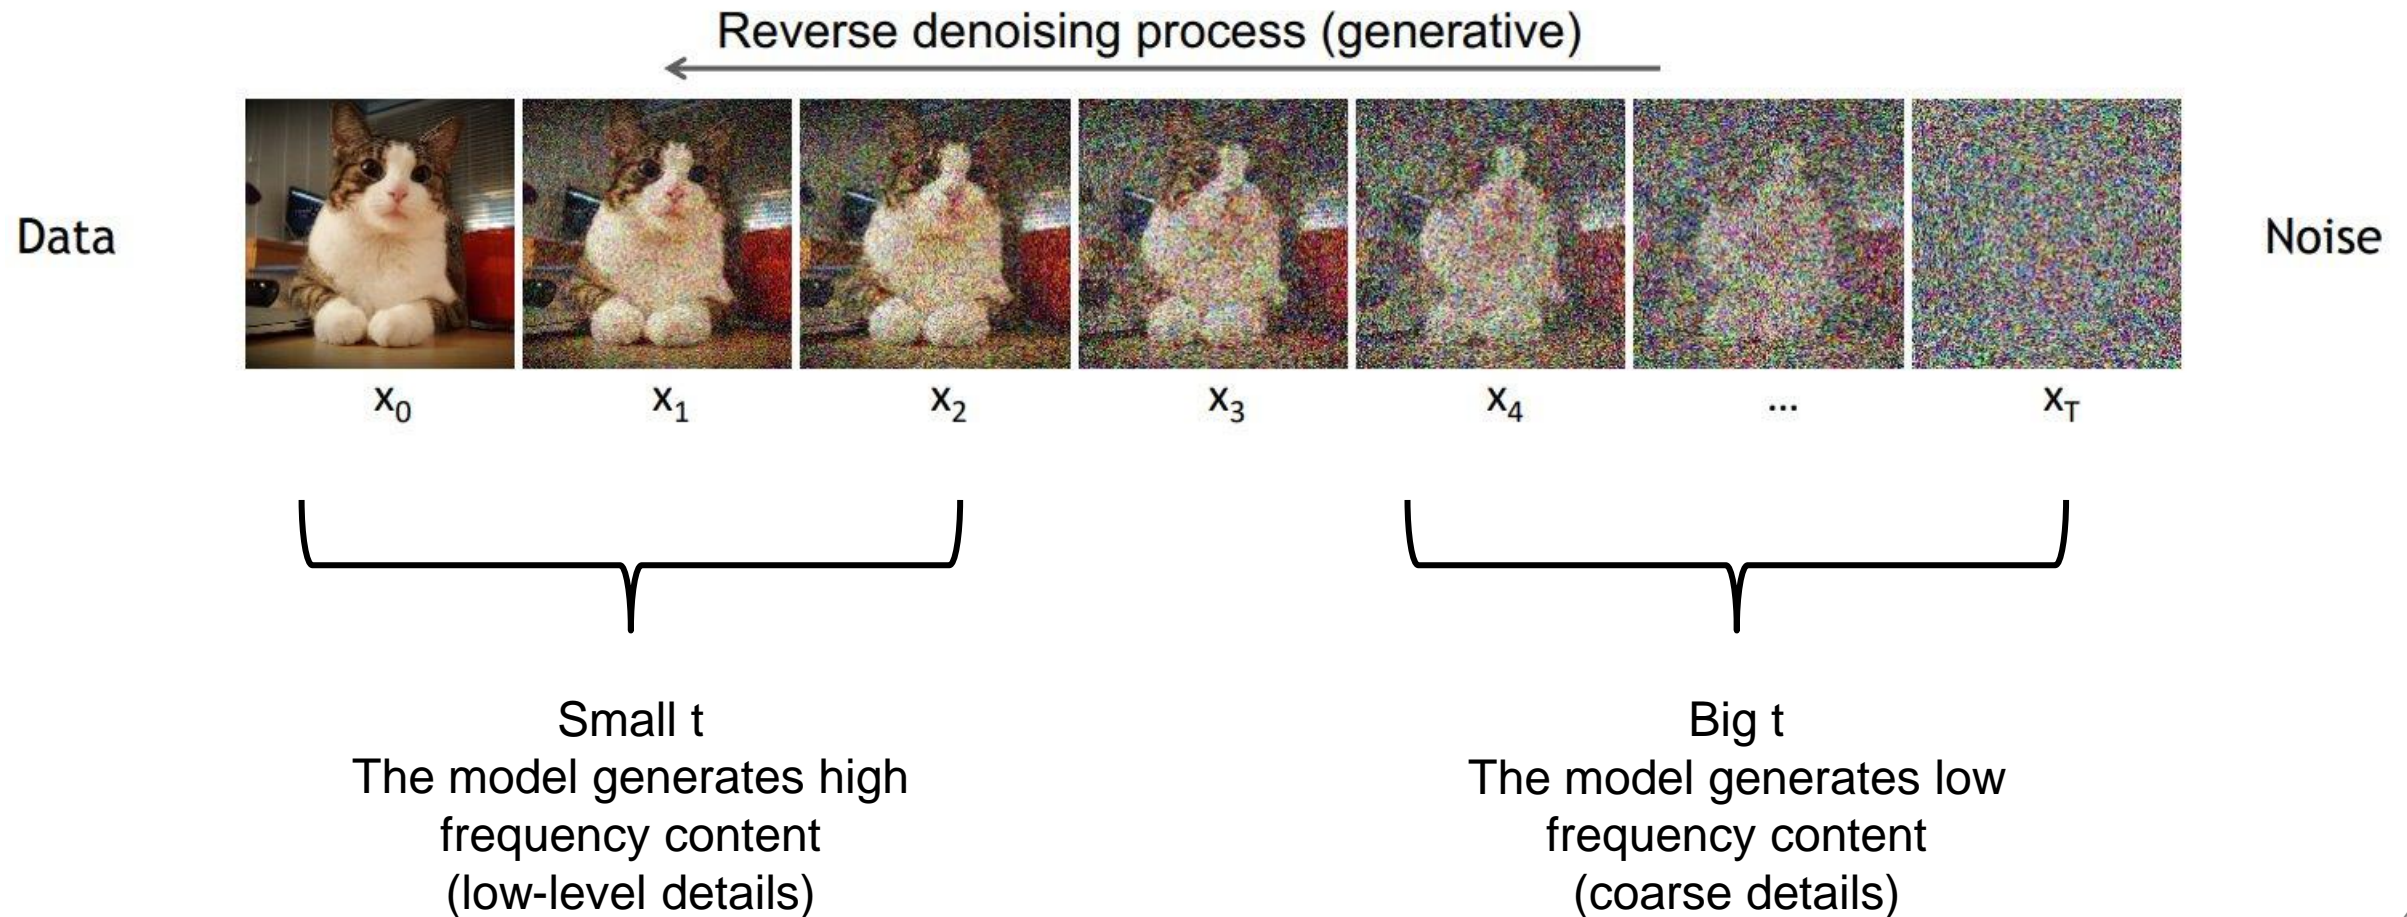
Small t
 $\bar{\alpha}_t \sim 1$

Big t
 $\bar{\alpha}_t \sim 0$


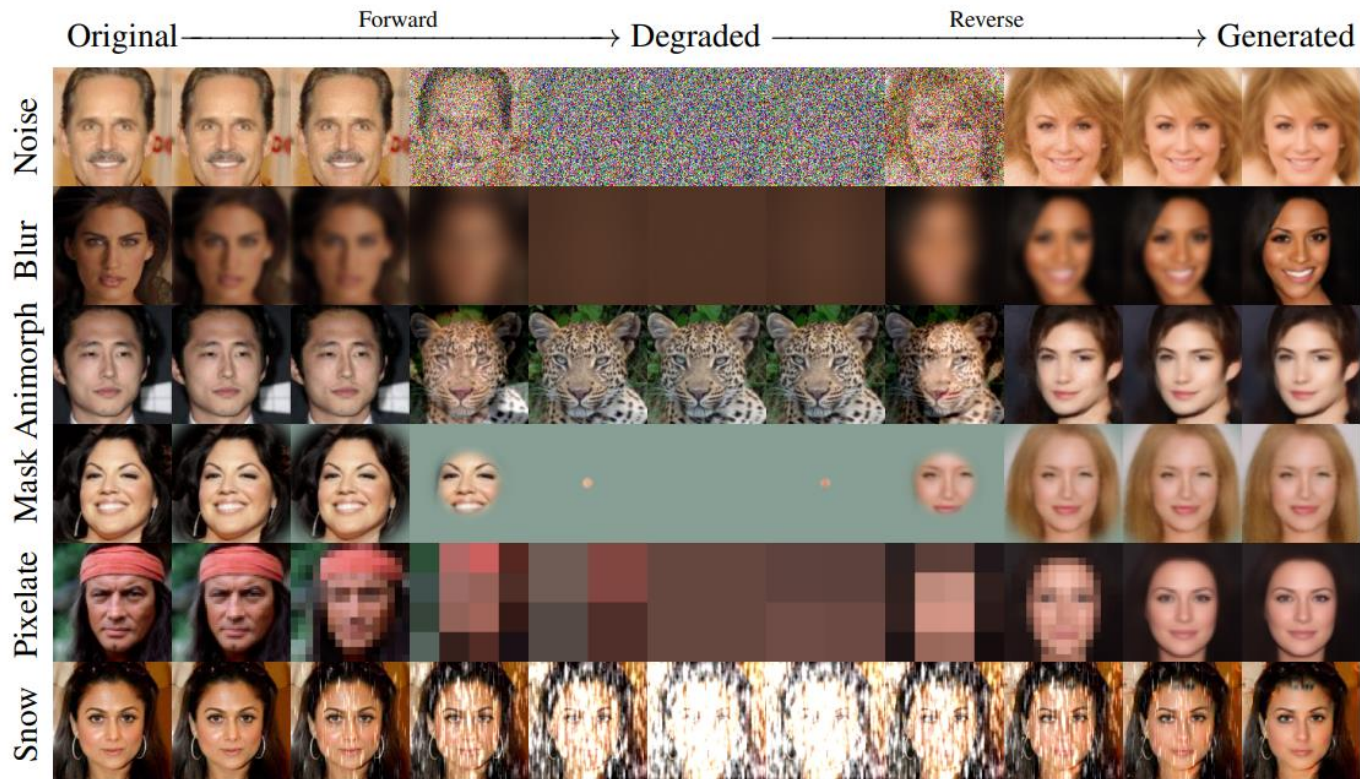


Slide inspired in CVPRs 2022 tutorial on diffusion models

Content – Detail Tradeoff




Gaussian Perturbation?



- [1] Daras, Giannis, et al. "Soft diffusion: Score matching for general corruptions." arXiv preprint arXiv:2209.05442 (2022).
- [2] Bansal, Arpit, et al. "Cold diffusion: Inverting arbitrary image transforms without noise." arXiv preprint arXiv:2208.09392 (2022).
- [3] Kascenas, Antanas, et al. "The role of noise in denoising models for anomaly detection in medical images." Medical Image Analysis (2023): 102963.

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



Unet!

Or transformers
Or VQ-VAEs
Or...

How do you think a diffusion model might compare to GANs in terms of output quality and stability?



How to participate?



1

Go to wooclap.com

2

Enter the event code in the top banner

Event code

FROCKY



Enable answers by SMS

[Copy participation link](#)

Tutorial Schedule

First Half

- ❑ Part 1: Introduction (30 mins)
 - ❖ What? Why? How?
 - ❖ Denoising Diffusion Models
 - ❖ Understanding and Intuition
- ❑ Part 2: Advanced Topics (30 mins)
 - ❖ Sampling Strategies
 - ❖ Conditioning




We are here!

Second Half

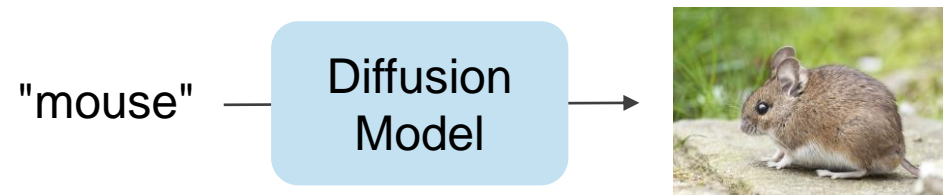
- ❑ Q&A
- ❑ Medical Imaging Applications (30 mins)
 - ❖ Synthesis
 - ❖ Segmentation
 - ❖ Anomaly Detection
 - ❖ Reconstruction
 - ❖ Registration

Part 2 – Advanced Topics

- Sampling Strategies

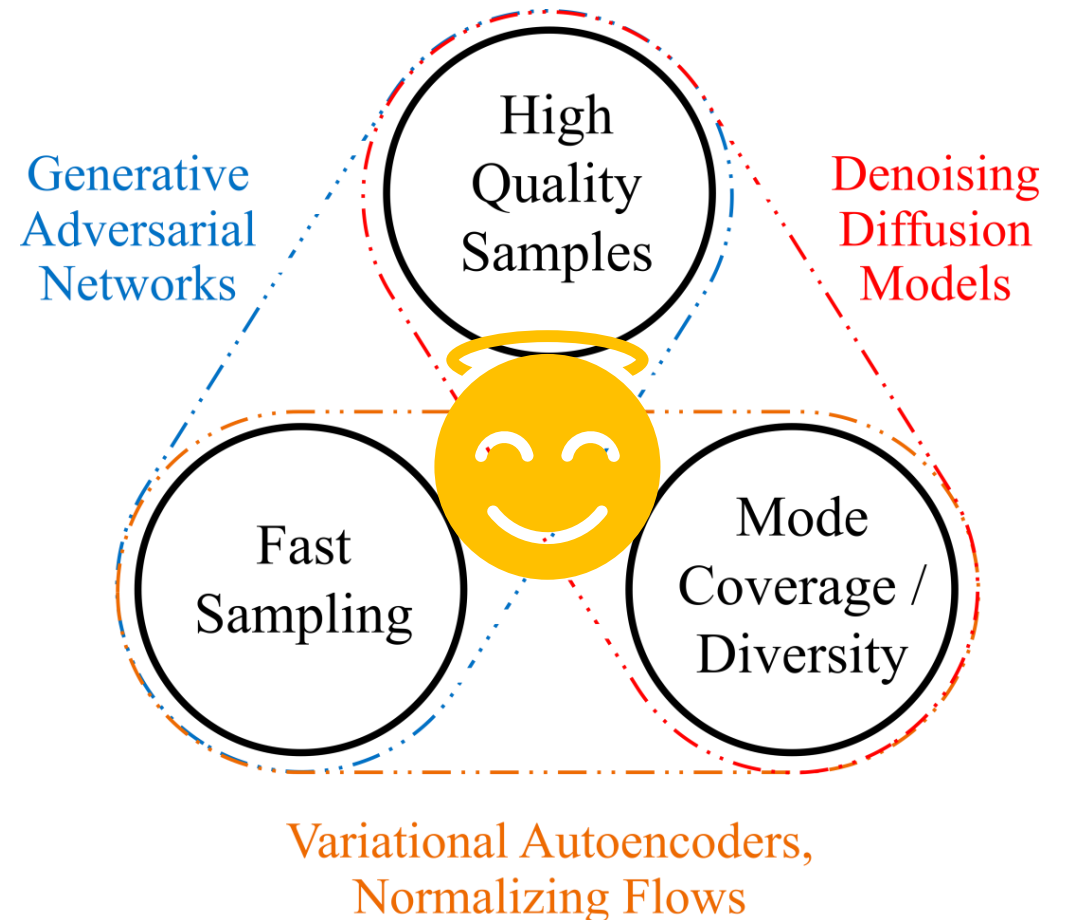


- Conditioning Mechanisms




An inherent drawback: slowness

- Facts about DMs
 - Training
 - 150-1000 V100 GPU days
 - Sampling
 - 50k samples, 5 A100 GPU days



How to accelerate diffusion models?




Naïve acceleration methods

- such as reducing diffusion time steps in training or sampling
 - Leading to immediate worse performance.

We need something clever!

Acceleration: Forward Process




$$\mathbf{x}_0 \rightarrow \cdots \mathbf{x}_t \rightarrow \mathbf{x}_{t+1} \rightarrow \cdots \rightarrow \mathbf{x}_T$$
$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$$


- Does it have to be a Markovian process?
- Is there any faster diffusion process?

Denoising Diffusion Implicit Models (DDIM)

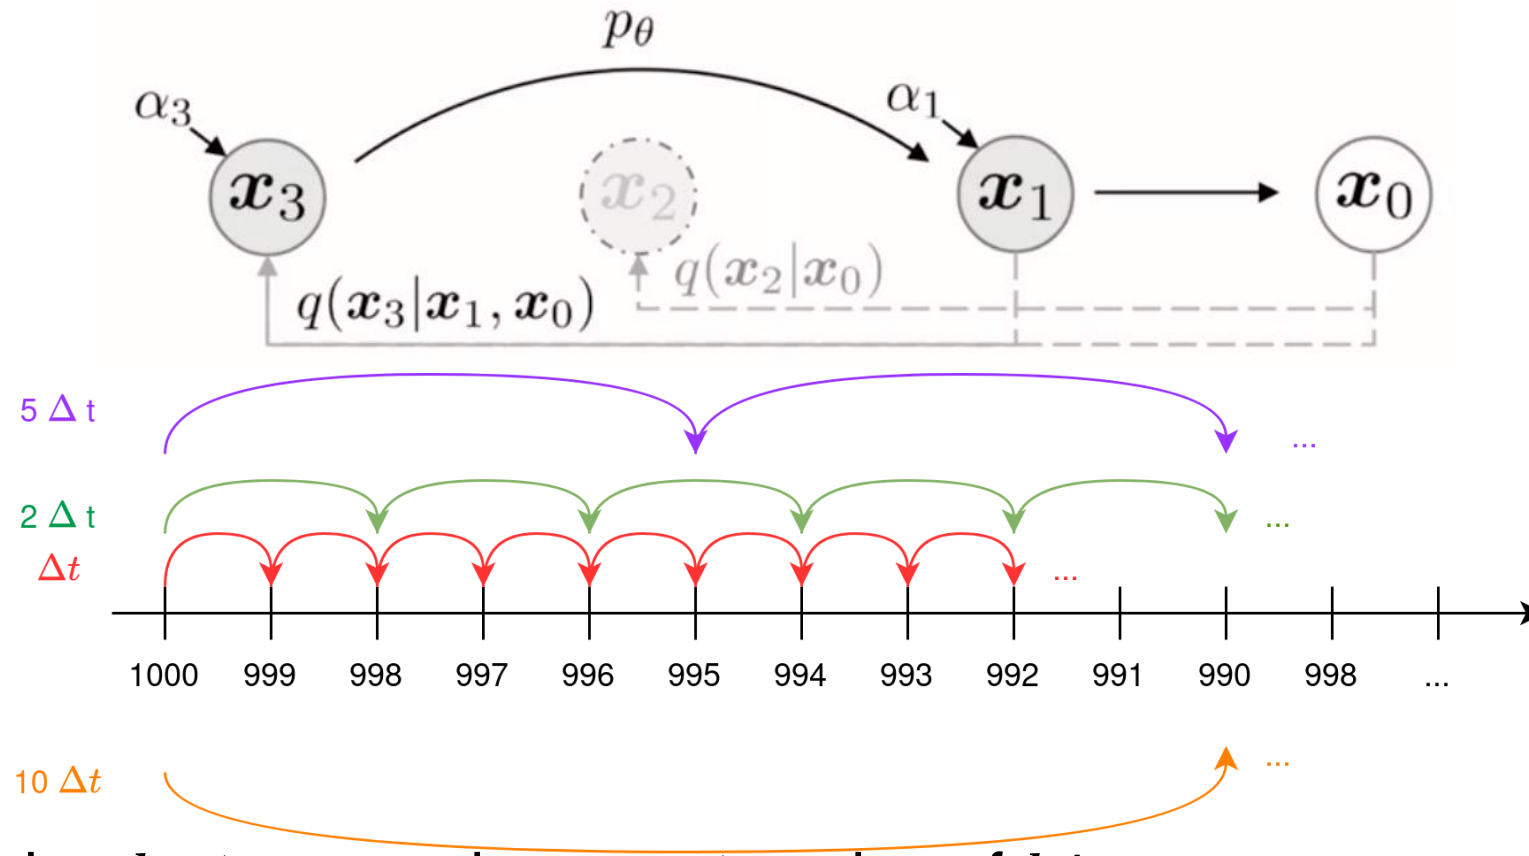
- Diffusion Model does not need to be Markovian!
 - DDPM forward process: $q_{\sigma}(\mathbf{x}_t | \mathbf{x}_{t-1})$
- The loss function for DDPM: $L(\theta) = [\|\epsilon - \epsilon_{\theta}(\mathbf{x}_t, t)\|^2]$



- Define a new forward process: $q_{\sigma}(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{x}_0)$



Acceleration: Skip steps



- By skipping k steps, we have a step size of $k\Delta t$.
- Sampling is k times faster.
- We trade image quality for speed.

DDIM Result

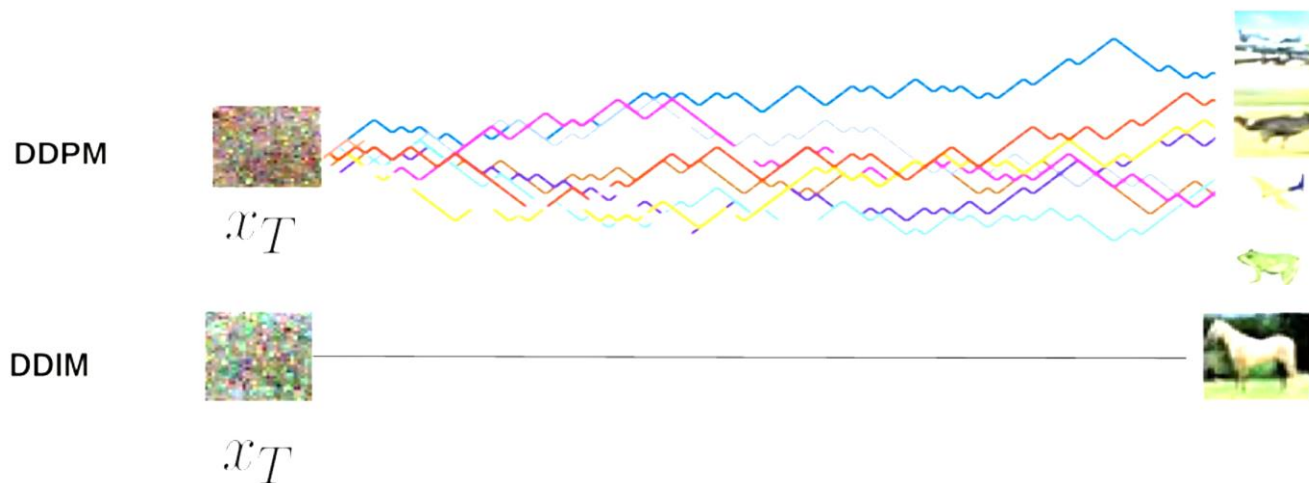
- DDIM: Non-Markovian process but 10-50x faster!!
 - We can directly use a pretrained DDPM to do sampling acceleration

Table 1: CIFAR10 and CelebA image generation measured in FID. $\eta = 1.0$ and $\hat{\sigma}$ are cases of **DDPM** (although Ho et al. (2020) only considered $T = 1000$ steps, and $S < T$ can be seen as simulating DDPMs trained with S steps), and $\eta = 0.0$ indicates **DDIM**.

S	CIFAR10 (32×32)					CelebA (64×64)				
	10	20	50	100	1000	10	20	50	100	1000
$\eta = 0.0$	13.36	6.84	4.67	4.16	4.04	17.33	13.73	9.17	6.53	3.51
$\eta = 0.2$	14.04	7.11	4.77	4.25	4.09	17.66	14.11	9.51	6.79	3.64
$\eta = 0.5$	16.66	8.35	5.25	4.46	4.29	19.86	16.06	11.01	8.09	4.28
$\eta = 1.0$	41.07	18.36	8.01	5.78	4.73	33.12	26.03	18.48	13.93	5.98
$\hat{\sigma}$	367.43	133.37	32.72	9.99	3.17	299.71	183.83	71.71	45.20	3.26


Denoising Diffusion Implicit Models (DDIM)

$$\mathbf{x}_{t-1} = \underbrace{\sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_{\theta}^{(t)}(\mathbf{x}_t)}{\sqrt{\alpha_t}} \right)}_{\text{“predicted } \mathbf{x}_0\text{”}} + \underbrace{\sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \epsilon_{\theta}^{(t)}(\mathbf{x}_t)}_{\text{“direction pointing to } \mathbf{x}_t\text{”}} + \underbrace{\sigma_t \epsilon_t}_{\text{random noise}}$$




Song, J., Meng, C., & Ermon, S. (2020). Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

Image Interpolation




↓ DDIM noise encoding




A

Linear Combination
 $(1-\alpha)A + \alpha B$




↓



B

↓ DDIM noise decoding

Output



α

0

0.1

0.2

0.4

0.5

0.6


0.8

1

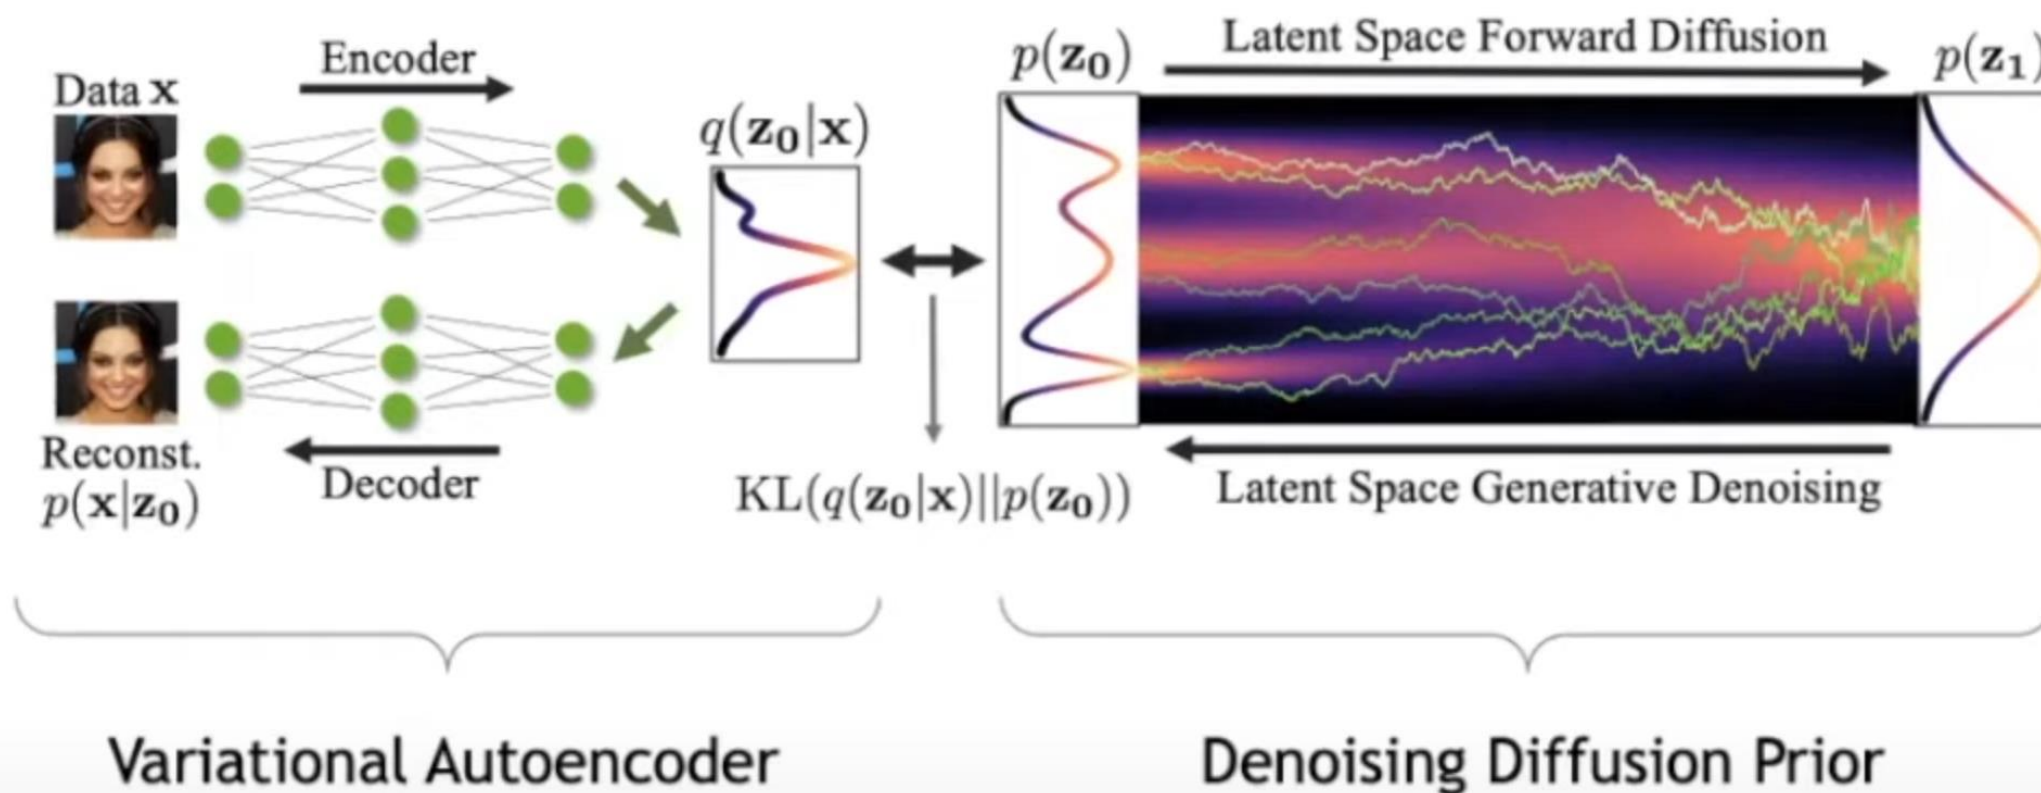
Progressive Distillation for Fast Sampling



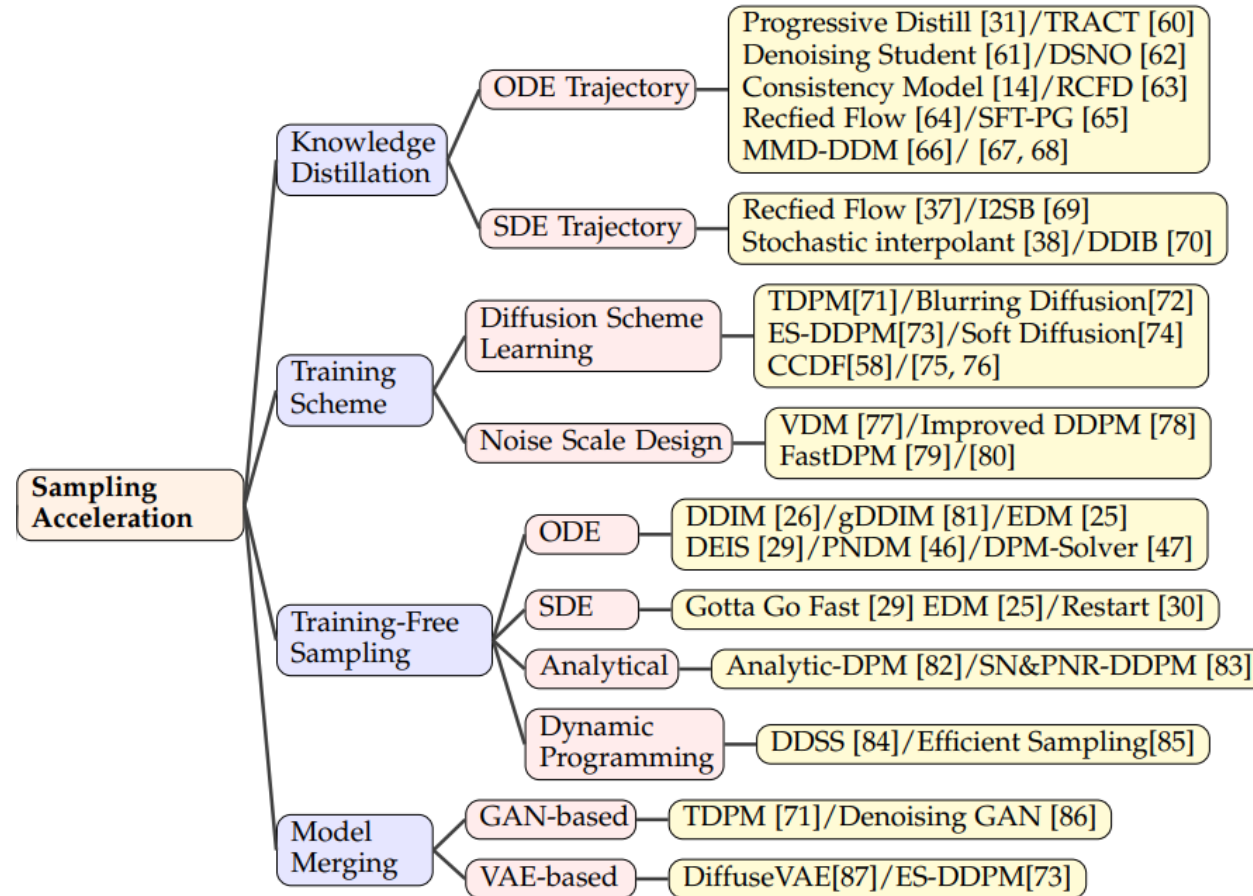
Progressive Distillation for Fast Sampling



Latent-space Diffusion Models



Taxonomy of Sampling Acceleration



Tutorial Schedule

First Half

- ❑ Part 1: Introduction (30 mins)
 - ❖ What? Why? How?
 - ❖ Denoising Diffusion Models
 - ❖ Understanding and Intuition
- ❑ Part 2: Advanced Topics (30 mins)
 - ❖ Sampling Strategies
 - ❖ Conditioning




We are here!

Second Half

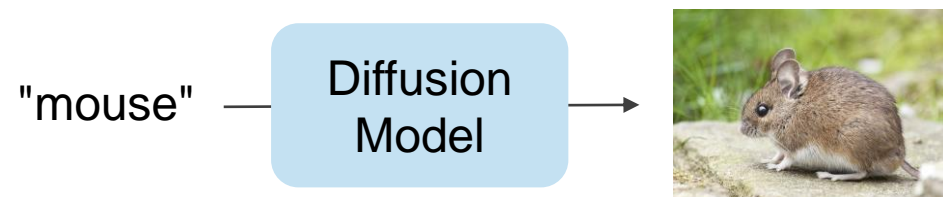
- ❑ Q&A
- ❑ Medical Imaging Applications (30 mins)
 - ❖ Synthesis
 - ❖ Segmentation
 - ❖ Anomaly Detection
 - ❖ Reconstruction
 - ❖ Registration

Part 2 – Advanced Topics

- Sampling Strategies



- Conditioning Mechanisms




Conditioning example: Image-to-image generation



Inpainting

Trends in Conditioning



The number of papers on controllable generation based on T2I diffusion models.

Cao P, Zhou F, Song Q, et al. Controllable Generation with Text-to-Image Diffusion Models: A Survey[J]. arXiv preprint arXiv:2403.04279, 2024.

How to evaluate generated samples?

- How to evaluate a generative model?
 - Quality
 - Diversity
- Low FID
 - Not real enough
- High diversity
 - Means hard to control



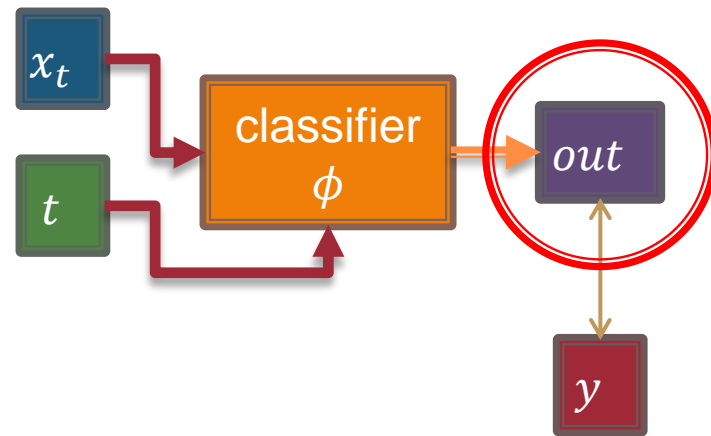
- We want to introduce condition so that the generation can be controllable

Generation Direction

- $\nabla_{x_t} \log p(\mathbf{x}_t)$: the direction of generation
 - SDE, generate with randomness
- What about a condition y ?
 - $\nabla_{x_t} \log p(\mathbf{x}_t|y)$: the direction of generation, conditioned on y
- We have $p(x_t|y) = \frac{p(y|x_t)p(x_t)}{p(y)}$, then we have
 - $\nabla_{x_t} \log p(\mathbf{x}_t|y) = \nabla_{x_t} \log p(y|\mathbf{x}_t) + \nabla_{x_t} \log p(\mathbf{x}_t)$
 - Direction of the condition
 - Direction of unconditional generation

Classifier Guidance - Post Editing

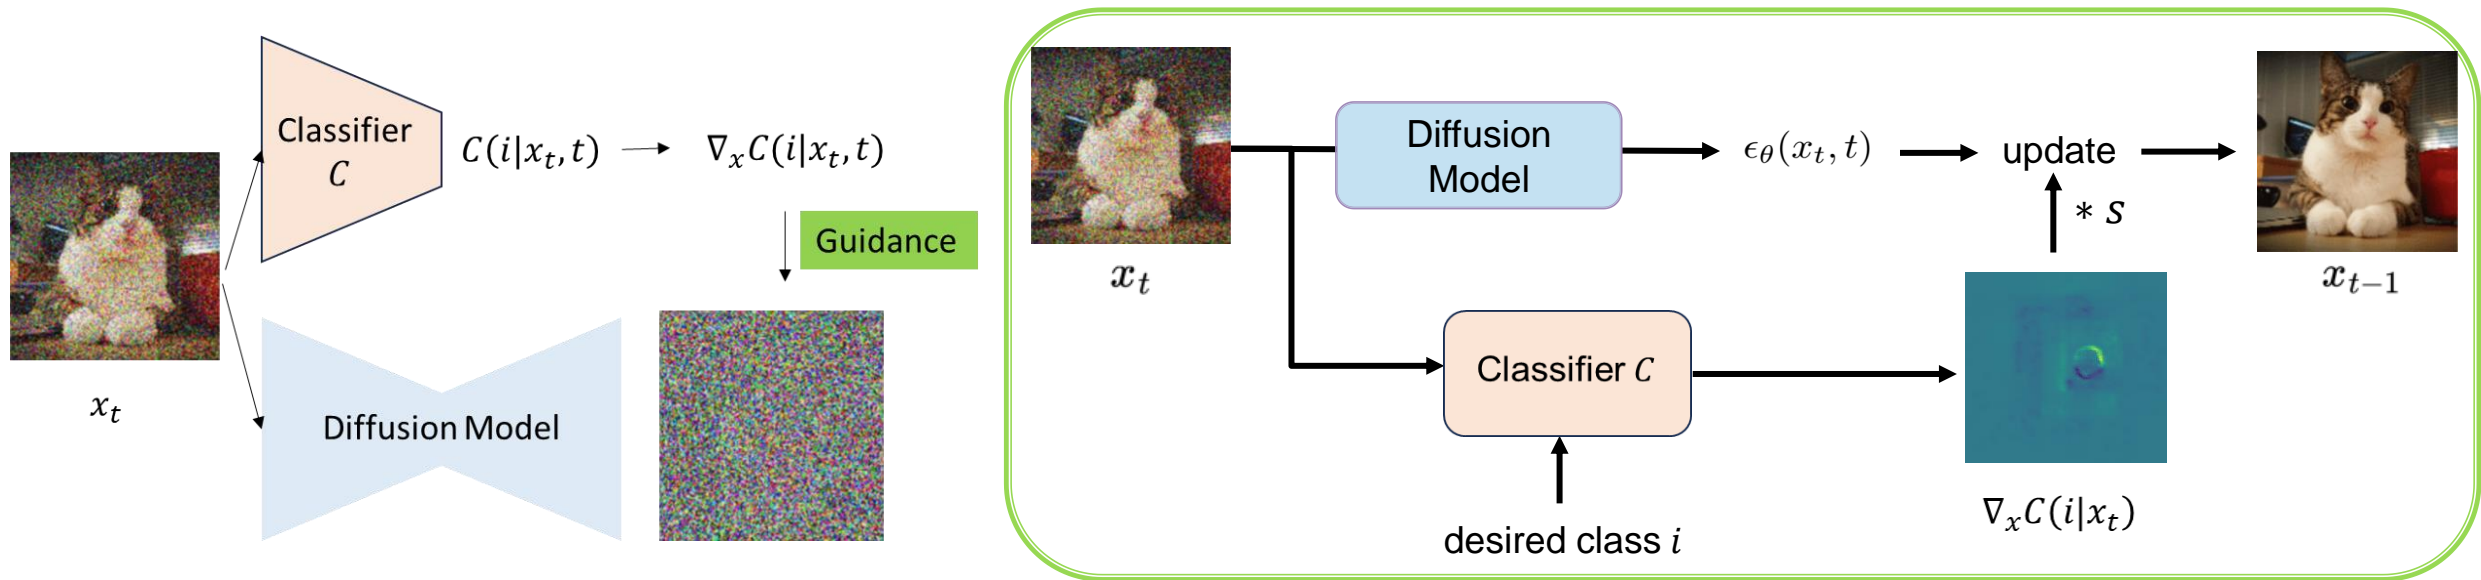
- Let's create a classifier $p_\phi(y|x_t)$



Take derivative to x_t and get
 $\nabla_{x_t} \log p(\hat{y}|x_t)$

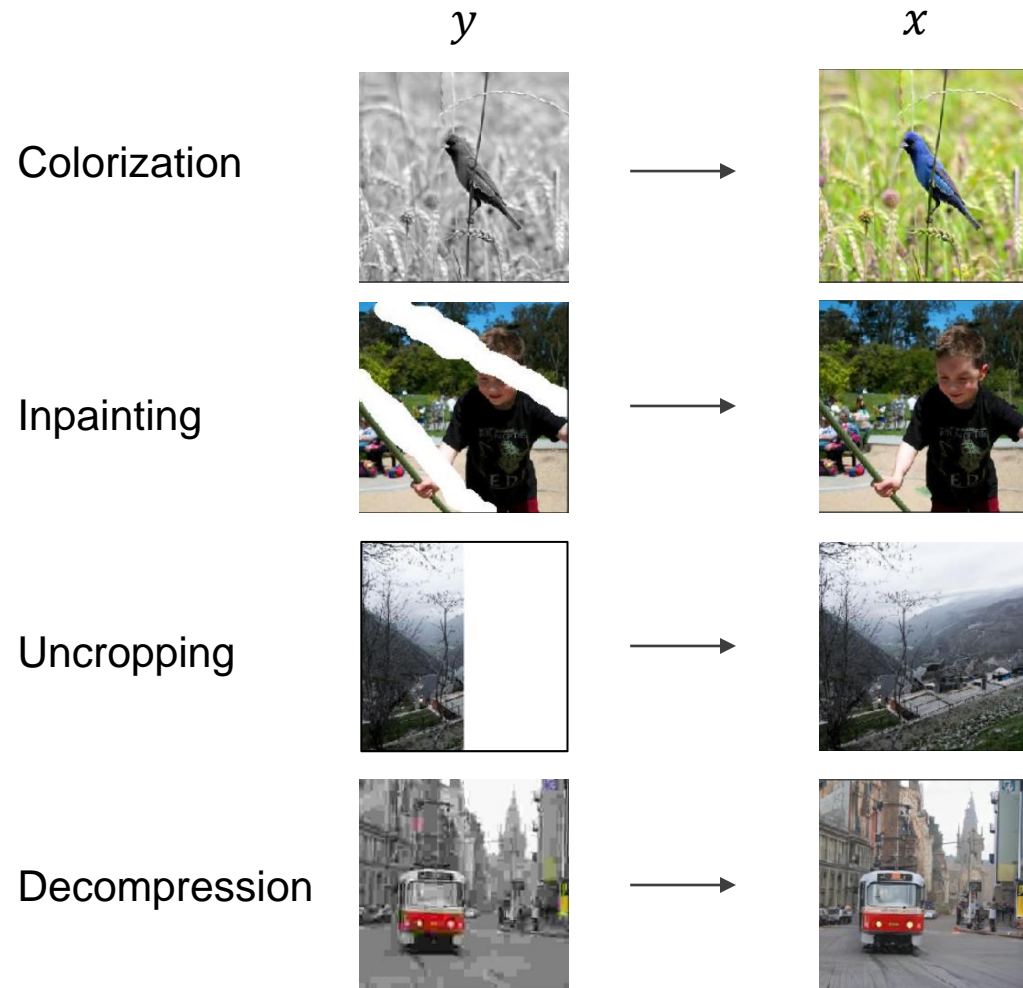
Classifier Guidance

We use the gradient to guide the generation process towards a desired class.



Gradient guidance is not restricted to classification models. Other models (e.g., regression, segmentation, ...) work just in the same way.

Classifier Free Conditioning – Pre-training



- For image generation of a fake image x , we can use a conditioning image y .
- This requires **paired** training.
- During training and sampling, we add information of the conditioning image x through **channel-wise concatenation**.

Comparison

- Conditioning
 - Pre-training: classifier free
 - Post-editing: classifier guidance
- Pros and Cons
 - Classifier guidance
 - Low training costs
 - Less detailed generation
 - Classifier free
 - Great details; Training with y , the more input samples the easier to train
 - Every time you want to add another signal/condition, needs to retrain the model

Evolutions of Text-to-Image Diffusions

- Pixel Space Diffusion Models

- GLIDE (ICML 2022)

- Transforms the input text c into a token sequence via a transformer
 - Replace the *class-embedding* with the *pooled text feature* and **concatenate** to the attention layers

- Imagen (NeurIPS 2022)

- Use a pre-trained LLM as its text encoder
 - “Cross-attention is **the most effective technique**”

- DALL.E 2, 3(arXiv 2022)

- Bridges the gap between CLIP text and the image latent space $p(z_i|z_t)$

- Latent Space Diffusion Models

- LDM (CVPR 2022)

- Enhance the underlying Unet with cross-attention

- Stable Diffusion (SD) v1, v2(CVPR 2022)

- SD XL (ICLR 2024)

Glide: Text Conditioning

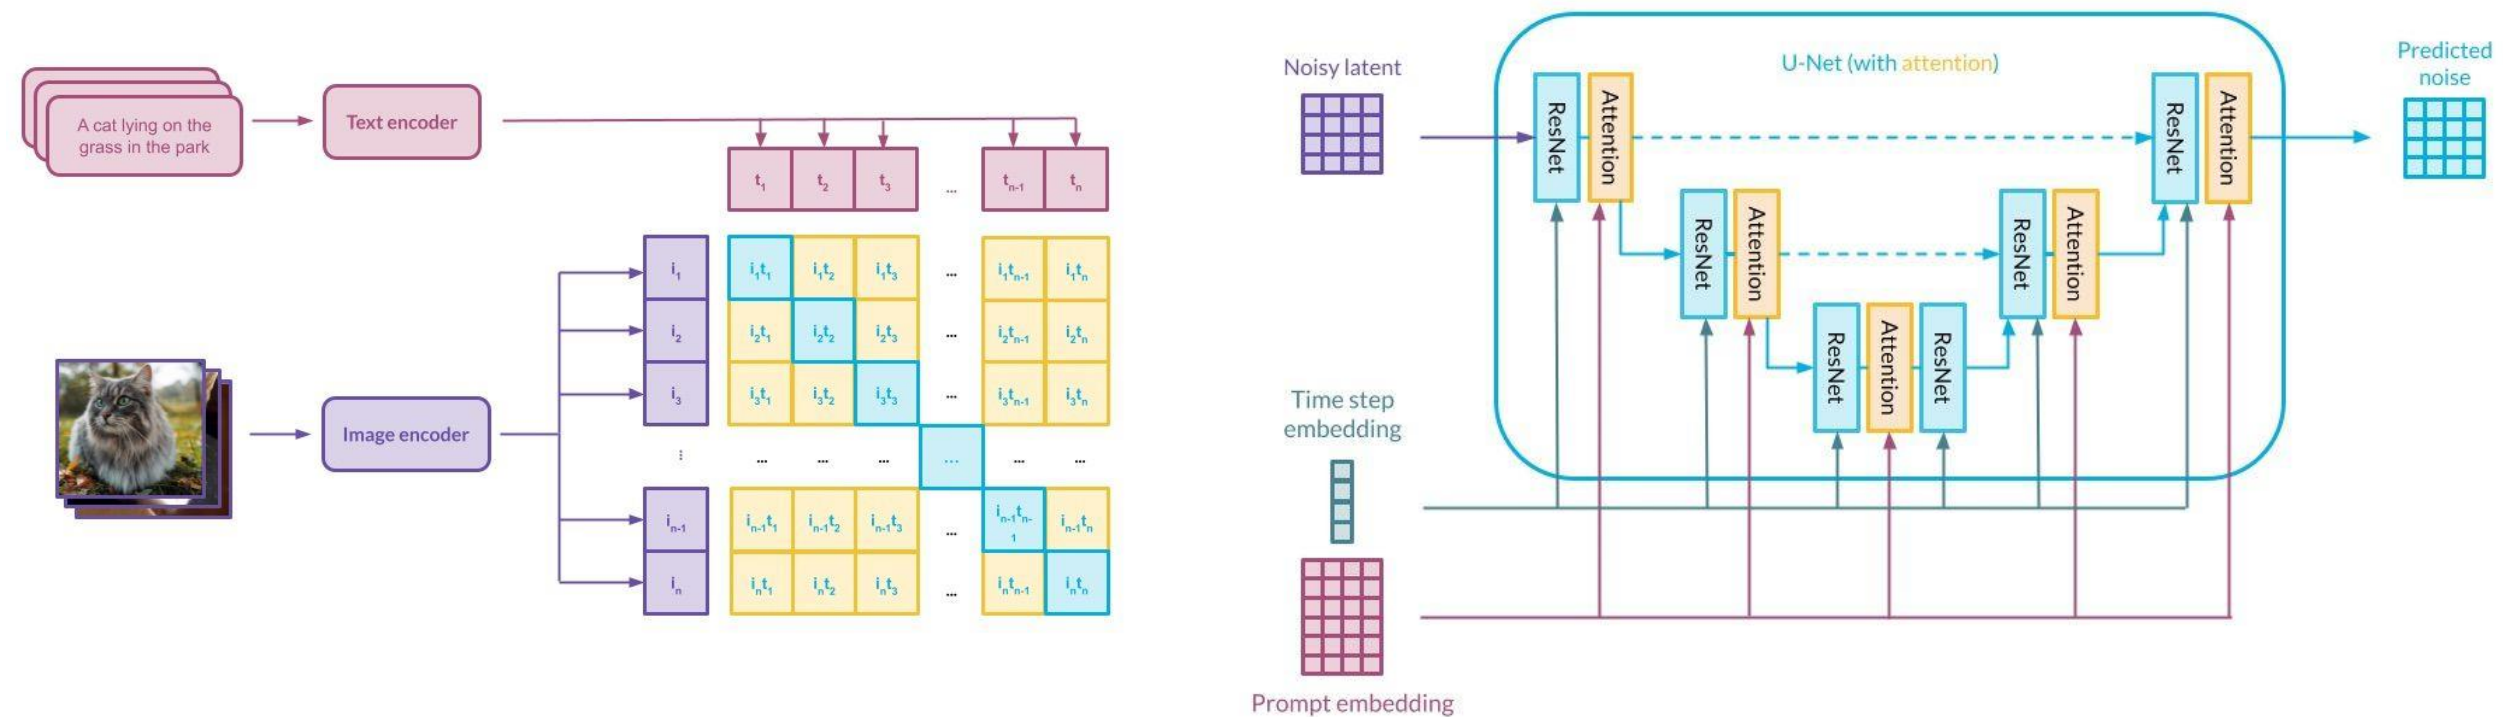
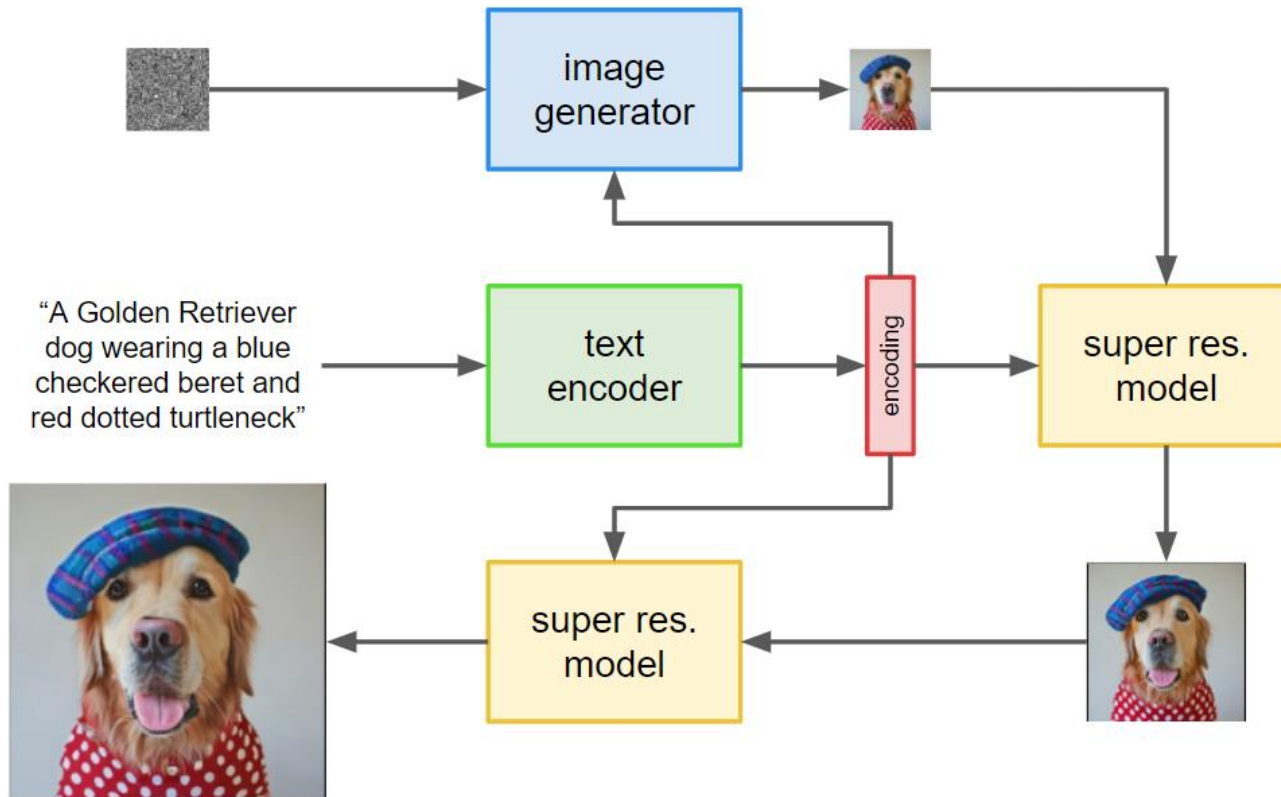


Imagen: Text Conditioning



Teddy bears swimming at the Olympics 400m Butterfly event.



A cute corgi lives in a house made out of sushi.

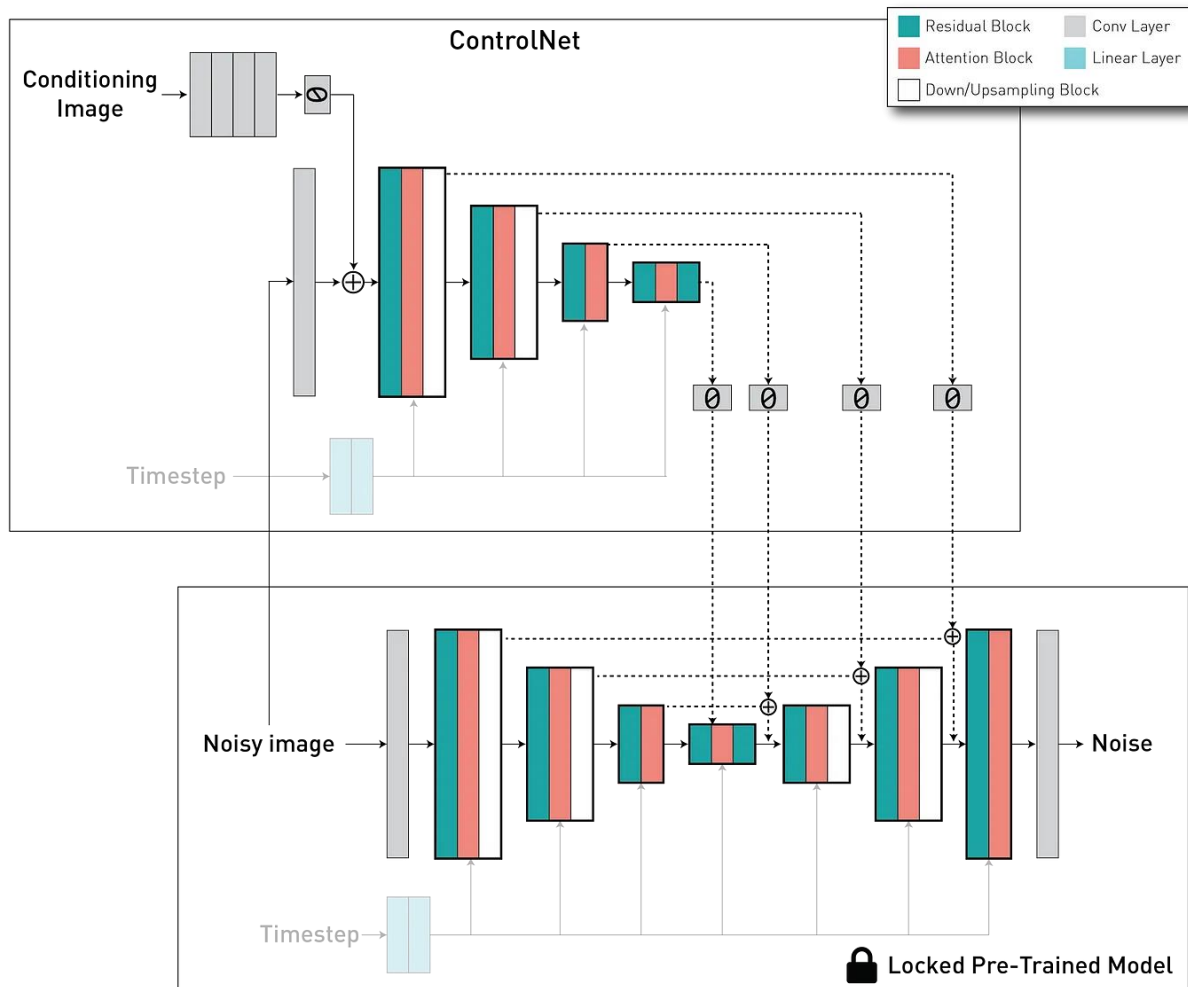


A brain riding a rocketship heading towards the moon.

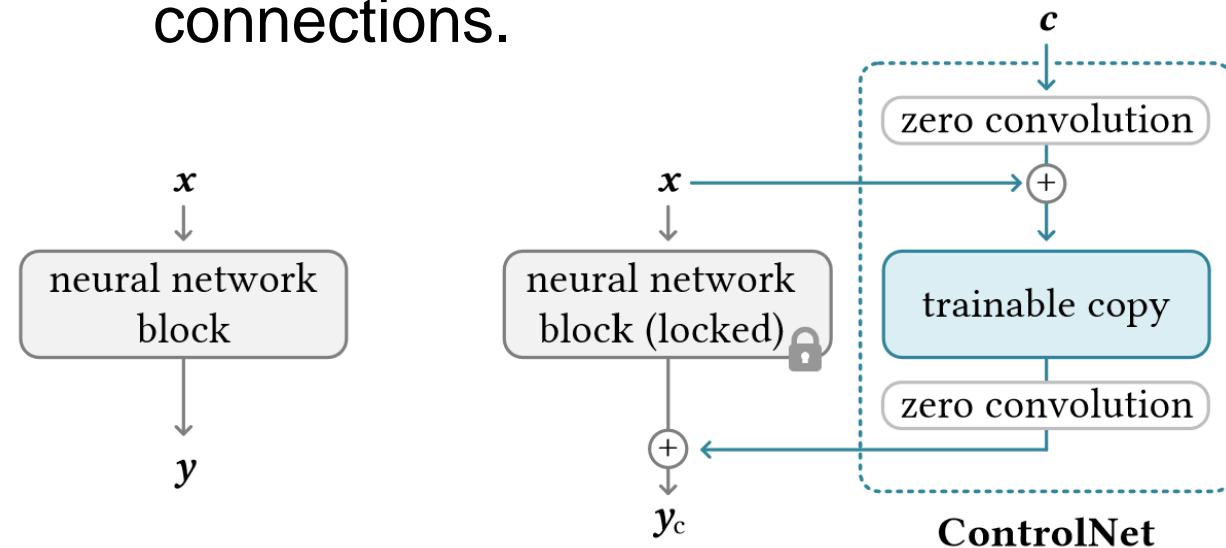


A dragon fruit wearing karate belt in the snow.

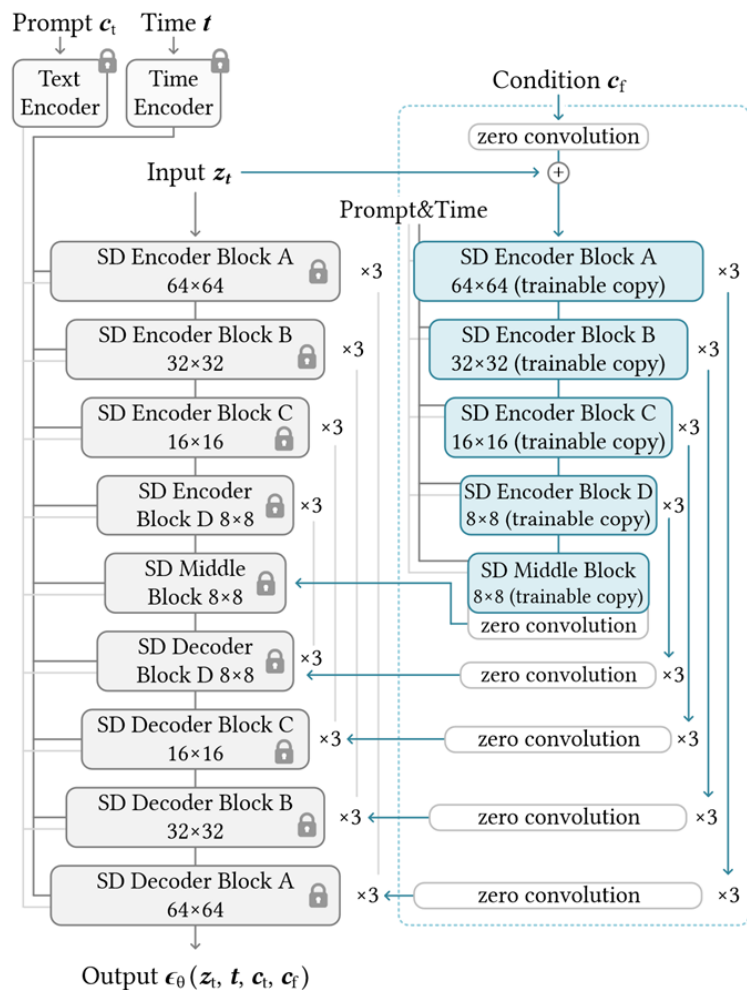
ControlNet



- We pretrain a diffusion model with text prompts.
- We freeze this model.
- We fine-tune a copy conditioned on c .
- We pass information through skip connections.



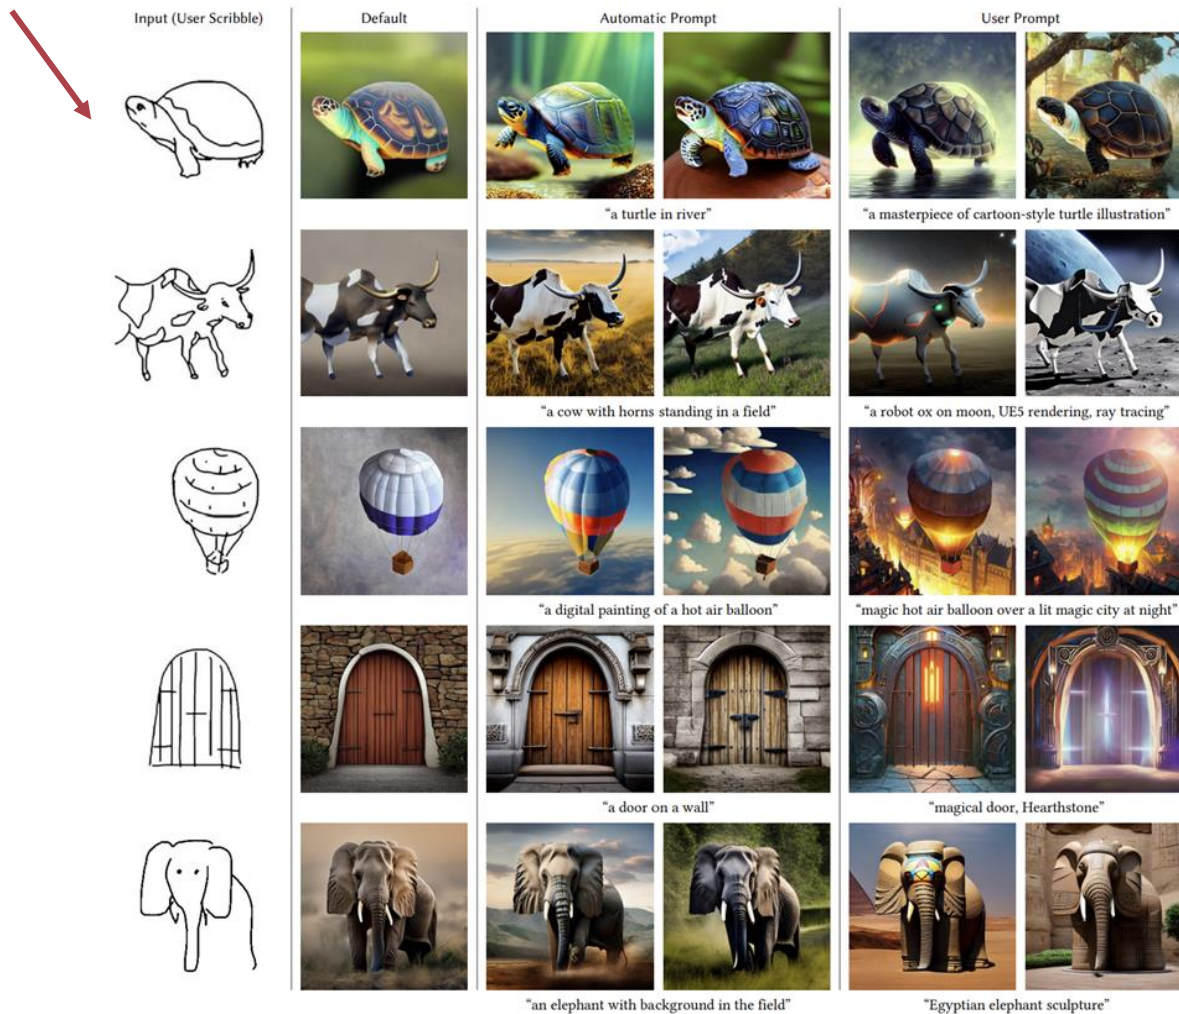
ControlNet



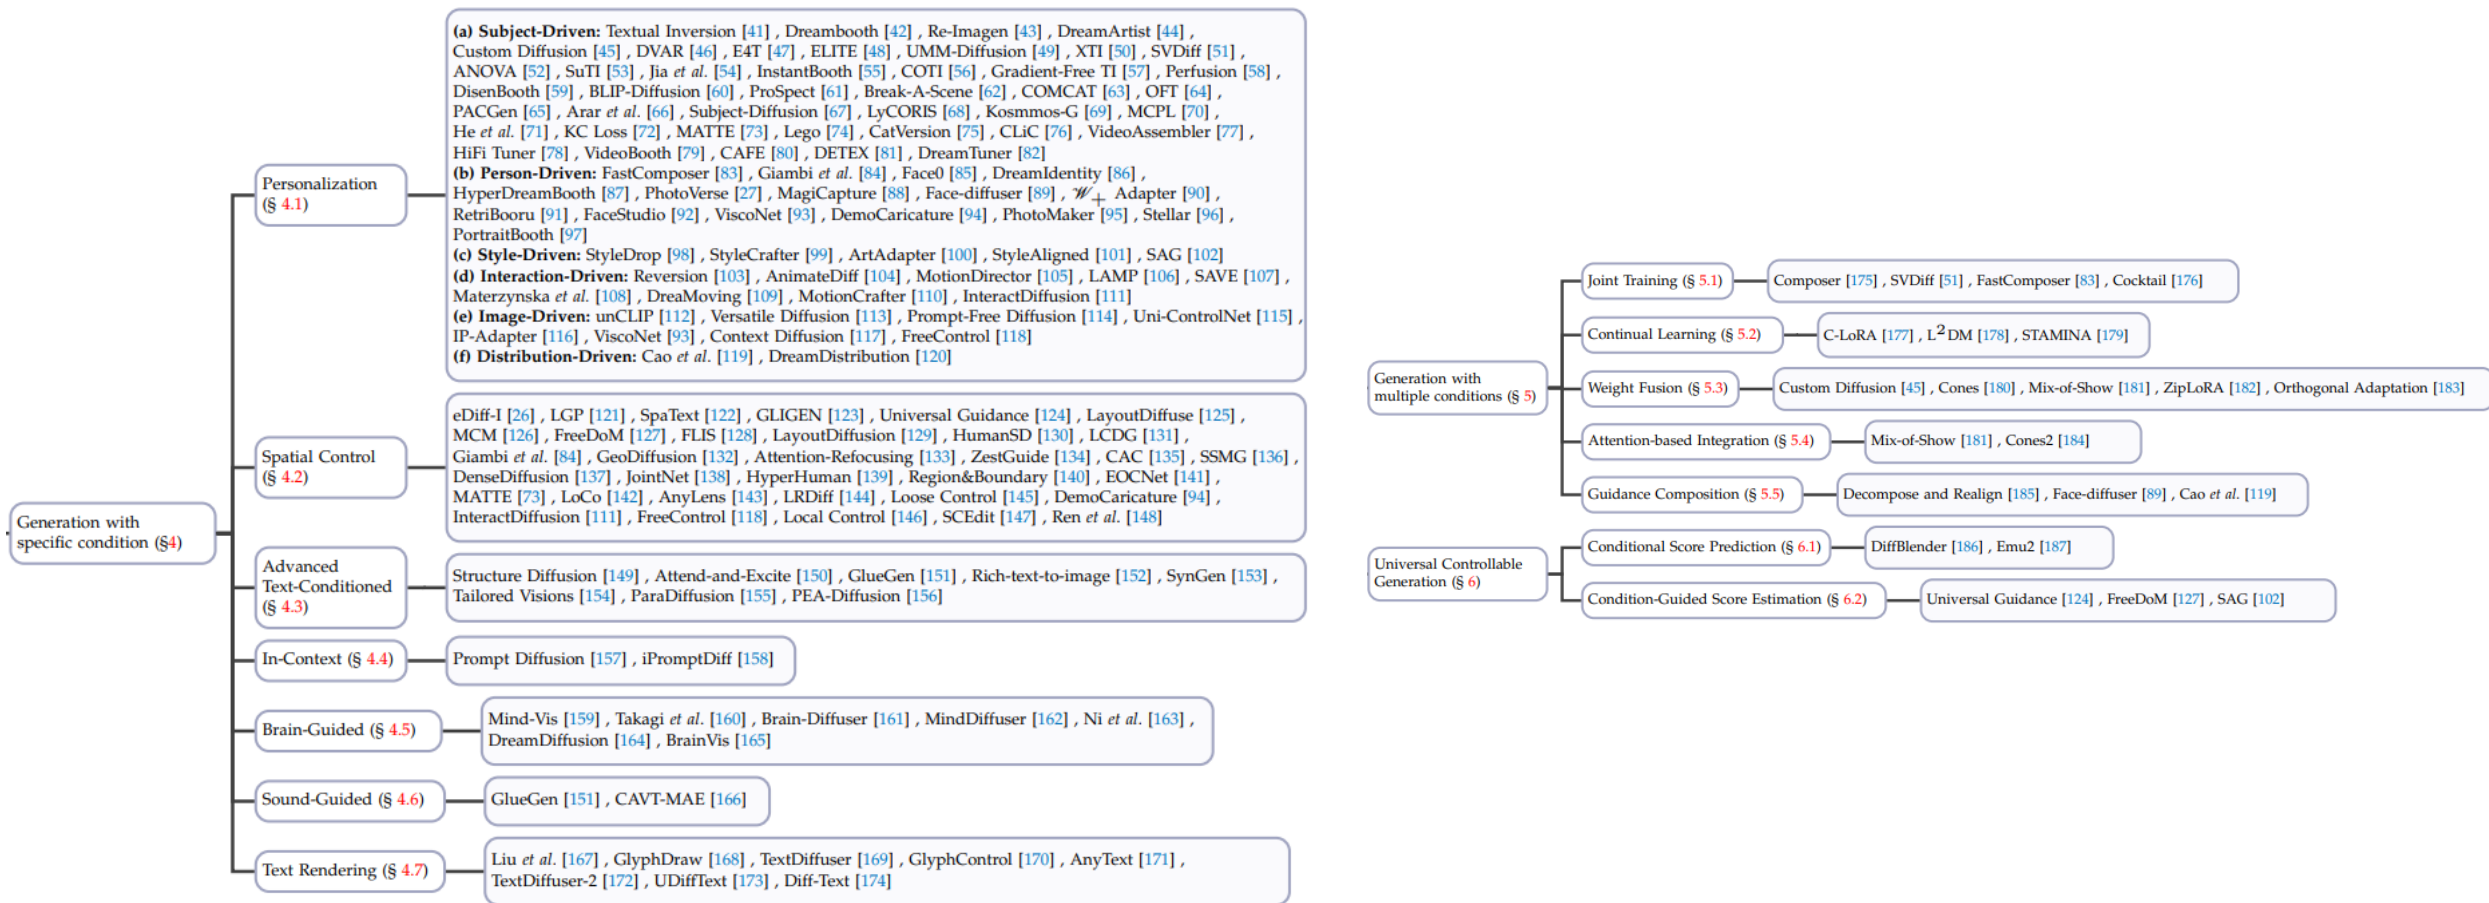
(a) Stable Diffusion

(b) ControlNet

conditioning image



Taxonomy of Controllable Generations



Why do you think diffusion models exhibit such performance?



How to participate?



- 1 Go to wooclap.com
- 2 Enter the event code in the top banner

Event code
FROCKY

 Enable answers by SMS

 [Copy participation link](#)

What types of applications in Medical Imaging are you most interested in applying generative models to?



How to participate?



- 1 Go to wooclap.com
- 2 Enter the event code in the top banner

Event code
FROCKY

 Enable answers by SMS

 [Copy participation link](#)

Q&A

Generated by DALL.E 3



Tutorial Schedule

First Half

- ❑ Part 1: Introduction (30 mins)
 - ❖ What? Why? How?
 - ❖ Denoising Diffusion Models
 - ❖ Understanding and Intuition
- ❑ Part 2: Advanced Topics (30 mins)
 - ❖ Sampling Strategies
 - ❖ Conditioning

Second Half

- ❑ Q&A
 - ❑ **Medical Imaging Applications (30 mins)**
 - ❖ Synthesis
 - ❖ Segmentation
 - ❖ Anomaly Detection
 - ❖ Reconstruction
 - ❖ Registration
- We are here!

Part 2 – Medical Image Applications

Image Reconstruction

Image Registration

Image Segmentation

Anomaly Detection


Image Synthesis

Image synthesis

Examples from the community

The simple setup of the problem

Real




Synthetic

Figure by Song et al ICLR 2022.
Copyright rests with the authors.

PAPERS

- Pinaya et al (2022) Brain Imaging Generation with Latent Diffusion Models. MICCAI 2022 workshop
- Khader et al. (2022) Medical Diffusion -- Denoising Diffusion Probabilistic Models for 3D Medical Image Generation. Scientific Reports, 2023
- Chambon, Pierre, et al. (2022) RoentGen: vision-language foundation model for chest x-ray generation.
- Ye, Jiarong, et al. (2023) Synthetic Augmentation with Large-scale Unconditional Pre-training. MICCAI 2023
- Fernandez, V.et al. (2022). Can segmentation models be trained with fully synthetically generated data? MICCAI 2022 Workshop
- Packhäuser et al. (2022) Generation of Anonymous Chest Radiographs Using Latent Diffusion Models for Training Thoracic Abnormality Classification Systems. ISBI 2023
- Fernandez, V. et al (2023). Privacy Distillation: Reducing Re-identification Risk of Multimodal Diffusion Models. MICCAI 2023 Workshop
- Sagers, Luke W., et al. (2023) Augmenting medical image classifiers with synthetic data from latent diffusion models.
- Frisch, Yannik, et al. (2023) Synthesising Rare Cataract Surgery Samples with Guided Diffusion Models. MICCAI 2023.
- Kim et al. (2022) Diffusion Deformable Model for 4D Temporal Medical Image Generation. MICCAI 2022
- Ali et al. (2022) Spot the fake lungs: Generating Synthetic Medical Images using Neural Diffusion Models. AICS 2022
- Rouzrokh et al. (2022) Multitask Brain Tumor Inpainting with Diffusion Models: A Methodological Report.
- Chambon et al (2022) Adapting Pretrained Vision-Language Foundational Models to Medical Imaging Domains. NeurIPS2022
- Lyu et al. (2022) Conversion Between CT and MRI Images Using Diffusion and Score-Matching Models.
- Ozbey et al. (2023) Unsupervised Medical Image Translation with Adversarial Diffusion Models. IEEE Transactions on Medical Imaging 2023
- Ktena, Ira, et al. (2024) Generative models improve fairness of medical classifiers under distribution shifts. Nature Medicine, 2024
- Dominik J. E. Waibel, et al. (2023) Diffusion Model Predicts 3D Shapes from 2D Microscopy Images. ISBI 2023.
- P.Huy, et al. (2023) Denoising Diffusion Medical Models. ISBI 2023.

Why? Medical Image Data is Scarce




Use of Synthetic Data



- Full “private” training
- Data augmentation
- Test-time augmentation
- Testing edge cases

Evaluation Criteria

- Realism
- Diversity
- Privacy




Generating high-resolution 3D brain data



Synthetic

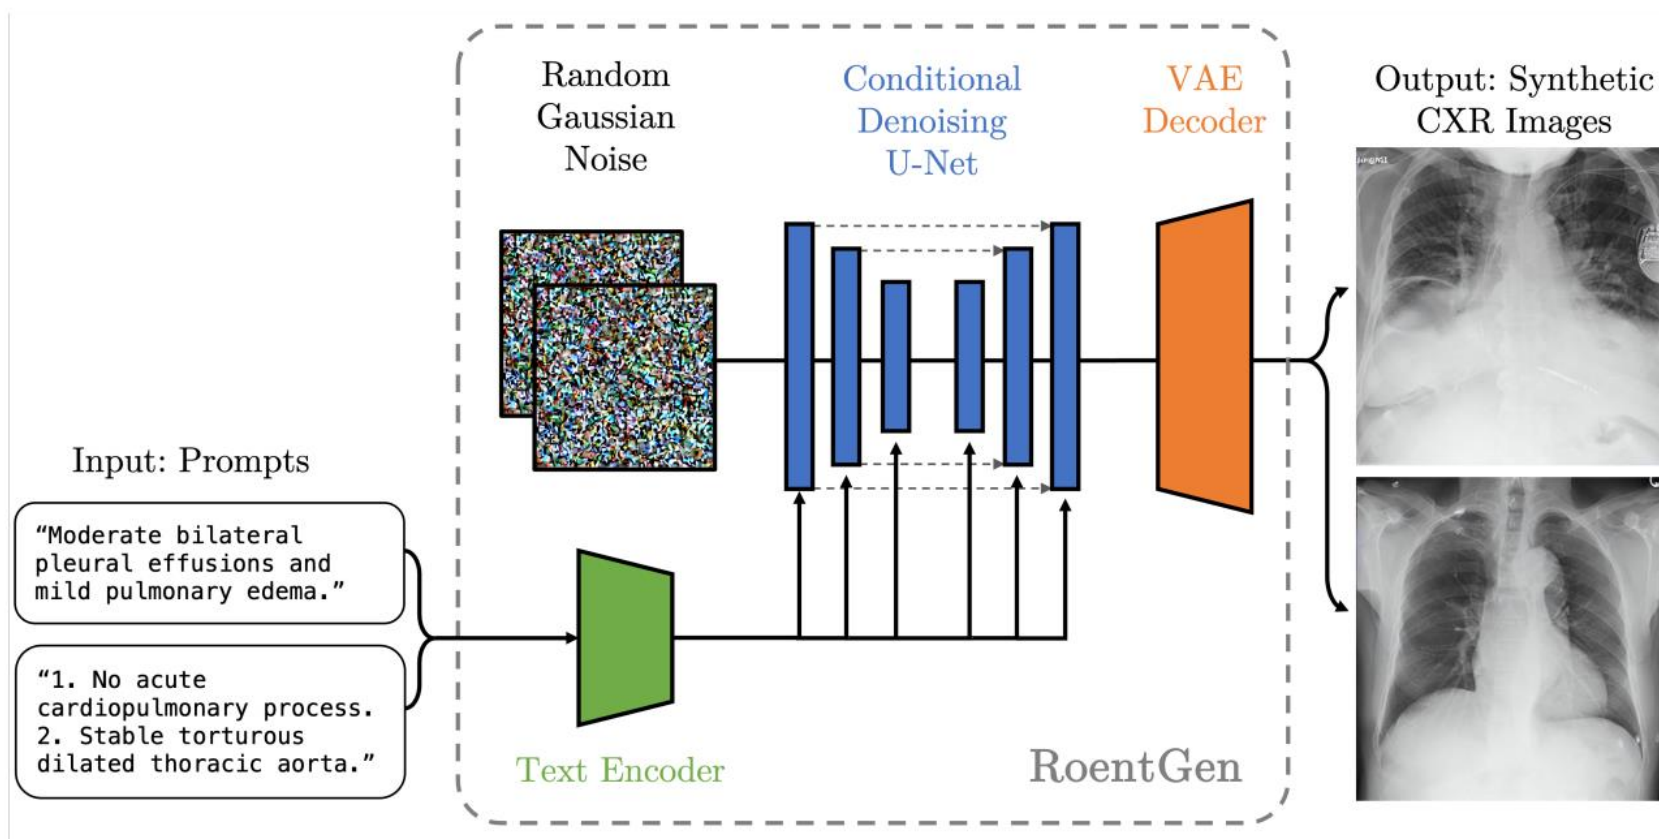
- Latent Diffusion Models trained on data from UK Biobank (N = 31,740)
 - T1 MRI brain images with 1 mm³ voxel size (160 × 224 × 160 voxels)
- Conditioned on covariates, such as:
 - Age
 - Gender
 - Ventricular and Brain volumes



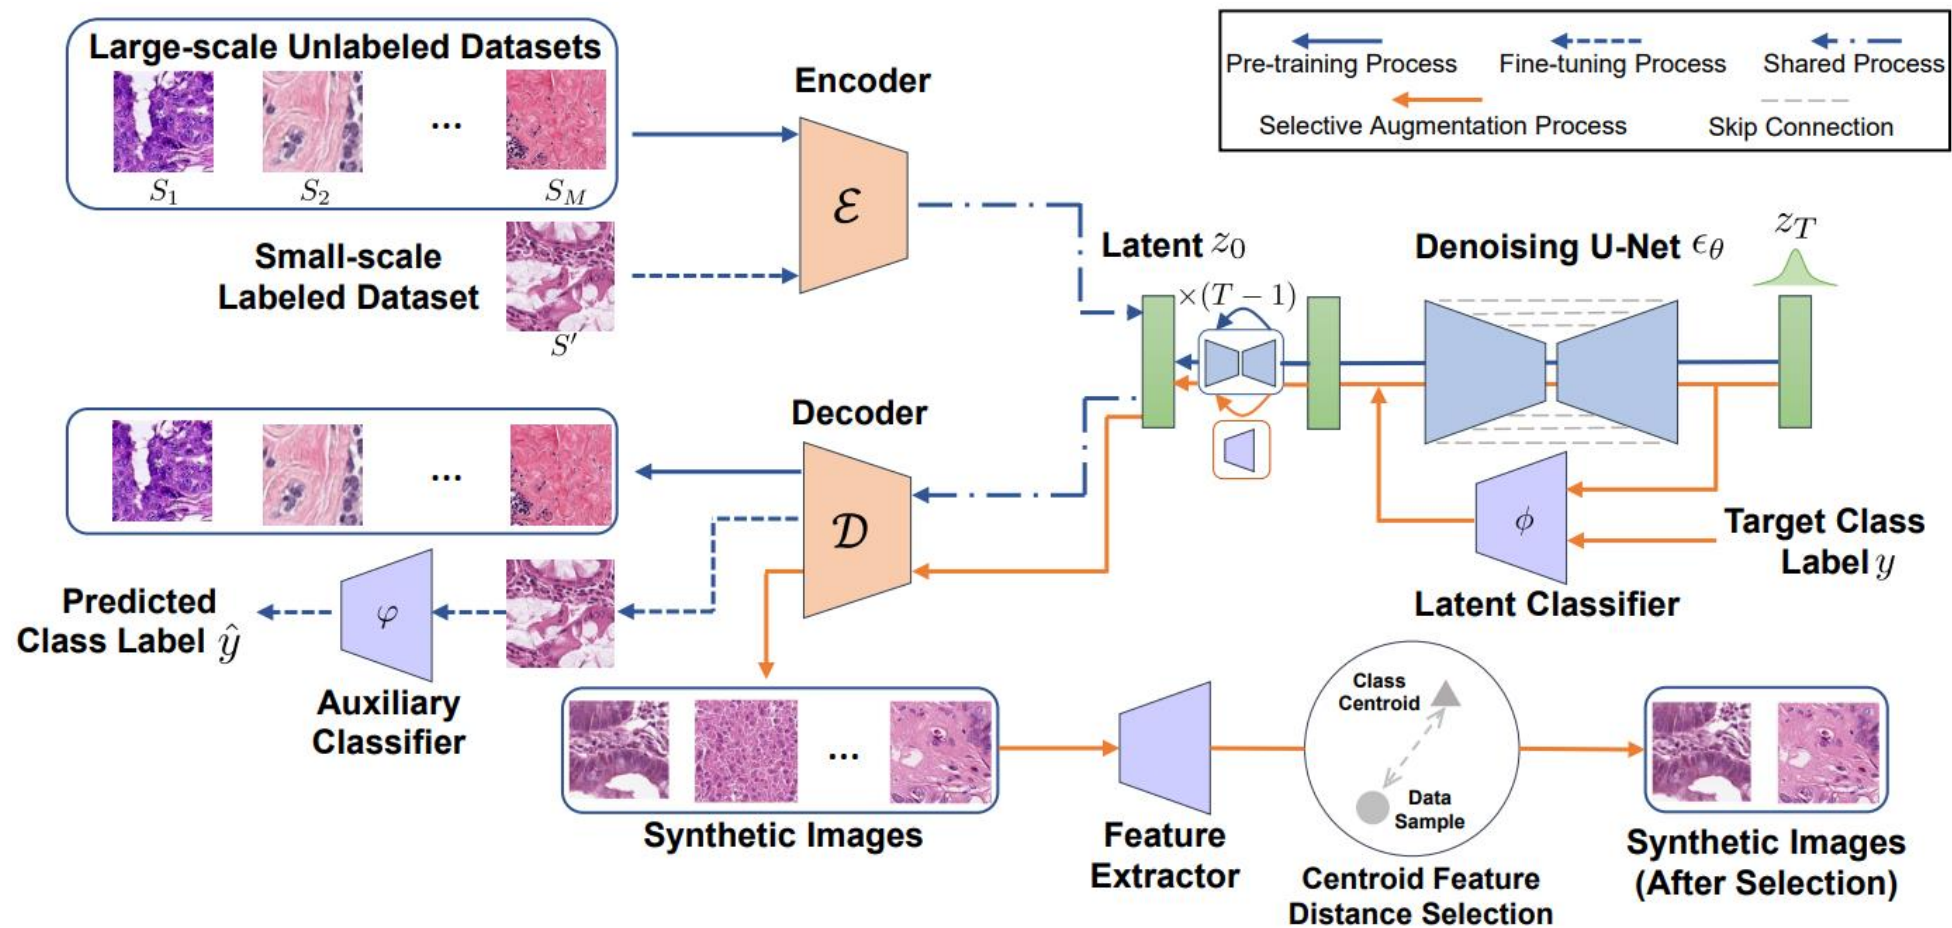
Pinaya et al (2022). Brain Imaging Generation with Latent Diffusion Models. MICCAI 2022 workshop

Khader et al. (2023) Medical Diffusion - Denoising Diffusion Probabilistic Models for 3D Medical Image Generation. Scientific Reports, 2023

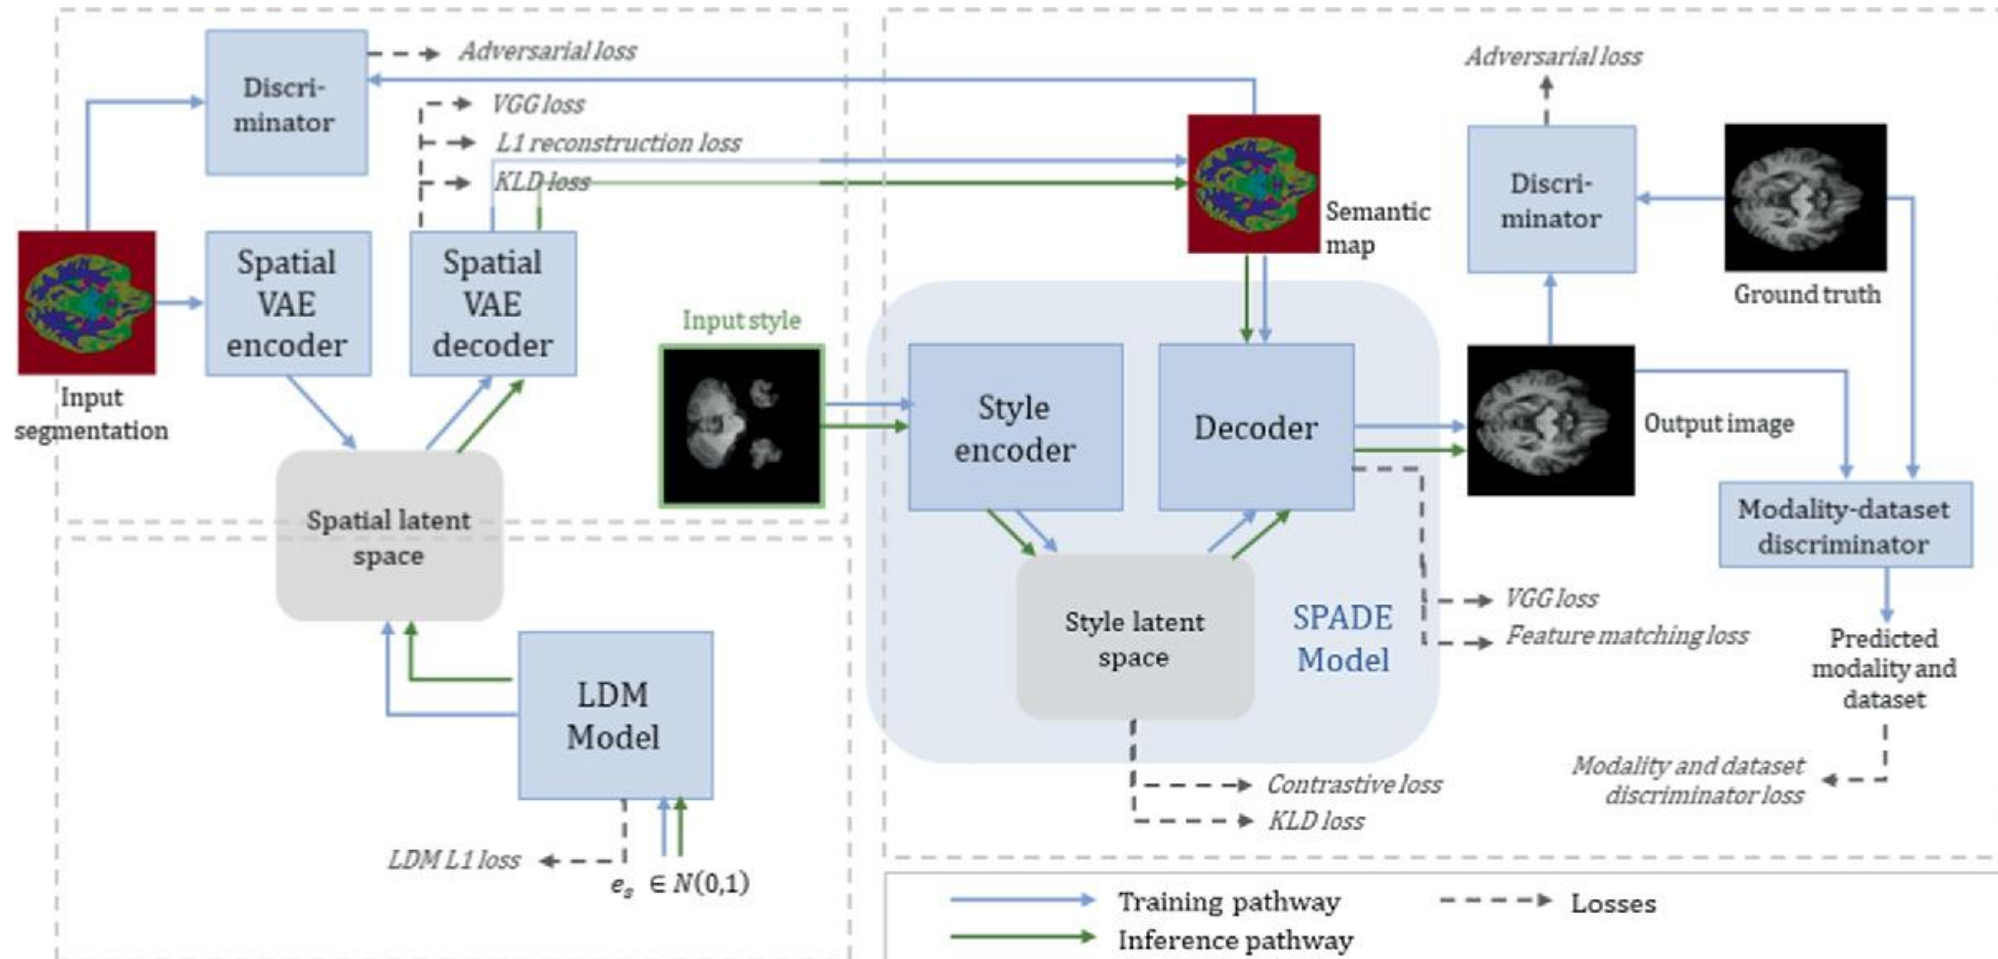
Fine-tuning Stable Diffusion



Unlabelled Pre-training



Generating Segmentation Masks



Generation of Anonymous Chest Radiographs

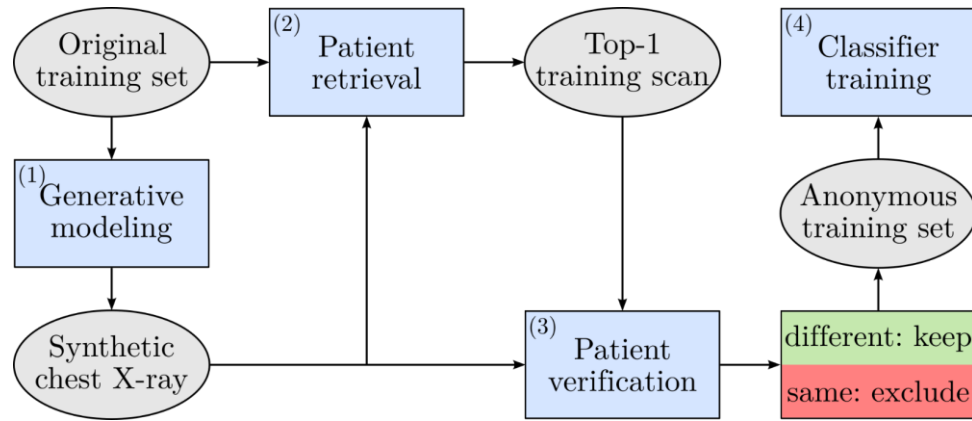


Fig. 1: Proposed privacy-enhancing image sampling strategy. Image taken from [1].




Fig. 2: Comparison of the classification performance of CheXNet.


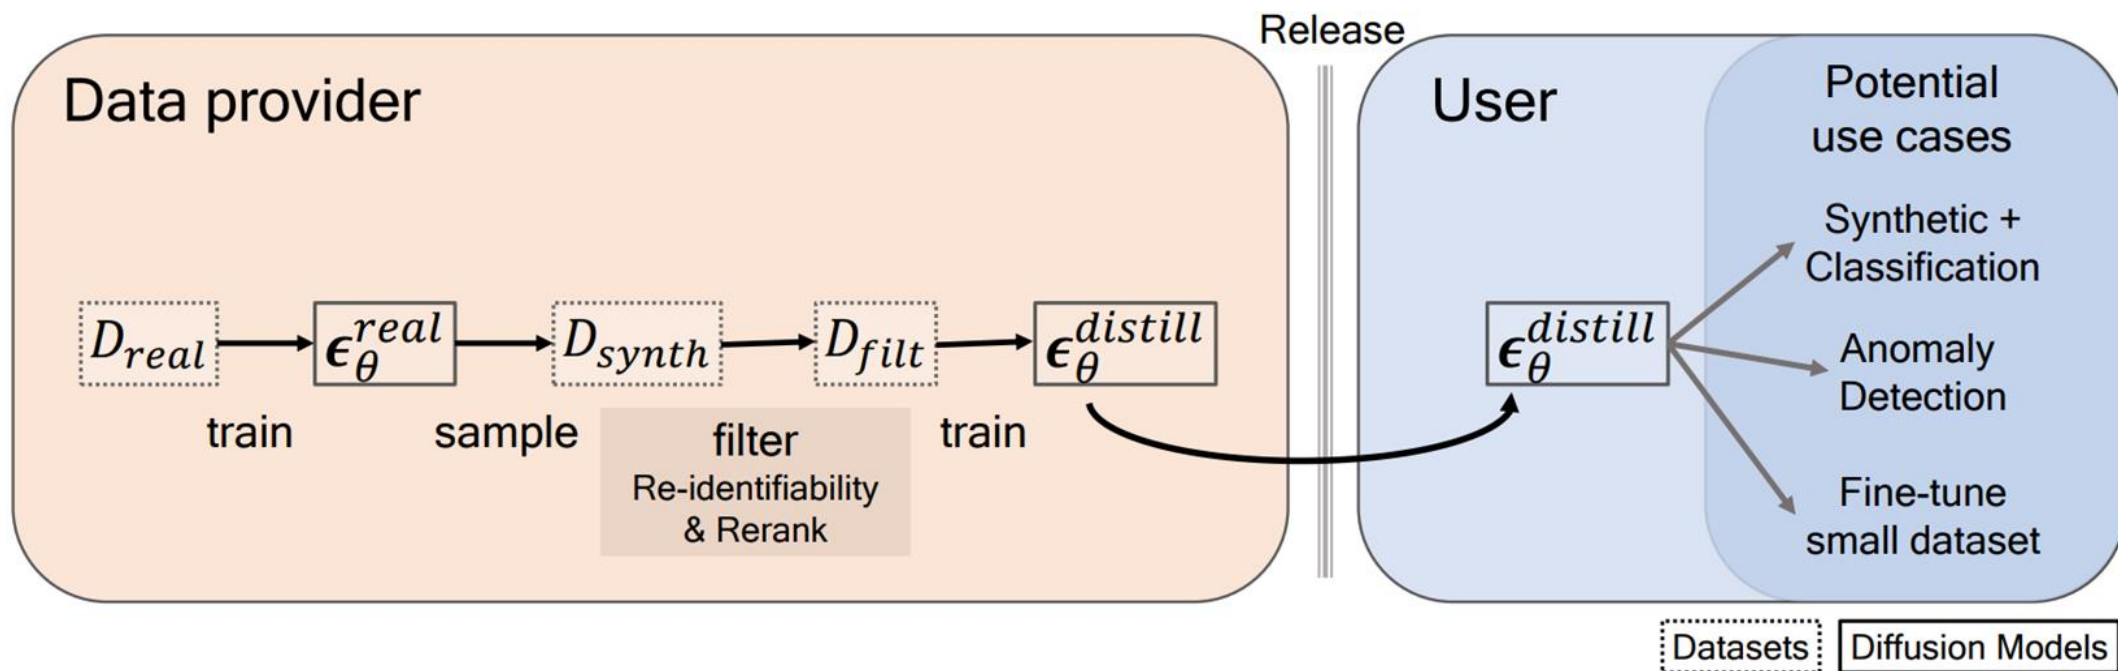


Fig. 3: Randomly selected images generated by the trained LDM. Images taken from [1].

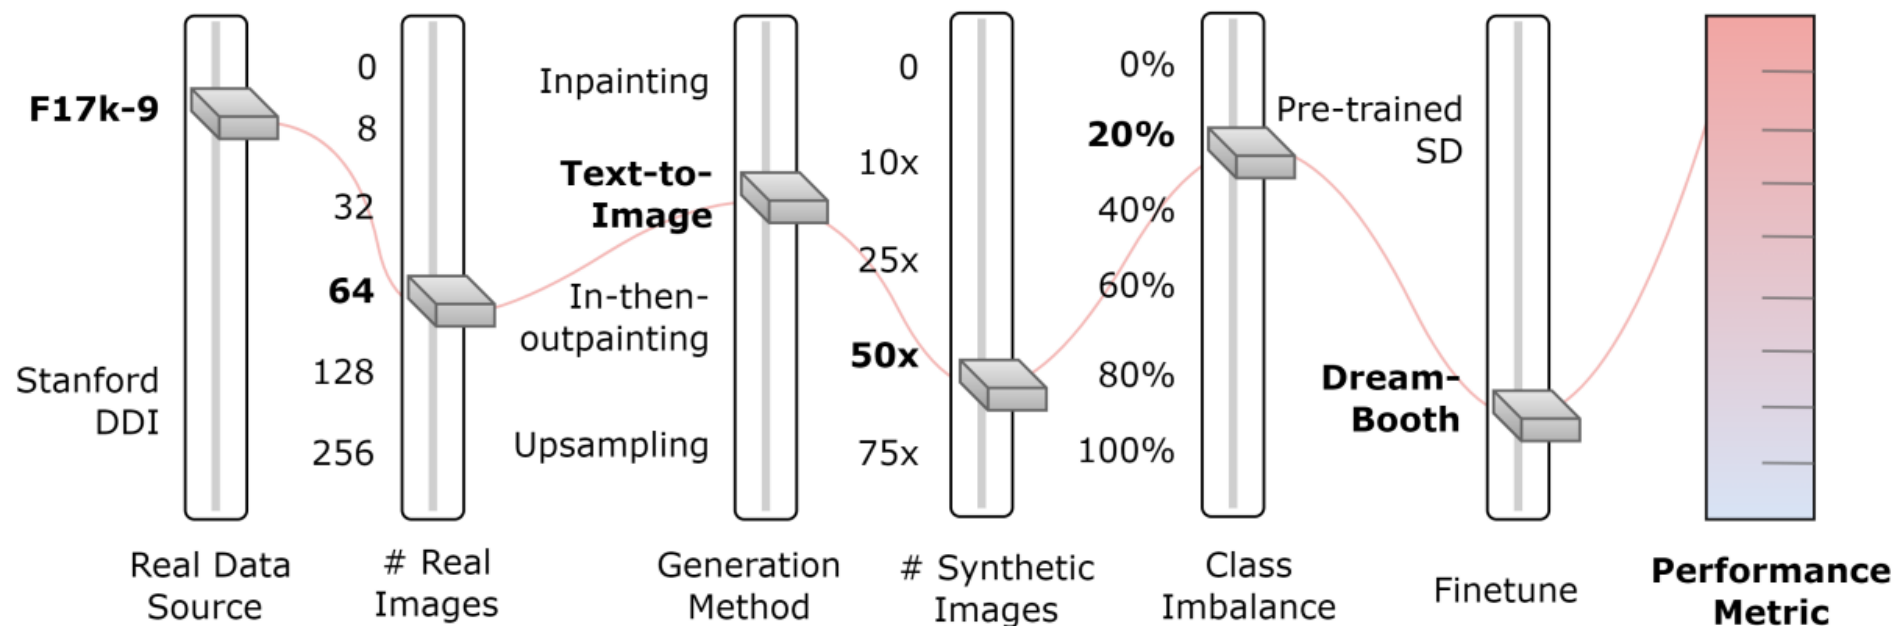
Slides courtesy of Kai Packhäuser

Privacy Distillation

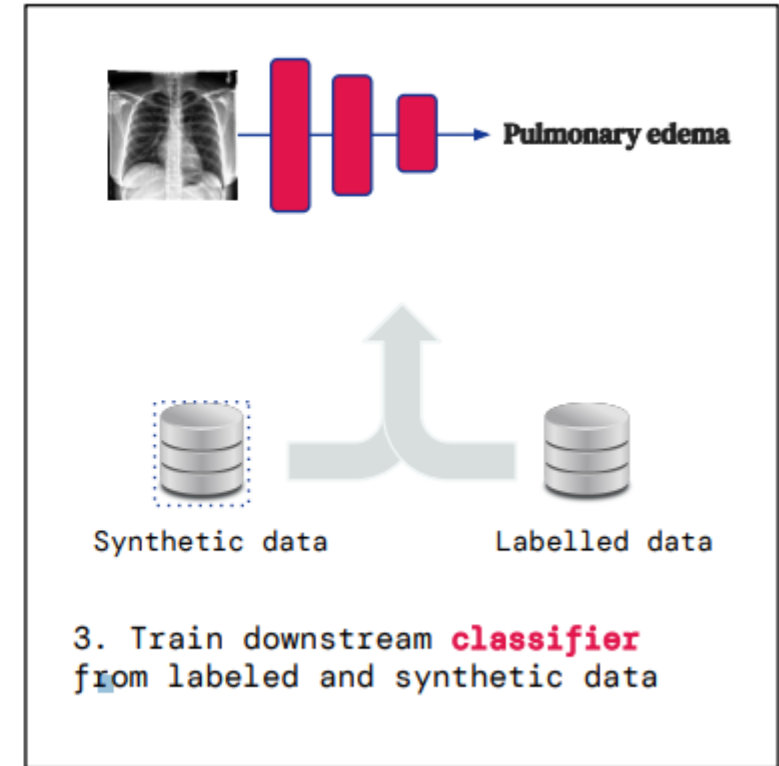
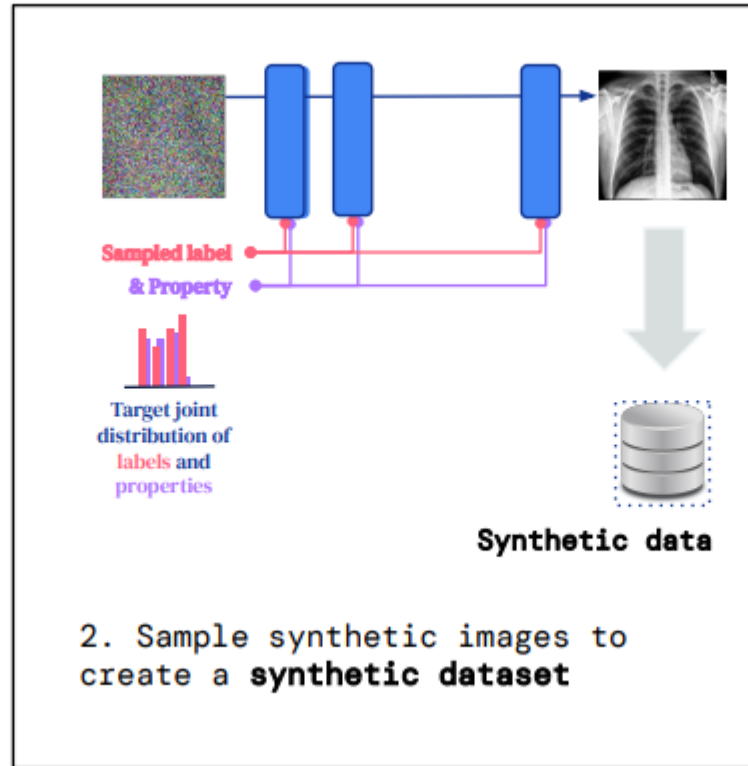
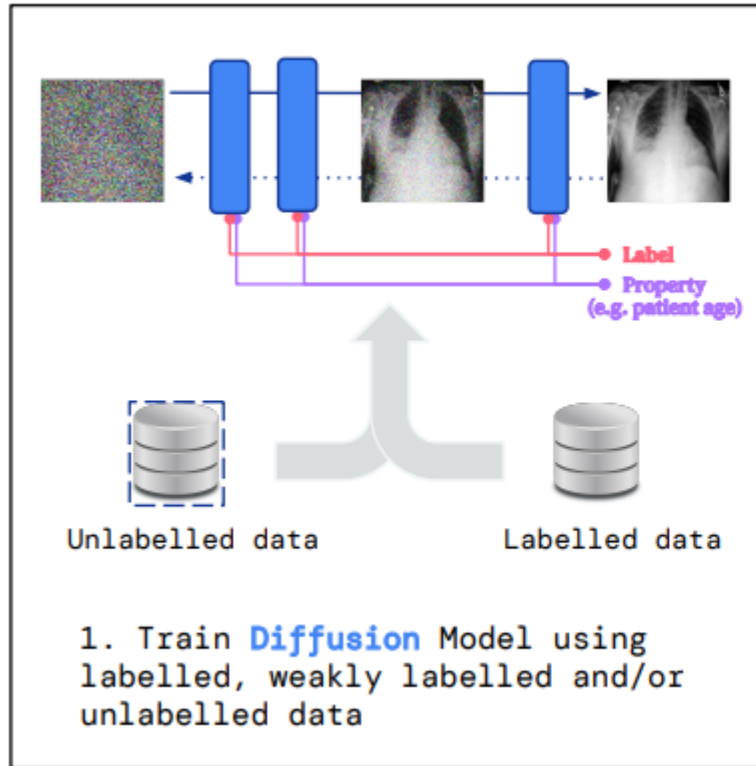


Synthetic Image Augmentation

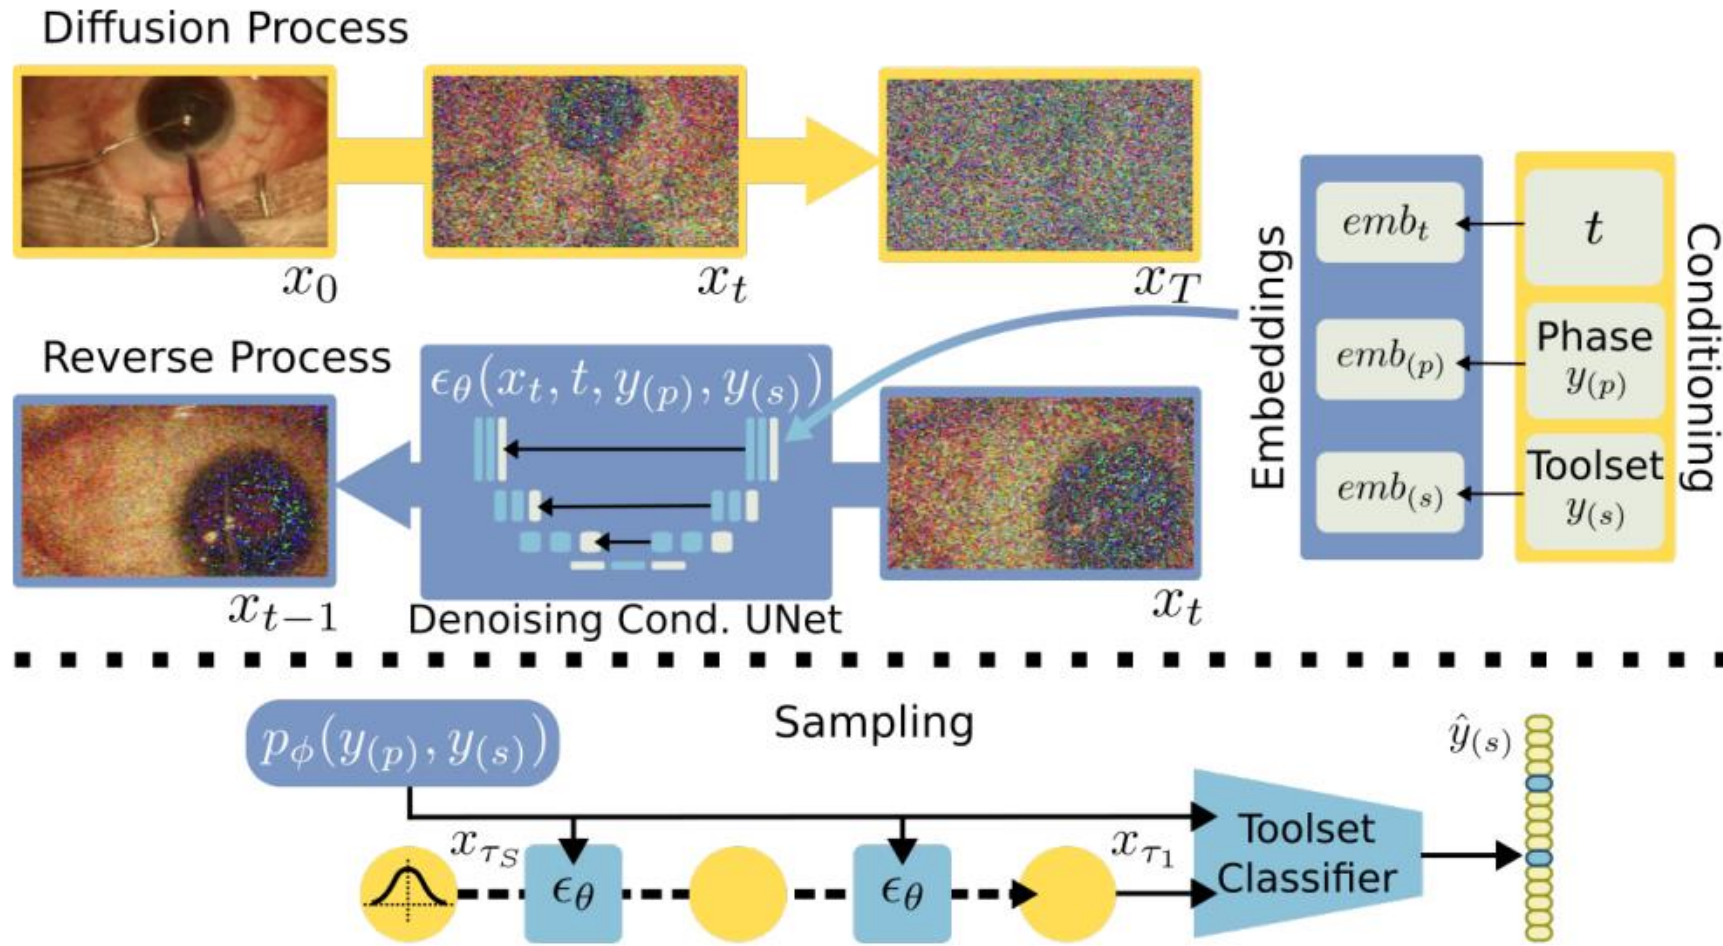
Synthetic-to-real ratio of 10:1



Synthetic Data for Distribution Shifts


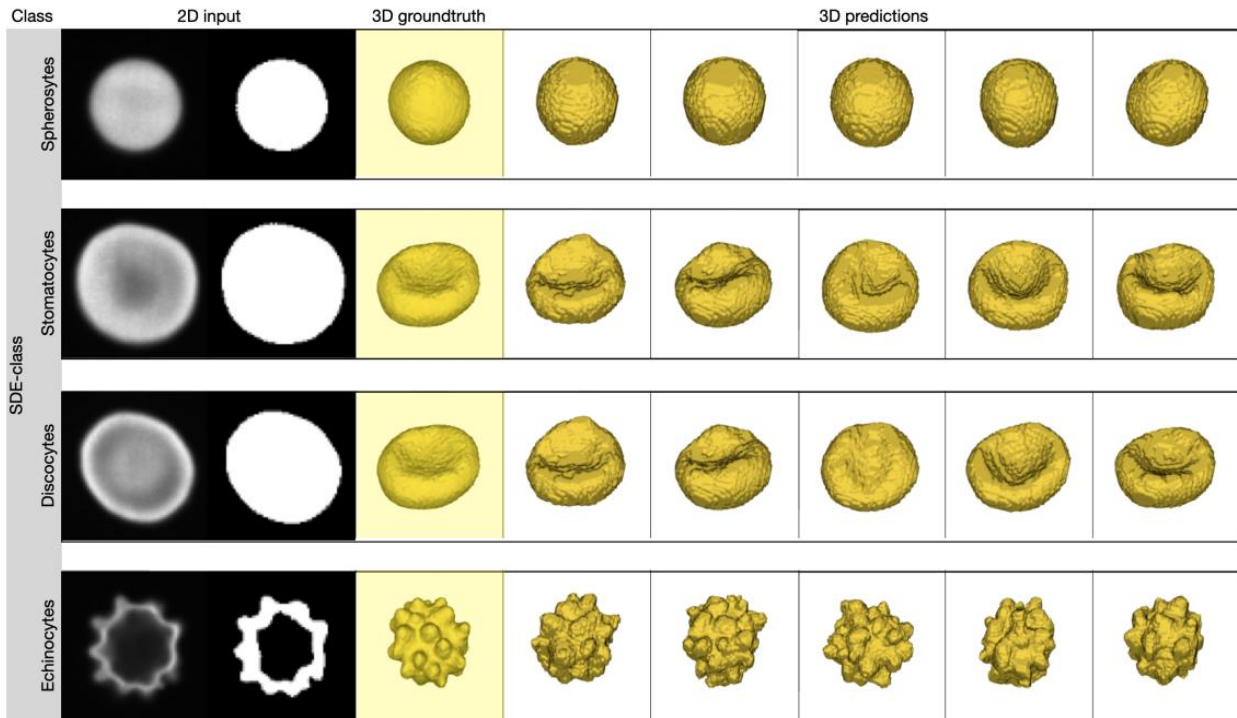


Synthesising Rare Samples



Diffusion based Shape PRediction, **DISPR**


3D shell shapes generation conditioned on 2D microscopy images



the 2D image is concatenated in each timestep of diffusion process

Denoising Diffusion Medical Models

DDMM

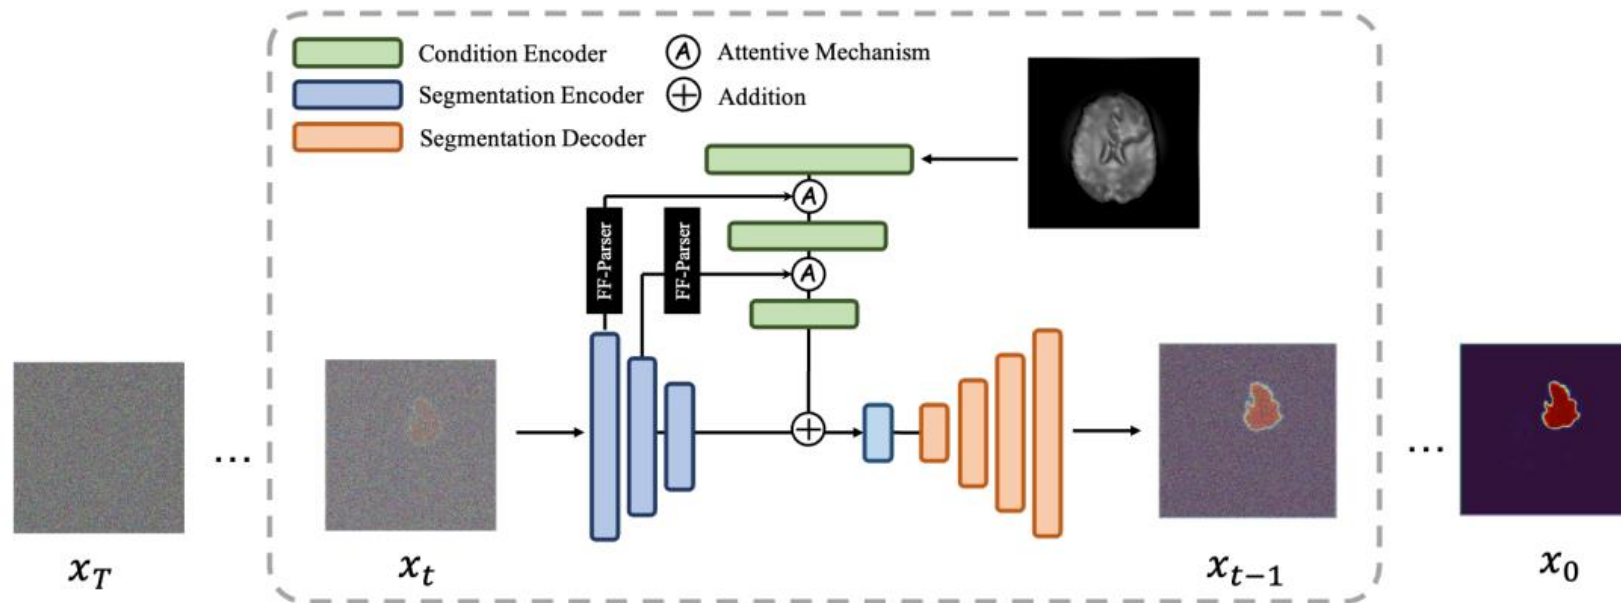


The DDMM model is built from one or more branches of DDPM (**radiographs** and **segmentation branches**) that share the same noise scheduler and latent code, which enforce **semantic consistency**

Image segmentation

Examples from the community

Setup

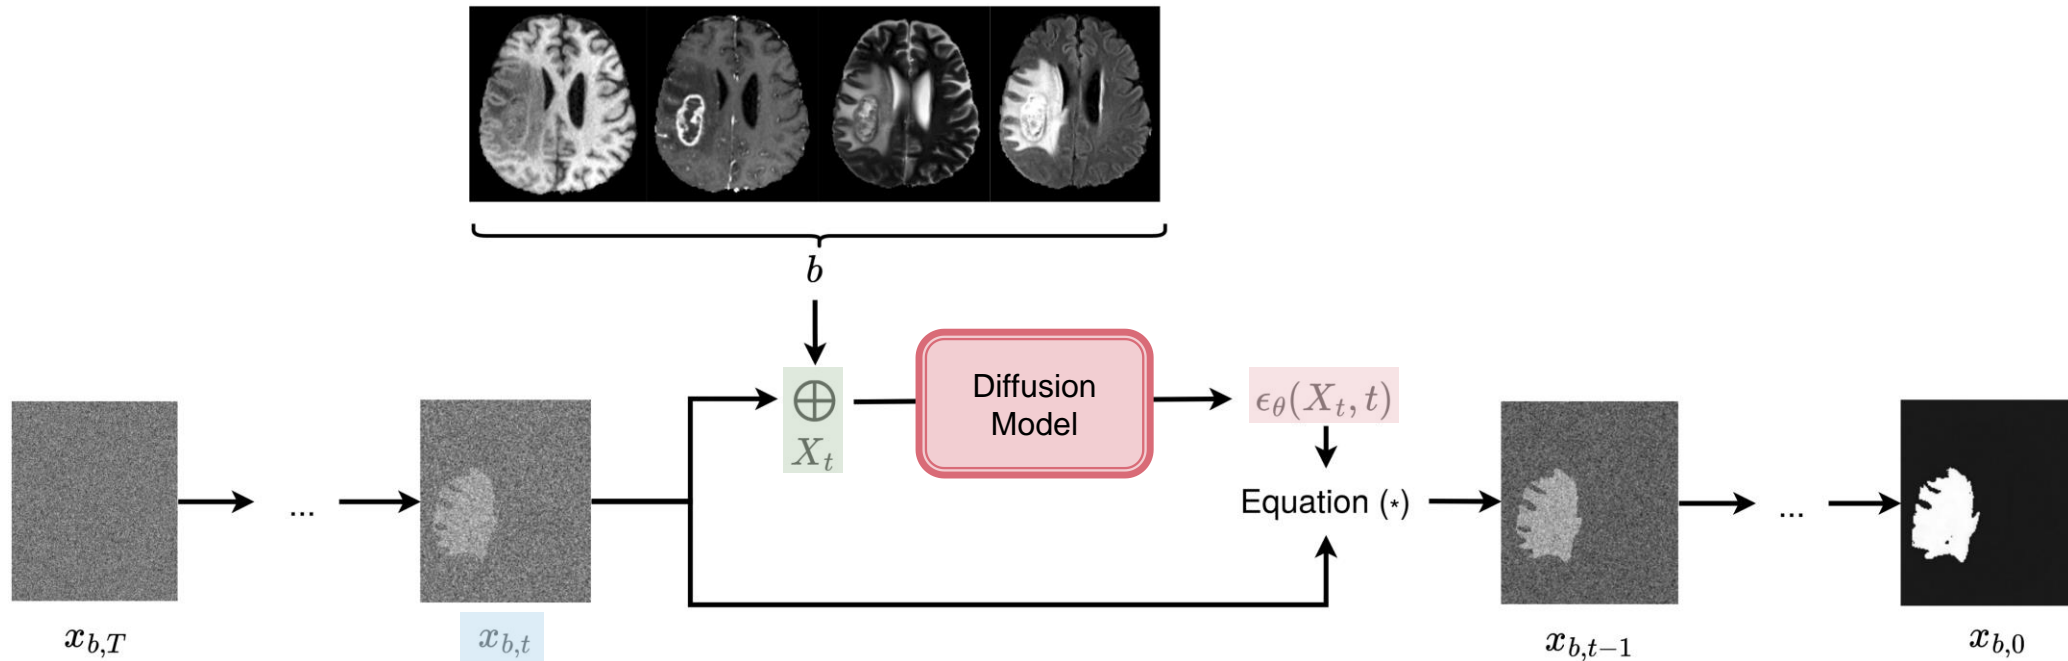


PAPERS

- Wolleb et al (2022). Diffusion Models for Implicit Image Segmentation Ensembles, MIDL 2022
- Bieder et al. (2023) Memory-Efficient 3D Denoising Diffusion Models for Medical Image Processing. MIDL 2023
- Rahman, Aimon, et al. (2023) Ambiguous medical image segmentation using diffusion models. CVPR 2023
- Rousseau et al. (2023) Pre-Training with Diffusion models for Dental Radiography segmentation. MICCAI 2023
- Guo et al. (2023) Accelerating Diffusion Models Via Pre-Segmentation Diffusion Sampling for Medical Image Segmentation. ISBI 2023
- La Barbera et al. (2022) Anatomically constrained CT image translation for heterogeneous blood vessel segmentation. BMVC 2022
- Kim et al. (2022) Diffusion Adversarial Representation Learning for Self-supervised Vessel Segmentation. ICLR 2023
- Wu et al (2022) MedSegDiff: Medical Image Segmentation with Diffusion Probabilistic Model. MIDL 2023

Figure by Song et al ICLR 2022.
Copyright rests with the authors.

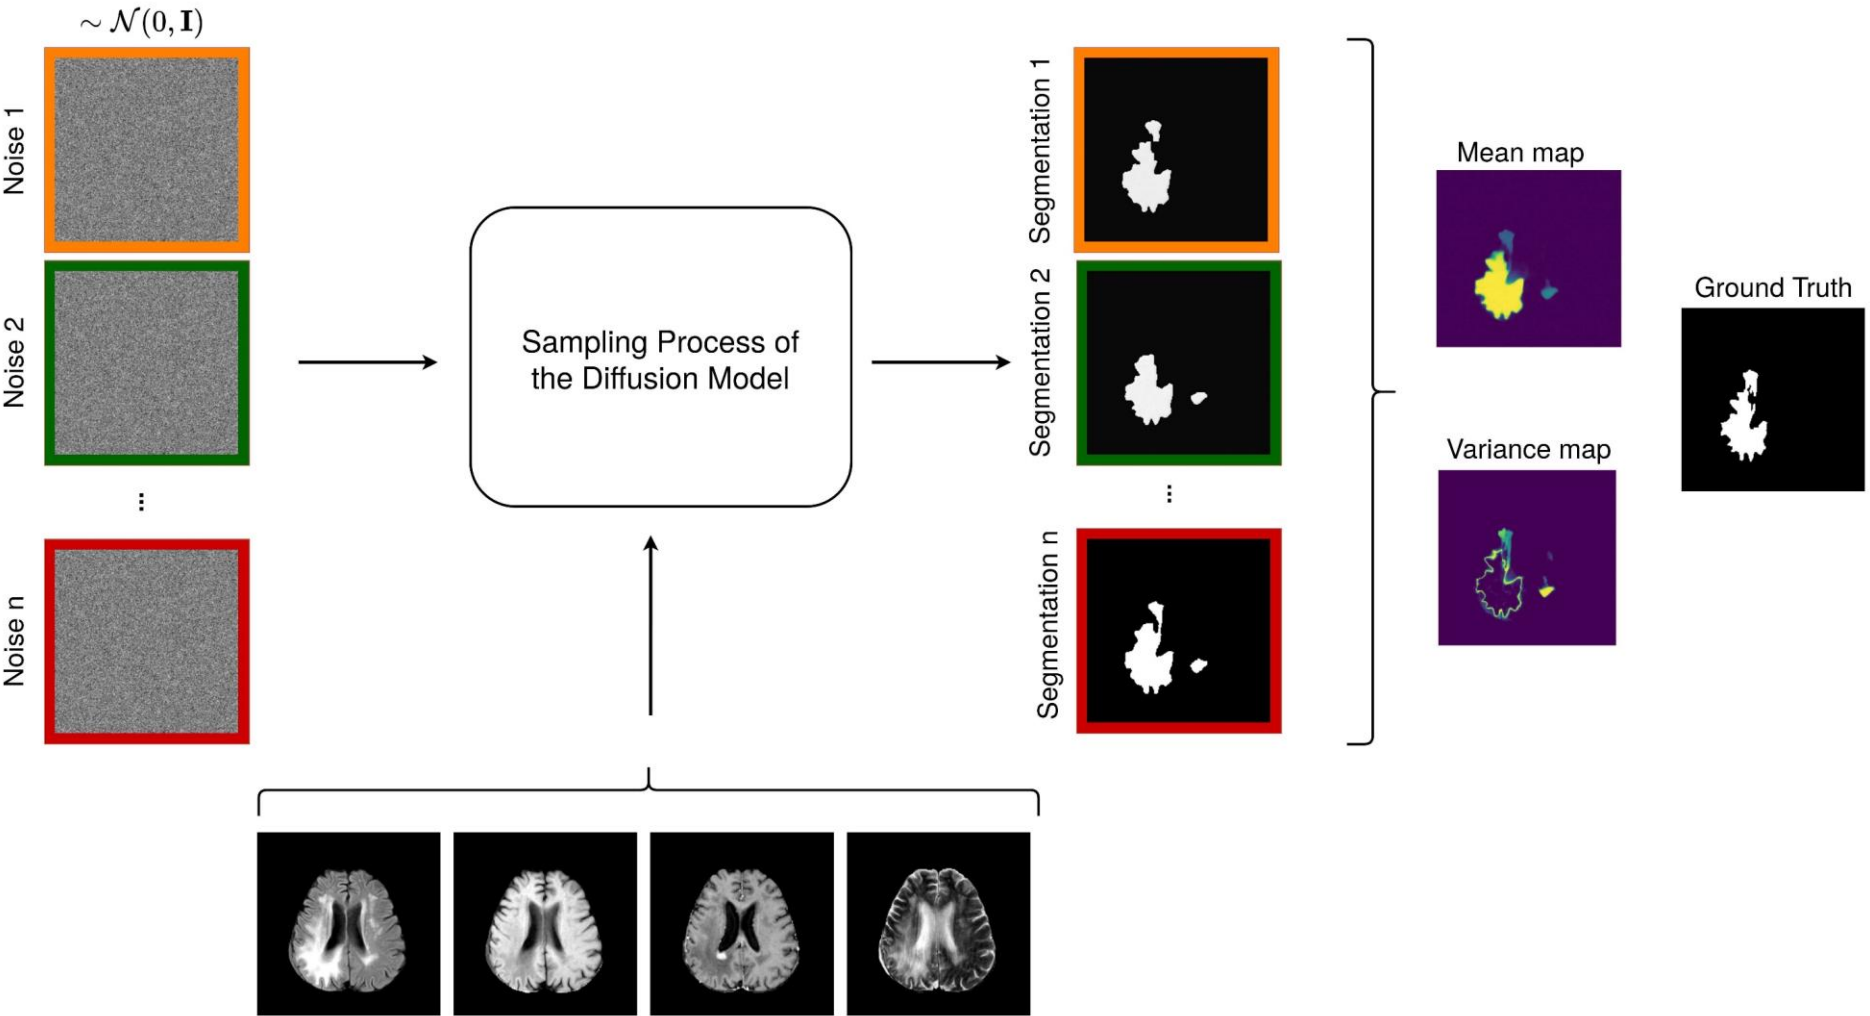
Diffusion Models for Segmentation Mask Generation



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

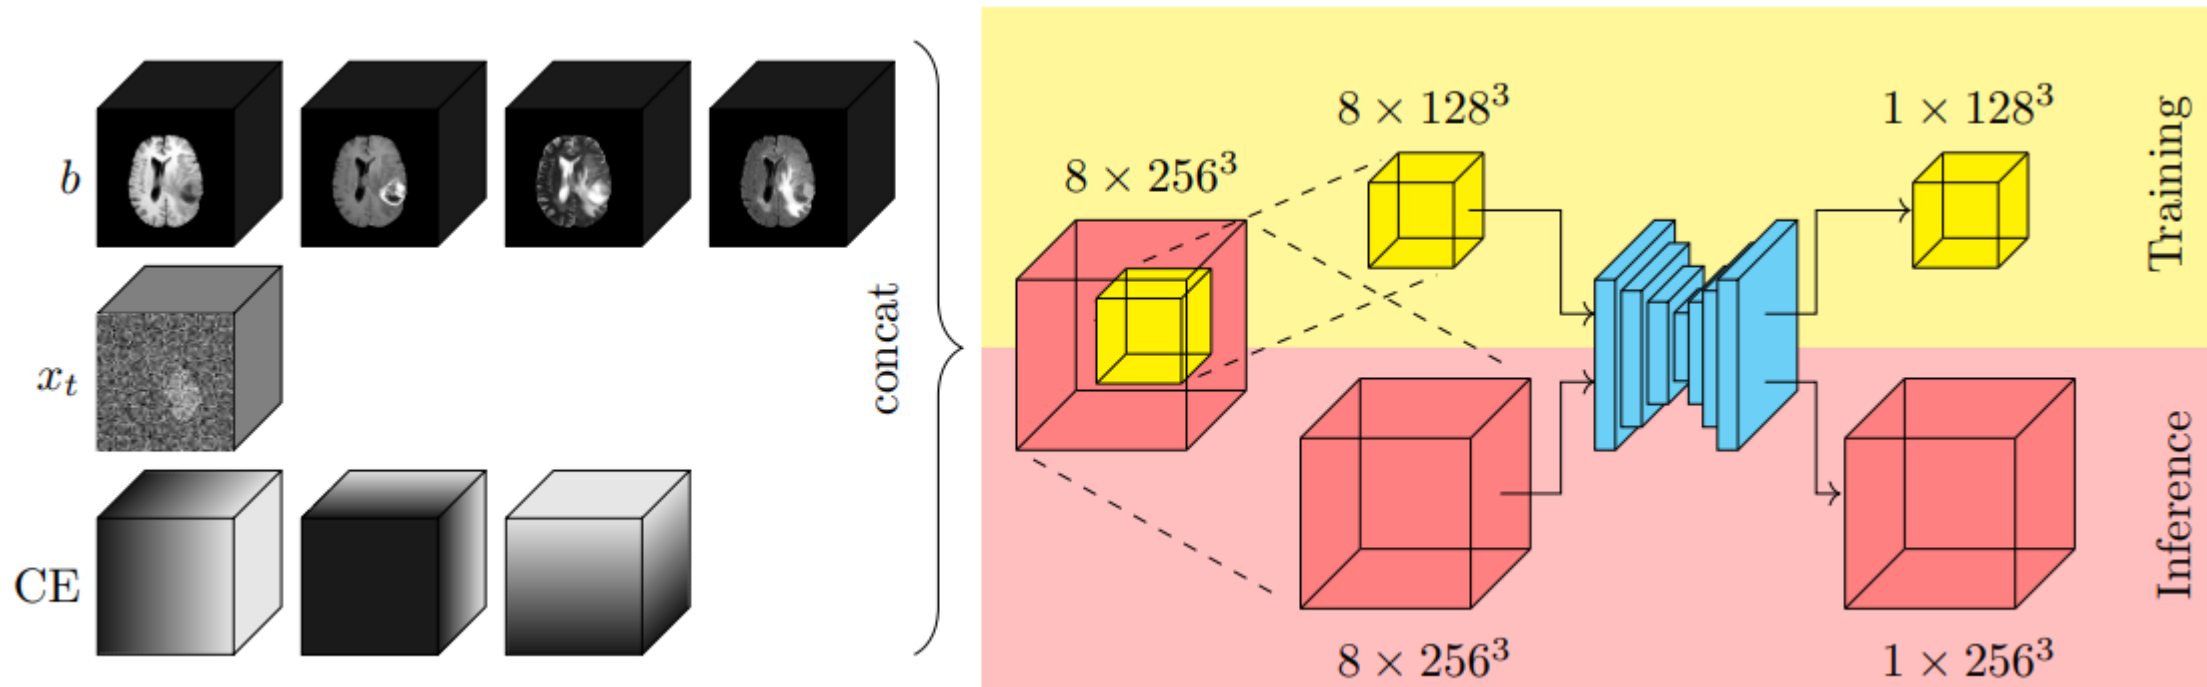
The anatomical information is added by **concatenating** the input images b to the noisy segmentation mask $x_{b,t}$ in every step t .

Generation of Segmentation Ensembles



Corresponding brain MR image b

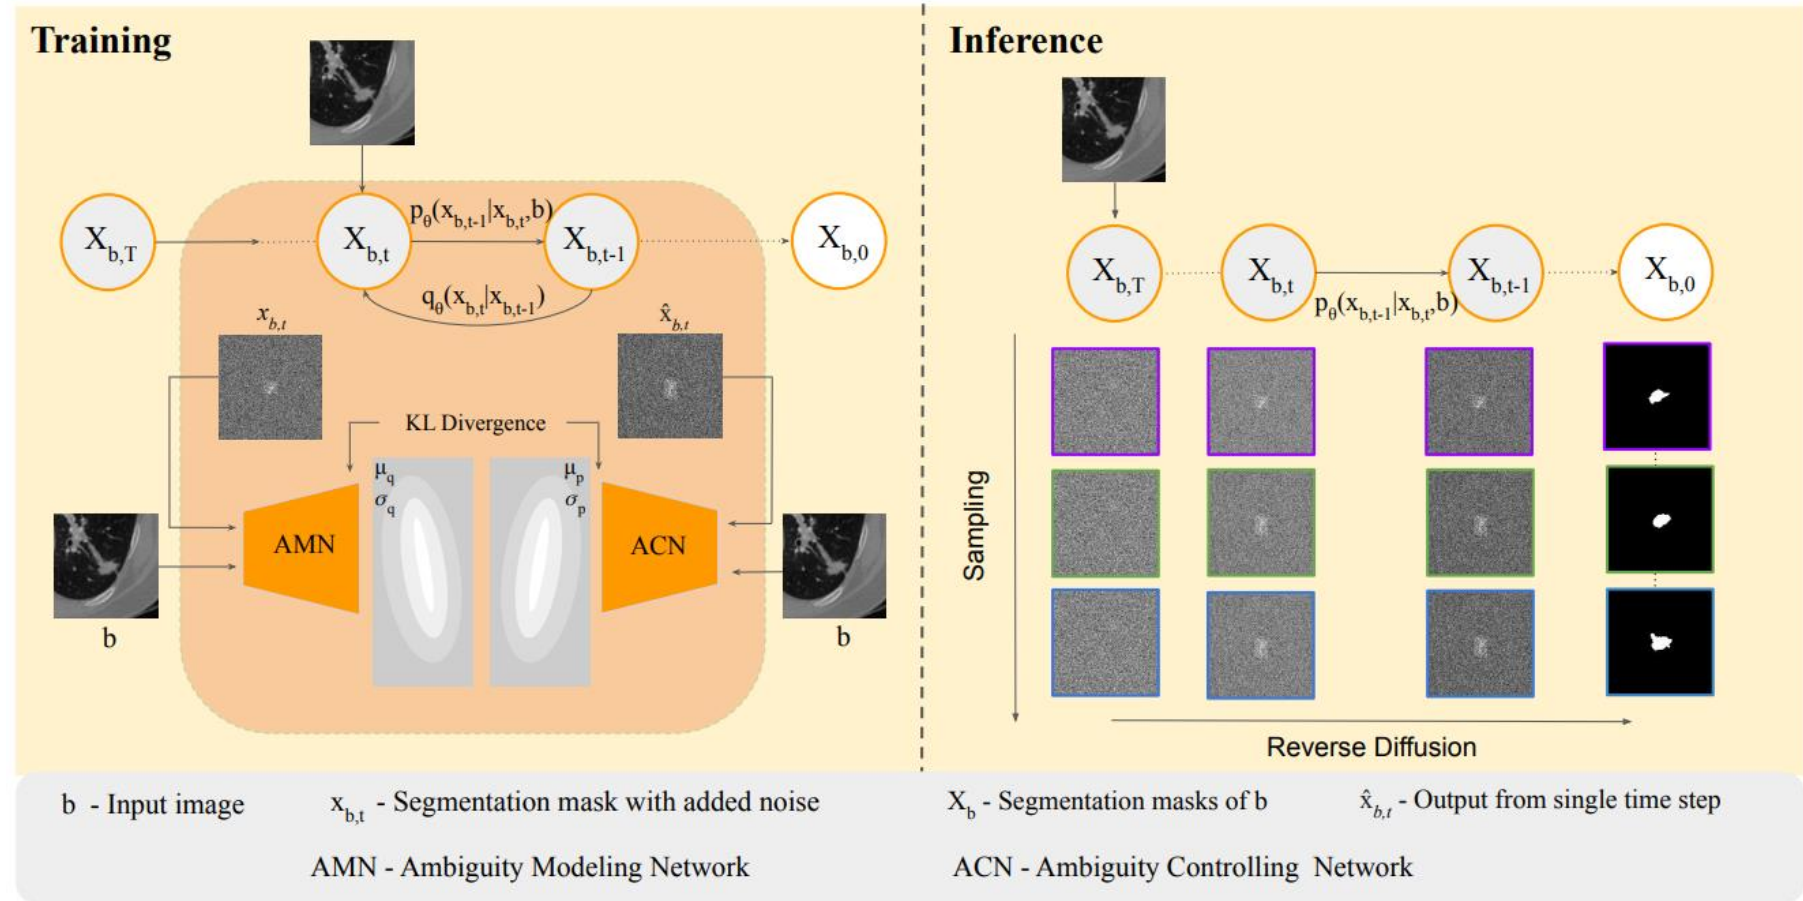
3D Segmentation with PatchDDM



- We add a position encoding in all 3 spatial dimensions.
- Training is on patches only, and saves memory and training time.
- Inference runs over the whole 3D volume.


Ambiguous Segmentation

- Ambiguity Modelling Network (AMN) models the distribution of ground truth masks given an input image.
- Ambiguity Controlling Network (ACN) models the noisy output from the diffusion model conditioning on an input image.




Segmentation with Diffusion Pre-training


Diffusion
 $x_0 \in X_1$
 $\epsilon \sim \mathcal{N}_{0,I}$
 $t \sim \mathcal{U}_{1,T}$
 $\nabla_{\theta} \|\epsilon_{\theta}(x_t, t) - \epsilon\|^2$



Few label
 $X_1 \cap X_2 = \emptyset$
 $(x, y) \in X_2 \times Y$
 $\hat{y} = \epsilon_{\theta}(x)$
 $\nabla_{\theta} Loss(\hat{y}, y)$



Acceleration of Diffusion Segmentation




- They obtain pre-segmentation predictions x_{pre} based on a separately trained segmentation network.
- They start with noisy predictions and use fewer reverse steps T' to generate segmentation results.

With a significantly smaller number of reverse sampling steps, PD-DDPM outperforms the vanilla DDPM

Anomaly detection

Examples from the community

The simple setup of the problem



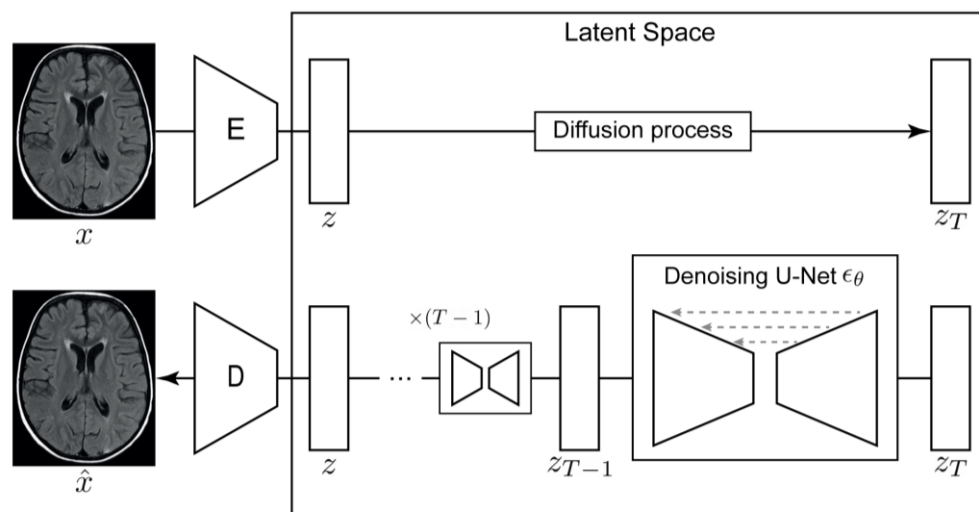
PAPERS

- Pinaya et al (2022) Fast Unsupervised Brain Anomaly Detection and Segmentation with Diffusion Models. MICCAI 2022
- Wyatt et al (2022) AnoDDPM: Anomaly Detection with Denoising Diffusion Probabilistic Models using Simplex Noise. CVPR 2022 workshop
- Kascenas et al (2023) The role of noise in denoising models for anomaly detection in medical images. Medical Image Analysis 2023
- Behrendt, Finn, et al. (2023) Patched diffusion models for unsupervised anomaly detection in brain mri. MIDL 2023
- Liang, Ziyun, et al. (2023) Modality Cycles with Masked Conditional Diffusion for Unsupervised Anomaly Segmentation in MRI. MICCAI 2023
- Wolleb et al (2022). Diffusion Models for Medical Anomaly Detection, MICCAI 2022
- Sanchez et al. (2022) What is Healthy? Generative Counterfactual Diffusion for Lesion Localization. MICCAI 2022 workshop

Figure by Song et al ICLR 2022.
Copyright rests with the authors.



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 being part of the healthy dataset

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

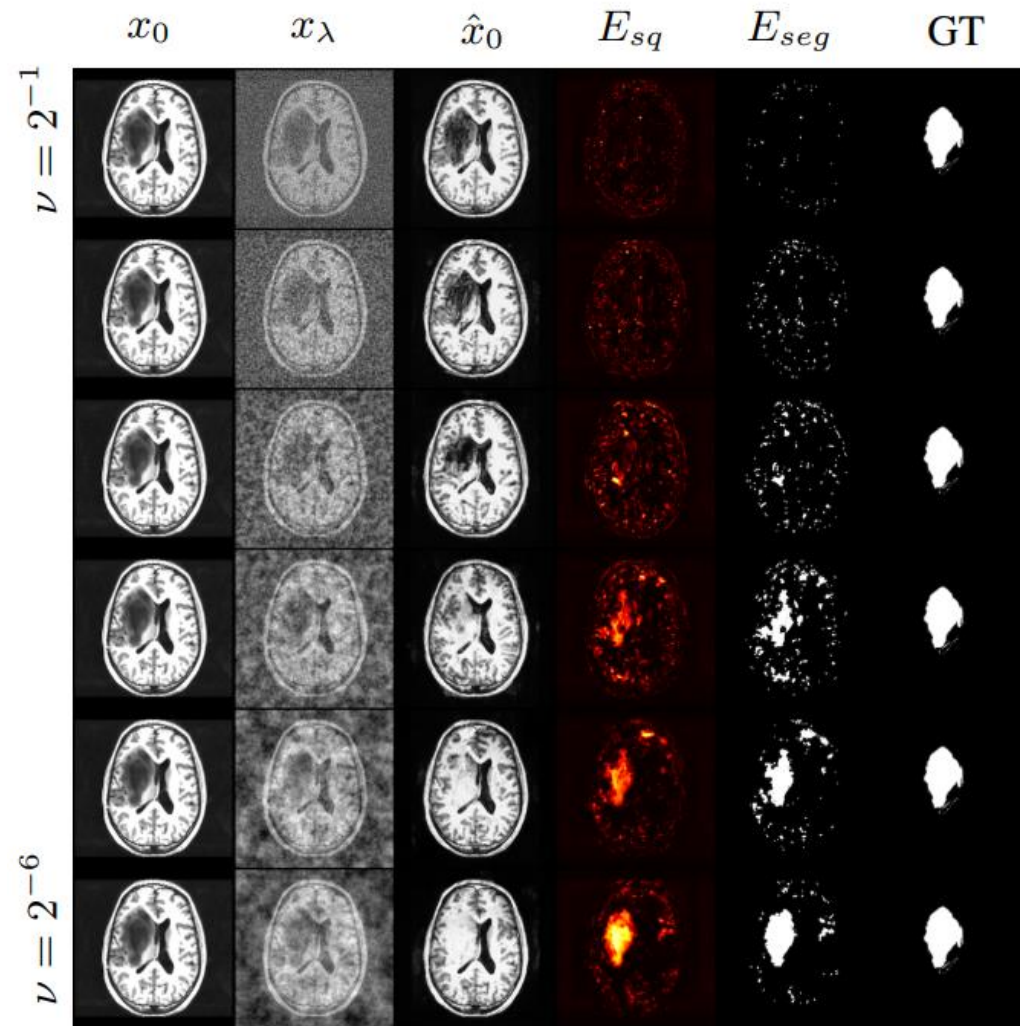


Anomaly Detection with Simplex Noise

- Typical Gaussian noise is found to be insufficient for anomaly detection.
- Therefore, we explore the use of simplex noise for the corruption and sample generation of medical images.

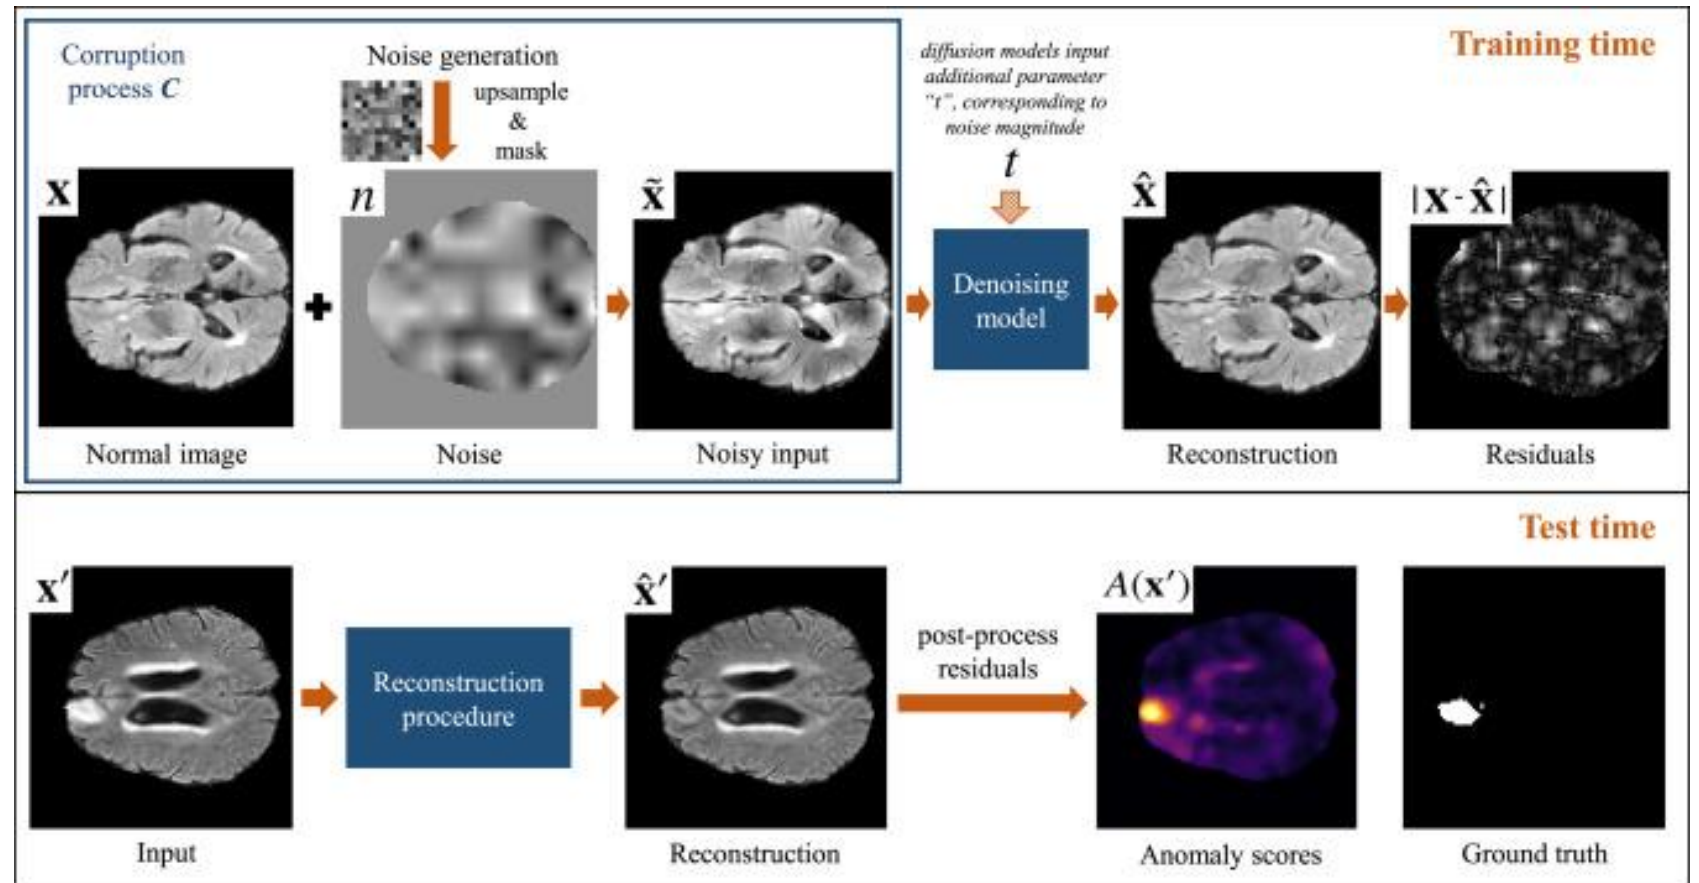



(a) Structures of simplex noise




simplex noise scale *controls* target anomaly size


Anomaly Detection with Coarse Noise



Anomaly Detection from Patches




Anomaly Detection from Modality Cycles





Weakly Supervised Lesion Detection

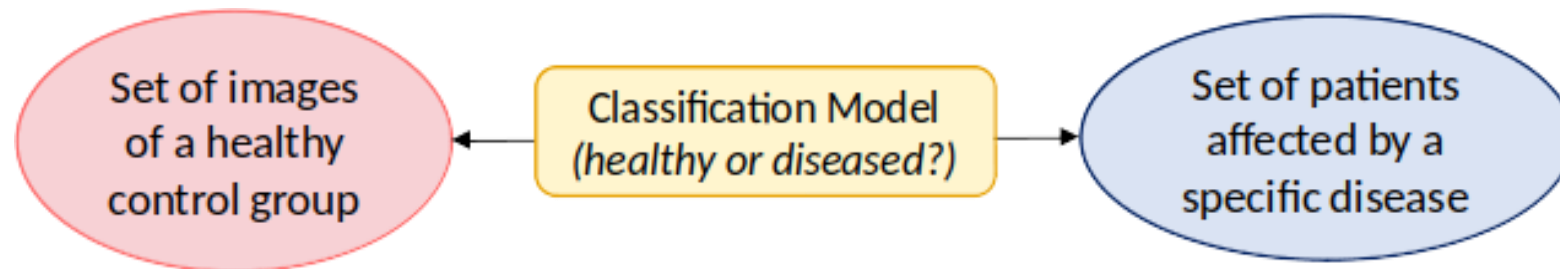
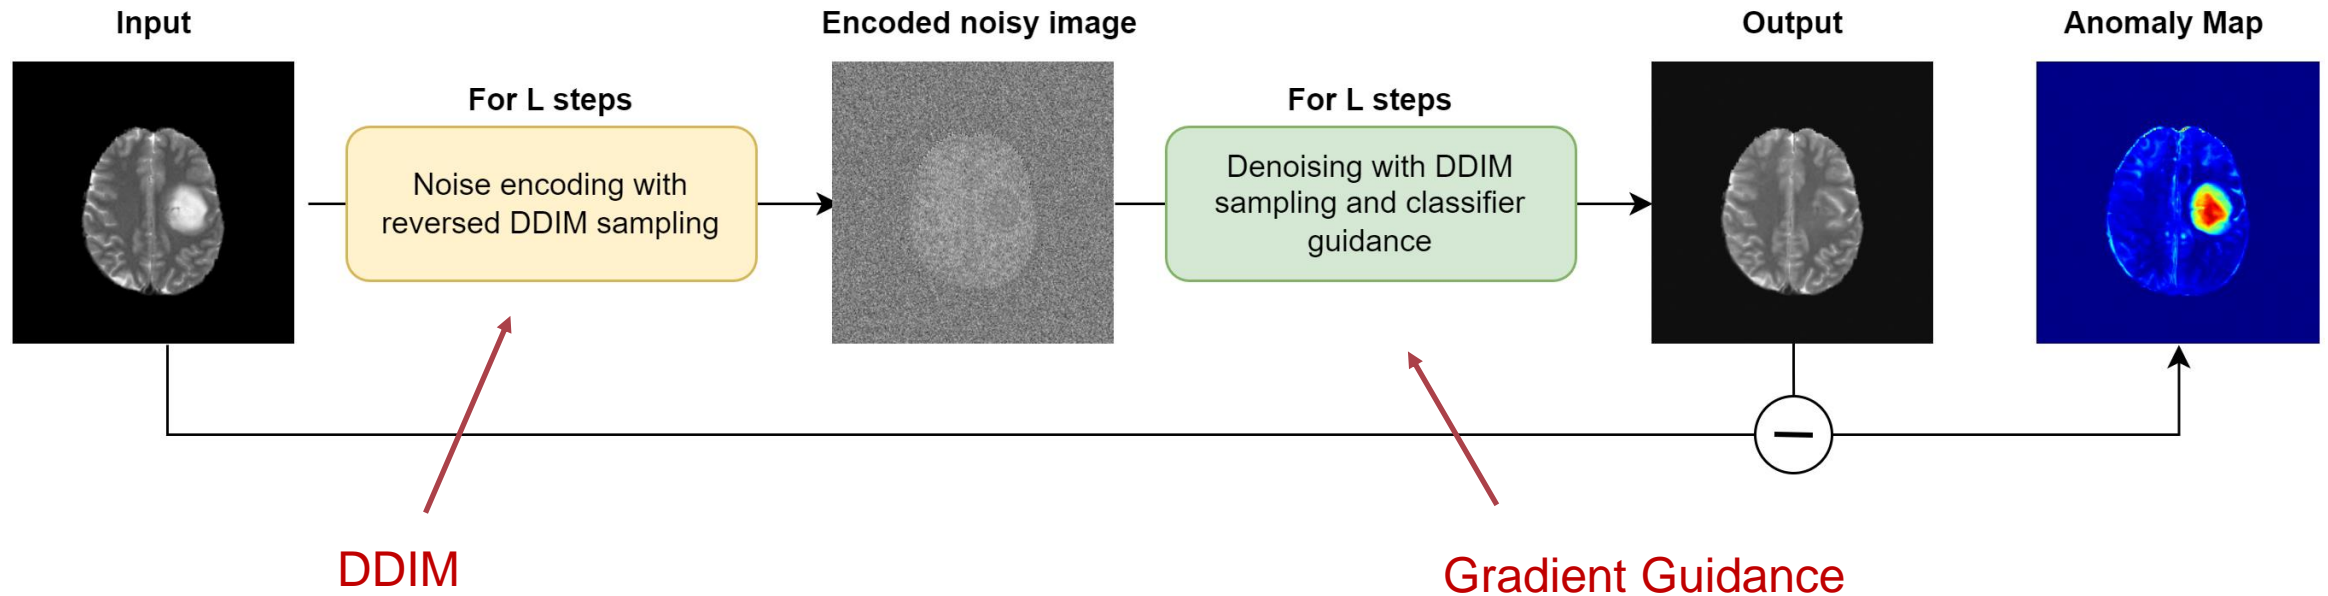
- **Goal:** Pixel-wise anomaly detection using image-level labels only



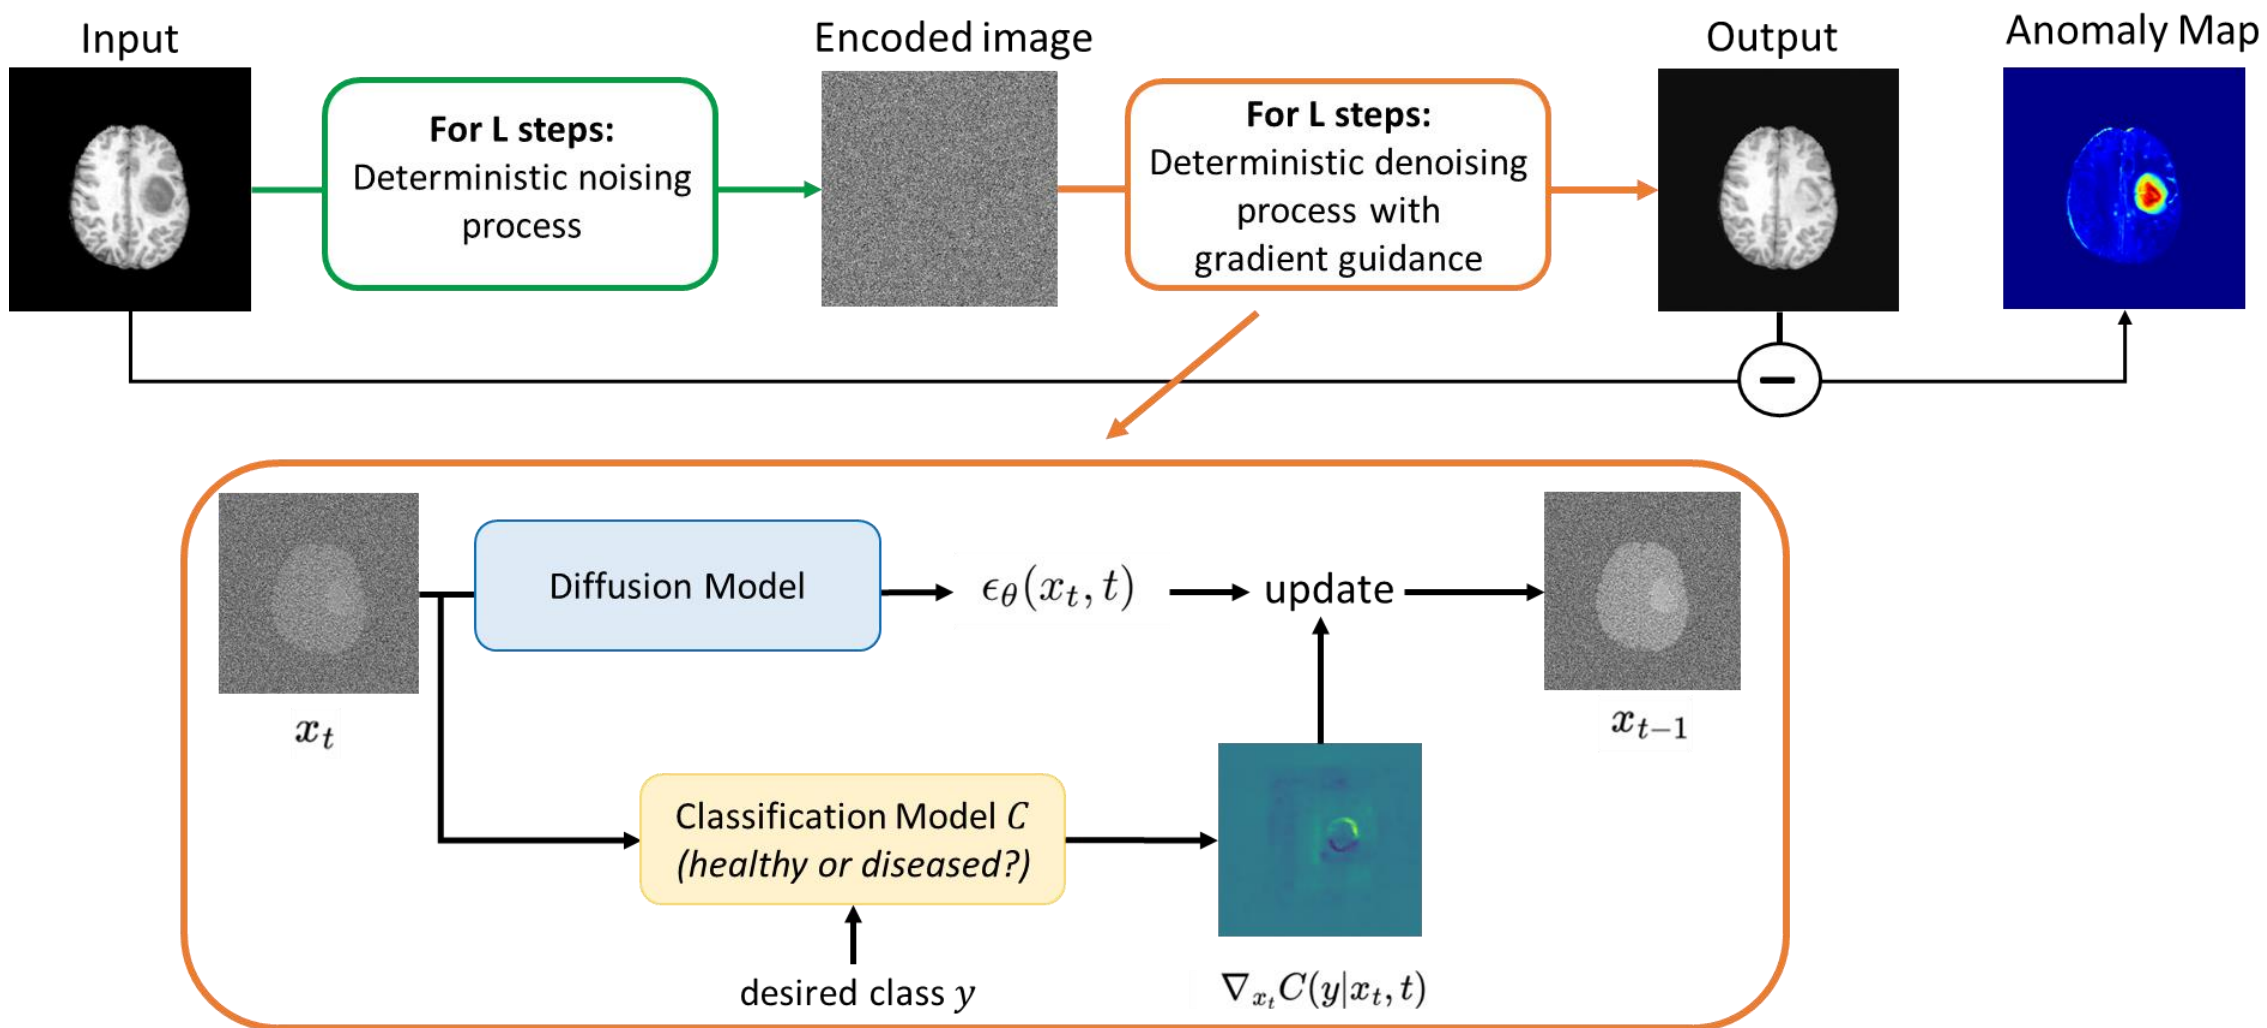
Weakly Supervised Lesion Detection



Weakly Supervised Lesion Detection



Gradient Guidance



Lesion Localization with Diffusion Models

Classifier-free guidance

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

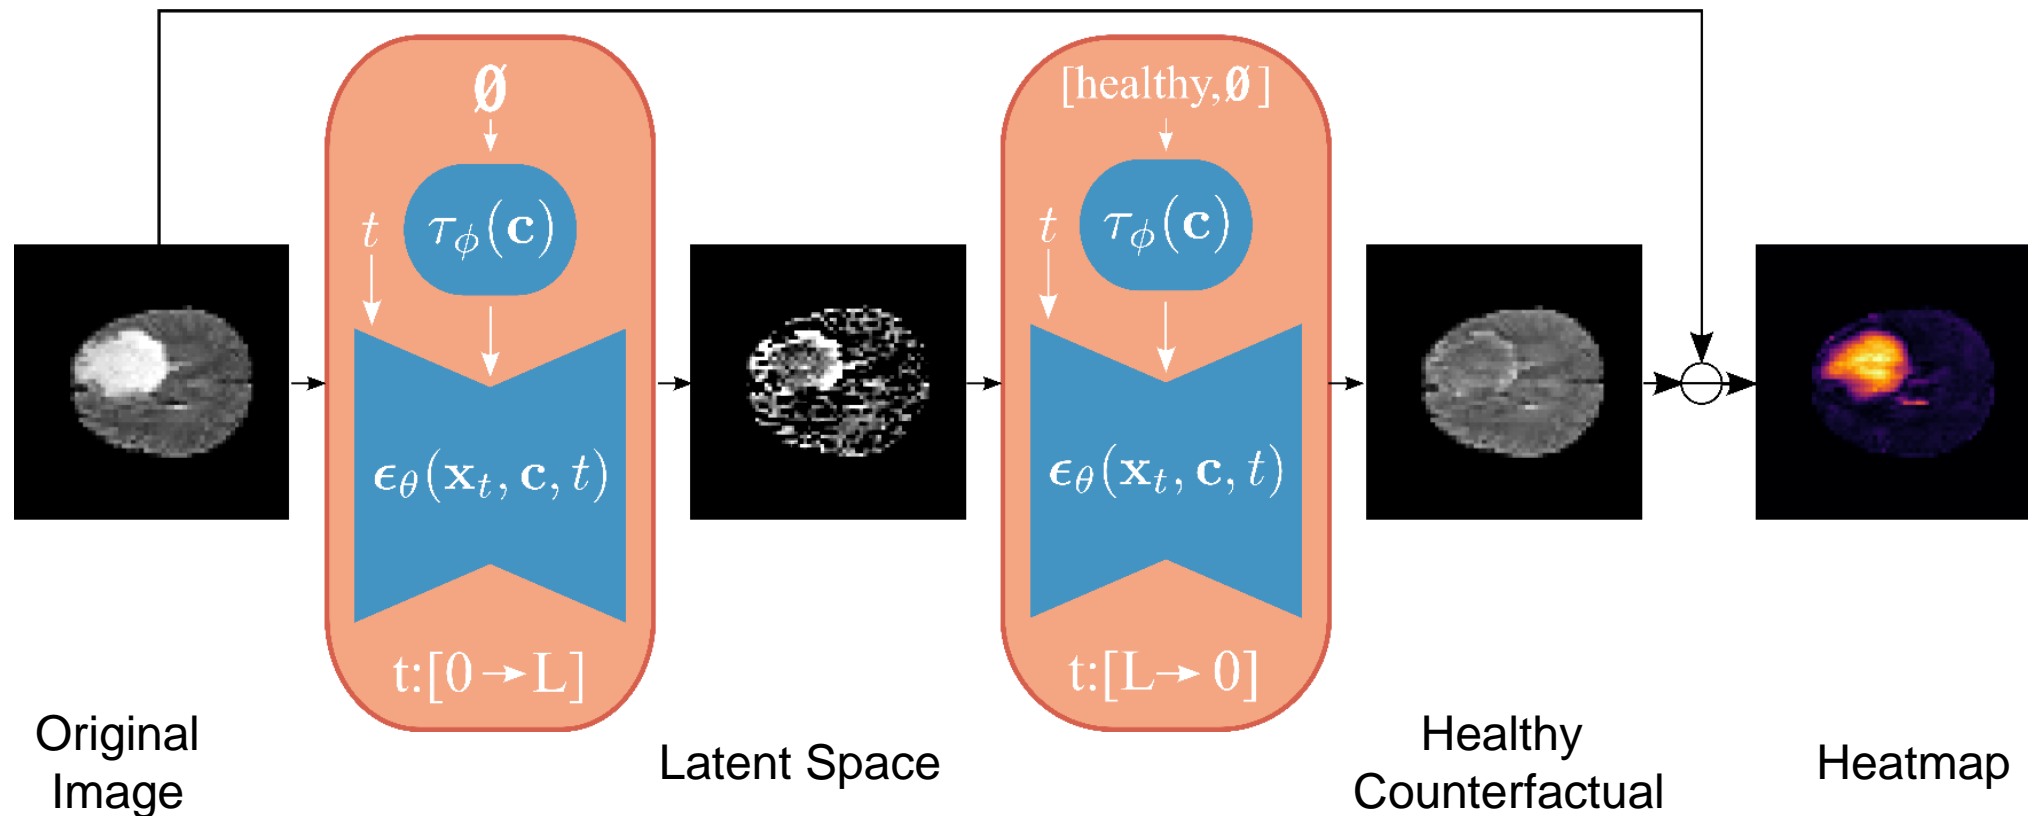




Image reconstruction

Examples from the community

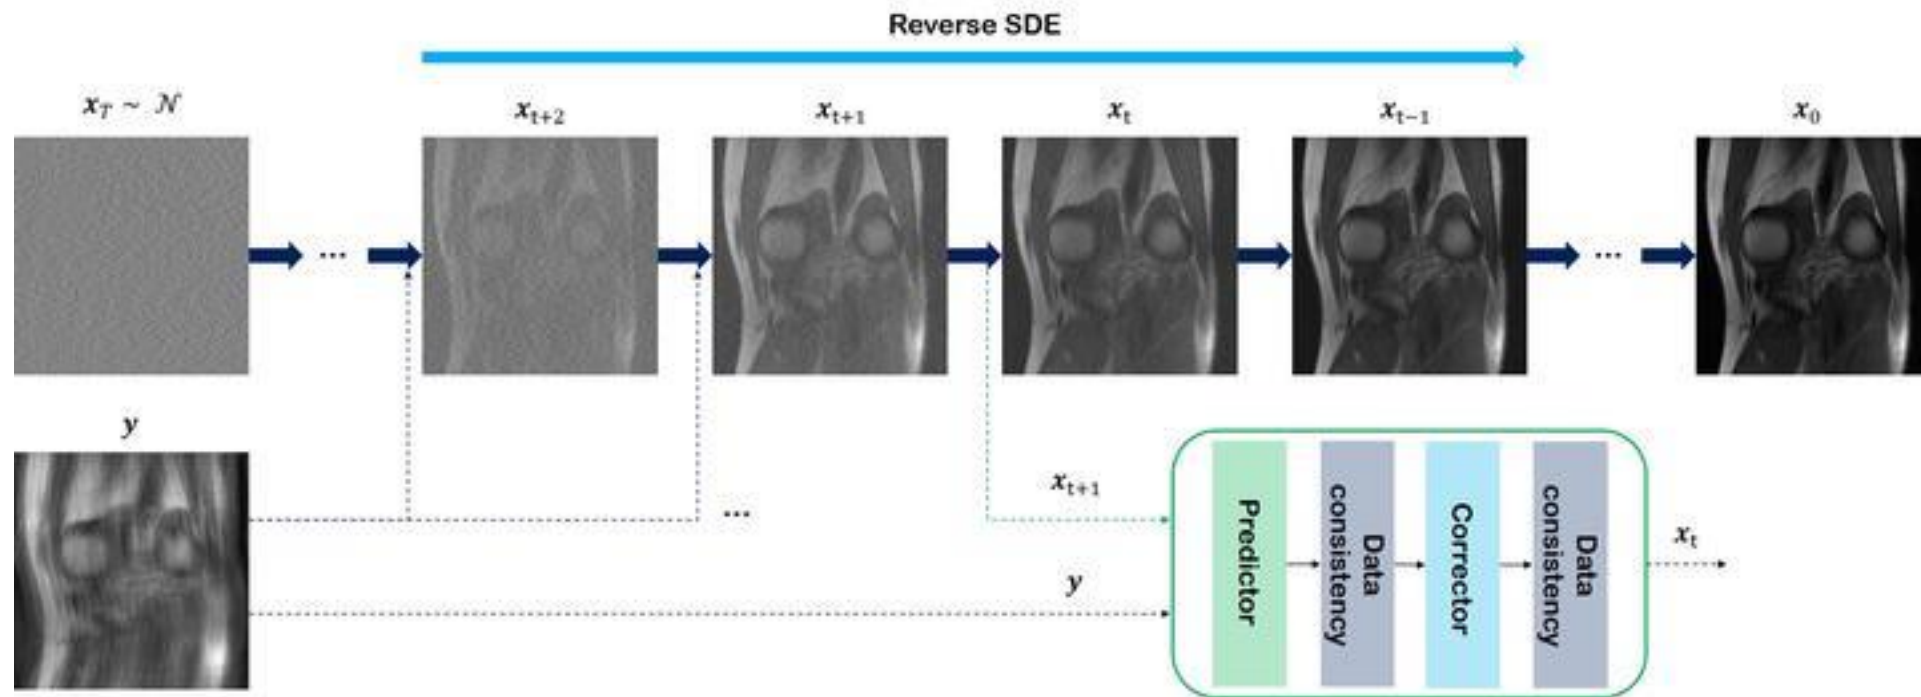
Solving Inverse Problems in Medical Imaging with Score-Based Generative Models

- No need for paired data
- Learn from the prior
 - using a dataset of complete medical images to learn the underlying data distribution
- Conditional Sampling for inverse problem
 - $y = Ax$, y is the measurement and x is the image.
 - Equivalent results for supervised methods



Score-based diffusion models for accelerated MRI

An unconditional diffusion prior is trained on fully-sampled MR acquisitions



Add a **data consistency** term at each sampling step:

$$\mathbf{x}_i \leftarrow \mathbf{x}_i + \lambda A^*(\mathbf{y} - A\mathbf{x}_i)$$

General Inverse Problems

$$\mathbf{y} = \mathcal{A}(\mathbf{x}_0) + \mathbf{n}, \quad \mathbf{y}, \mathbf{n} \in \mathbb{R}^n, \quad \mathbf{x} \in \mathbb{R}^d$$

Algorithm 1 DPS - Gaussian

Require: $N, \mathbf{y}, \{\zeta_i\}_{i=1}^N, \{\tilde{\sigma}_i\}_{i=1}^N$

1: $\mathbf{x}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

2: **for** $i = N - 1$ **to** 0 **do**

3: $\hat{\mathbf{s}} \leftarrow \mathbf{s}_\theta(\mathbf{x}_i, i)$

4: $\hat{\mathbf{x}}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_i}} (\mathbf{x}_i + (1 - \bar{\alpha}_i) \hat{\mathbf{s}})$

5: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

6: $\mathbf{x}'_{i-1} \leftarrow \frac{\sqrt{\bar{\alpha}_i}(1 - \bar{\alpha}_{i-1})}{1 - \bar{\alpha}_i} \mathbf{x}_i + \frac{\sqrt{\bar{\alpha}_{i-1}}\beta_i}{1 - \bar{\alpha}_i} \hat{\mathbf{x}}_0 + \tilde{\sigma}_i \mathbf{z}$

7: $\mathbf{x}_{i-1} \leftarrow \mathbf{x}'_{i-1} - \zeta_i \nabla_{\mathbf{x}_i} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_2^2$

8: **end for**

9: **return** $\hat{\mathbf{x}}_0$

Linear


(a) Inpainting




(c) Gaussian deblur



(b) Super-resolution



(d) Motion deblur



Non-linear

(e) Phase retrieval



(f) Non-uniform deblur

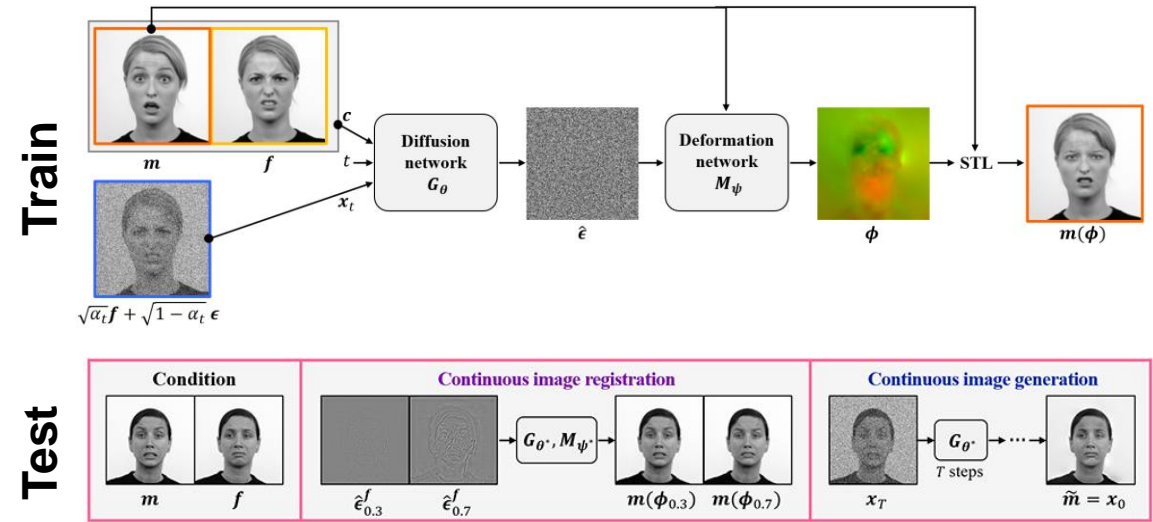


Image registration

Examples from the community

DiffuseMorph

- To perform image registration along the continuous trajectory
- Diffusion network:** To estimate a conditional score function
- Deformation network:** To yield the registration fields & provide the deformed image



Loss function

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

$$\left[\begin{array}{l} L_{diffusion}(c, x_t, t) = \mathbb{E}_{\epsilon, x_t, t} \|G_\theta(c, x_t, t) - \epsilon\|_2^2 \\ L_{regist}(m, f) = -(m(\phi) \otimes f) + \lambda_\phi \sum \|\nabla \phi\|^2 \end{array} \right.$$

Algorithm 1 Continuous image registration

- Input:** Conditional images, $c = (m, f)$
- Output:** Deformed moving image, $m(\phi_\eta)$
- Set the latent feature $\hat{\epsilon}^f = G_{\theta^*}(c, f, 0)$
- for** $\eta \in [0, 1]$ **do**
- $\hat{\epsilon}_\eta^f \leftarrow \eta \cdot \hat{\epsilon}^f$
- $\phi_\eta \leftarrow M_{\psi^*}(m, \hat{\epsilon}_\eta^f)$
- end for**
- return** $m(\phi_\eta)$

Algorithm 2 Synthetic image generation process

- Input:** Conditional images, $c = (m, f)$
- Output:** Synthetic deformed image, x
- Set $T \in (0, T_{train})$
- Sample $x_T = \sqrt{\alpha_T}m + \sqrt{1 - \alpha_T}\epsilon$, where $\epsilon \sim \mathcal{N}(0, I)$
- for** $t = T, T - 1, \dots, 1$ **do**
- $z \sim \mathcal{N}(0, I)$
- $x_{t-1} \leftarrow \frac{1}{\sqrt{1 - \beta_t}}(x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}}G_{\theta^*}(c, x_t, t)) + \sigma_t z$
- end for**
- return** x_0

Useful key references, gits to watch etc

- Surveys
 - <https://arxiv.org/abs/2404.07771>
 - <https://arxiv.org/abs/2209.02646>
 - <https://arxiv.org/abs/2209.00796>
- Github
 - <https://github.com/heejkoo/Awesome-Diffusion-Models>
- Tutorial
 - <https://arxiv.org/pdf/2403.18103.pdf>
 - <https://cvpr2022-tutorial-diffusion-models.github.io>
 - <https://huggingface.co/blog/annotated-diffusion>
 - <https://huggingface.co/docs/diffusers>



THE UNIVERSITY
of EDINBURGH



DiMEDIA: Diffusion Models in Medical Imaging and Analysis

<https://vios.science/tutorials/DiMEDIA-2024>

Nefeli Gkouti
Archimedes/Athena RC,
University of Athens



Prof. Sotirios A. Tsaftaris
Archimedes/Athena RC,
University of Edinburgh





DiMEDIA 2024