

Inria

ESR  
LABORATOIRE  
HUBERT CURIEN  
UNIVERSITÉ PARIS-SACLAY

Mila

CIFAR

# Tutorial on Diffusion Models

Tutorial on ~~Diffusion~~ Conditional Flow Matching

Quentin Bertrand

[QB3.github.io](https://github.com/QB3)

Inria | Laboratoire Hubert Curien | Mila Affiliated Member | CIFAR Global Scholar


anr<sup>®</sup>  
agence nationale  
de la recherche



IVADO


# What is Generative Modeling? A 2D-Example

# What is Generative Modeling? A 2D-Example




# What is Generative Modeling? A 2D-Example

Density  $p_{\text{data}}$  + Samples  $x_1, \dots, x_n \sim p_{\text{data}}$



# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$




# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$

## Problem Setting:

- Access to  $n$  samples  $x_1, \dots, x_n \in \mathbb{R}^2$
- $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$




# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$

## Problem Setting:

- Access to  $n$  samples  $x_1, \dots, x_n \in \mathbb{R}^2$
- $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$



## Goal:

- Draw new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$



# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$

**Goal:** Generate Samples  $x_1^{\text{new}}, \dots, x_n^{\text{new}} \sim p_{\text{data}}$

## Problem Setting:

- Access to  $n$  samples  $x_1, \dots, x_n \in \mathbb{R}^2$
- $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$



## Goal:

- Draw new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$


# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$

**Goal:** Generate Samples  $x_1^{\text{new}}, \dots, x_n^{\text{new}} \sim p_{\text{data}}$

## Problem Setting:

- Access to  $n$  samples  $x_1, \dots, x_n \in \mathbb{R}^2$
- $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$



## Goal:

- Draw new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$


# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$

**Goal:** Generate Samples  $x_1^{\text{new}}, \dots, x_n^{\text{new}} \sim p_{\text{data}}$

## Problem Setting:

- Access to  $n$  samples  $x_1, \dots, x_n \in \mathbb{R}^2$
- $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$



## Goal:

- Draw new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$


# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$

**Goal:** Generate Samples  $x_1^{\text{new}}, \dots, x_n^{\text{new}} \sim p_{\text{data}}$

## Problem Setting:

- Access to  $n$  samples  $x_1, \dots, x_n \in \mathbb{R}^2$
- $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$



## Goal:

- Draw new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$



# What is Generative Modeling? A 2D-Example

Samples  $x_1, \dots, x_n \sim p_{\text{data}}$

**Goal:** Generate Samples  $x_1^{\text{new}}, \dots, x_n^{\text{new}} \sim p_{\text{data}}$

## Problem Setting:

- Access to  $n$  samples  $x_1, \dots, x_n \in \mathbb{R}^2$
- $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$




## Goal:

- Draw new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$

# What is Generative Modeling? An Image Example


# What is Generative Modeling? An Image Example

Data,  $x_1, \dots, x_n \sim p_{\text{data}}$



# What is Generative Modeling? An Image Example

Data,  $x_1, \dots, x_n \sim p_{\text{data}}$



Data:

- Access to samples  $x_1, \dots, x_n$
- Drawn from  $p_{\text{data}}$ ,  $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$
- e.g., set of images, text, proteins, etc

# What is Generative Modeling? An Image Example

Data,  $x_1, \dots, x_n \sim p_{\text{data}}$



Data:

- Access to samples  $x_1, \dots, x_n$
- Drawn from  $p_{\text{data}}$ ,  $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$
- e.g., set of images, text, proteins, etc

Goal:


- Create new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$

# What is Generative Modeling? An Image Example

Data,  $x_1, \dots, x_n \sim p_{\text{data}}$



New Samples,  $x_1^{\text{new}}, \dots, x_n^{\text{new}} \sim p_{\text{data}}$



Data:

- Access to samples  $x_1, \dots, x_n$
- Drawn from  $p_{\text{data}}$ ,  $\underbrace{x_1, \dots, x_n}_{\text{known}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$
- e.g., set of images, text, proteins, etc


Goal:

- Create new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}}_{\text{unknown}}$

# Class-Conditional Generative Modeling


# Class-Conditional Generative Modeling

Data:  $(x_i, y_i), x_i \sim p_{\text{data}}(\cdot | y_i)$



# Class-Conditional Generative Modeling

Data:  $(x_i, y_i), x_i \sim p_{\text{data}}(\cdot | y_i)$




Data:

- Access to samples  $(\underbrace{x_1}_{\text{images}}, \underbrace{y_1}_{\text{label}}), \dots, (x_n, y_n)$
- Drawn from  $p_{\text{data}}, \underbrace{x_i}_{\text{known}} \sim \underbrace{p_{\text{data}}(\cdot | y_i)}_{\text{unknown}}$
- e.g., set of labelled images

# Class-Conditional Generative Modeling

Data:  $(x_i, y_i), x_i \sim p_{\text{data}}(\cdot | y_i)$



Data:

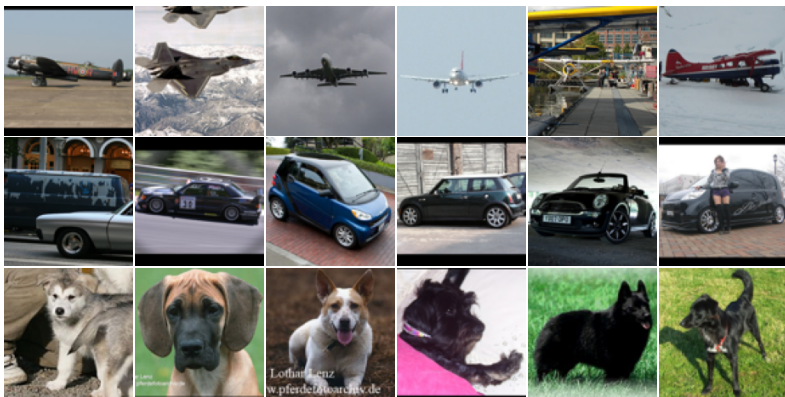
- Access to samples  $(\underbrace{x_1}_{\text{images}}, \underbrace{y_1}_{\text{label}}), \dots, (x_n, y_n)$
- Drawn from  $p_{\text{data}}, \underbrace{x_i}_{\text{known}} \sim \underbrace{p_{\text{data}}(\cdot | y_i)}_{\text{unknown}}$
- e.g., set of labelled images

Goal:


- Create new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}(\cdot | y_i)}_{\text{unknown}}$

# Class-Conditional Generative Modeling

Data:  $(x_i, y_i), x_i \sim p_{\text{data}}(\cdot | y_i)$



Goal:  $x_i^{\text{new}} \sim p_{\text{data}}(\cdot | y_i)$



Data:


- Access to samples  $(\underbrace{x_1}_{\text{images}}, \underbrace{y_1}_{\text{label}}), \dots, (x_n, y_n)$
- Drawn from  $p_{\text{data}}, \underbrace{x_i}_{\text{known}} \sim \underbrace{p_{\text{data}}(\cdot | y_i)}_{\text{unknown}}$
- e.g., set of labelled images

Goal:

- Create new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}(\cdot | y_i)}_{\text{unknown}}$

# Class-Conditional Generative Modeling


Data:  $(x_i, y_i), x_i \sim p_{\text{data}}(\cdot | y_i)$



Data:

- Access to samples  $(\underbrace{x_1}_{\text{images}}, \underbrace{y_1}_{\text{label}}), \dots, (x_n, y_n)$
- Drawn from  $p_{\text{data}}, \underbrace{x_i}_{\text{known}} \sim \underbrace{p_{\text{data}}(\cdot | y_i)}_{\text{unknown}}$
- e.g., set of labelled images

Goal:  $x_i^{\text{new}} \sim p_{\text{data}}(\cdot | y_i)$



Goal:

- Create new samples  $\underbrace{x_1^{\text{new}}, \dots, x_n^{\text{new}}}_{\text{Goal}} \sim \underbrace{p_{\text{data}}(\cdot | y_i)}_{\text{unknown}}$
- Comment:
- Generative learning:** given a class  $y_i$ , generate an image  $x_i^{\text{new}} \sim p_{\text{data}}(\cdot | y_i)$

# Text-to-Image Generative Modeling



MALICE team logo



An avocado chair



A house made of sushi


# Text-to-Image Generative Modeling



MALICE team logo



An avocado chair



A house made of sushi

Data:

- Access to samples  $(\underbrace{x_1}_{\text{image}}, \underbrace{t_1}_{\text{text}}), \dots, (x_n, t_n)$
- Drawn from an unknown distribution  $\underbrace{(x_i, t_i)}_{\text{known}} \sim$

$\underbrace{p_{\text{data}}(x, t)}_{\text{unknown}}$


# Text-to-Image Generative Modeling



MALICE team logo



An avocado chair



A house made of sushi

Data:

- Access to samples  $(\underbrace{x_1}_{\text{image}}, \underbrace{t_1}_{\text{text}}), \dots, (x_n, t_n)$
- Drawn from an unknown distribution  $\underbrace{(x_i, t_i)}_{\text{known}} \sim$

$\underbrace{p_{\text{data}}(x, t)}_{\text{unknown}}$

Goal:

- Generate images conditioned on **new** text
- Sample:  $\underbrace{x^{\text{new}}}_{\text{image}} \sim p_{\theta}(x \mid \underbrace{t}_{\text{unseen at training}})$


# Text-to-Image Generative Modeling



MALICE team logo



An avocado chair



A house made of sushi

Data:

- Access to samples  $(\underbrace{x_1}_{\text{image}}, \underbrace{t_1}_{\text{text}}), \dots, (x_n, t_n)$
- Drawn from an unknown distribution  $\underbrace{(x_i, t_i)}_{\text{known}} \sim$

$\underbrace{p_{\text{data}}(x, t)}_{\text{unknown}}$

Goal:

- Generate images conditioned on **new** text
- Sample:  $\underbrace{x^{\text{new}}}_{\text{image}} \sim p_{\theta}(x \mid \underbrace{t}_{\text{unseen at training}})$

**Comment:** text is harder

- Discrete
- Generalization to unseen prompts

# A Few Comments

# A Few Comments

## Text-to-Text Models

- Conditional text generation w.r.t. the previous text sequence

# A Few Comments

## Text-to-Text Models

- Conditional text generation w.r.t. the previous text sequence

"From Noise to Structure"

# A Few Comments

## Text-to-Text Models

- Conditional text generation w.r.t. the previous text sequence

## "From Noise to Structure"

- Just samples, no density
  - Too hard to sample from the set of natural images
  - Easy to sample from standard Gaussian noise
  - "Rebranding of sampling"

# A Few Comments

## Text-to-Text Models

- Conditional text generation w.r.t. the previous text sequence

## "From Noise to Structure"

- Just samples, no density
  - Too hard to sample from the set of natural images
  - Easy to sample from standard Gaussian noise
  - "Rebranding of sampling"

## Evaluation?

- Quality of the generated images?
  - No cross-validation
  - Very challenging for text-to-text models
  - Often use other models to assess quality (e.g. GPT)

# Zoo of Generative Models <sup>1</sup>

---

<sup>1</sup> F. Caetano, [GenerativeZoo: A Model Zoo for Generative Models](#), In: GitHub repository, 2024.

<sup>2</sup> D. Kingma et al., [Auto-encoding variational bayes](#), In: ICLR, 2014.

<sup>3</sup> C. Doersch, [Tutorial on variational autoencoders](#), In: arXiv preprint arXiv:1606.05908, 2016.

<sup>4</sup> I. Goodfellow et al., [Generative adversarial nets](#), In: NeurIPS, 2014.

<sup>5</sup> J. Zhu et al., [Unpaired image-to-image translation using cycle-consistent adversarial networks](#), In: ICCV, 2017.

<sup>6</sup> I. Gulrajani et al., [Improved training of Wasserstein gans](#), In: NeurIPS, 2017.

<sup>7</sup> A. Van Den Oord et al., [Pixel recurrent neural networks](#), In: International conference on machine learning, 2016.

<sup>8</sup> A. Van Den Oord et al., [Wavenet: A generative model for raw audio](#), In: arXiv preprint arXiv:1609.03499, 2016.

<sup>9</sup> J. Sohl-Dickstein et al., [Deep unsupervised learning using nonequilibrium thermodynamics](#), In: ICML, 2015.

<sup>10</sup> J. Ho et al., [Denosing diffusion probabilistic models](#), In: Advances in neural information processing systems, 2020.

<sup>11</sup> R. Gao et al., [Diffusion Meets Flow Matching: Two Sides of the Same Coin](#), In: , 2024.

<sup>12</sup> Y. Lipman et al., [Flow matching for generative modeling](#), In: ICLR, 2023.

<sup>13</sup> X. Liu et al., [Flow straight and fast: Learning to generate and transfer data with rectified flow](#), In: ICLR, 2023.

<sup>14</sup> A. Gagneux et al., [A Visual Dive into Conditional Flow Matching](#), In: ICLR Blogposts, 2025.

# Zoo of Generative Models <sup>1</sup>

- Variational Autoencoders (2014) <sup>2 3</sup>

---

<sup>1</sup> F. Caetano, [GenerativeZoo: A Model Zoo for Generative Models](#), In: GitHub repository, 2024.

<sup>2</sup> D. Kingma et al., [Auto-encoding variational bayes](#), In: ICLR, 2014.

<sup>3</sup> C. Doersch, [Tutorial on variational autoencoders](#), In: arXiv preprint arXiv:1606.05908, 2016.

<sup>4</sup> I. Goodfellow et al., [Generative adversarial nets](#), In: NeurIPS, 2014.

<sup>5</sup> J. Zhu et al., [Unpaired image-to-image translation using cycle-consistent adversarial networks](#), In: ICCV, 2017.

<sup>6</sup> I. Gulrajani et al., [Improved training of Wasserstein gans](#), In: NeurIPS, 2017.

<sup>7</sup> A. Van Den Oord et al., [Pixel recurrent neural networks](#), In: International conference on machine learning, 2016.

<sup>8</sup> A. Van Den Oord et al., [Wavenet: A generative model for raw audio](#), In: arXiv preprint arXiv:1609.03499, 2016.

<sup>9</sup> J. Sohl-Dickstein et al., [Deep unsupervised learning using nonequilibrium thermodynamics](#), In: ICML, 2015.

<sup>10</sup> J. Ho et al., [Denosing diffusion probabilistic models](#), In: Advances in neural information processing systems, 2020.

<sup>11</sup> R. Gao et al., [Diffusion Meets Flow Matching: Two Sides of the Same Coin](#), In: , 2024.

<sup>12</sup> Y. Lipman et al., [Flow matching for generative modeling](#), In: ICLR, 2023.

<sup>13</sup> X. Liu et al., [Flow straight and fast: Learning to generate and transfer data with rectified flow](#), In: ICLR, 2023.

<sup>14</sup> A. Gagneux et al., [A Visual Dive into Conditional Flow Matching](#), In: ICLR Blogposts, 2025.

# Zoo of Generative Models <sup>1</sup>

- Variational Autoencoders (2014) <sup>2 3</sup>
- Generative Adversarial Networks (GANs) <sup>4</sup> (2014)
  - Cycle GANs <sup>5</sup>
  - Wasserstein GANs <sup>6</sup>

---

<sup>1</sup> F. Caetano, [GenerativeZoo: A Model Zoo for Generative Models](#), In: GitHub repository, 2024.

<sup>2</sup> D. Kingma et al., [Auto-encoding variational bayes](#), In: ICLR, 2014.

<sup>3</sup> C. Doersch, [Tutorial on variational autoencoders](#), In: arXiv preprint arXiv:1606.05908, 2016.

<sup>4</sup> I. Goodfellow et al., [Generative adversarial nets](#), In: NeurIPS, 2014.

<sup>5</sup> J. Zhu et al., [Unpaired image-to-image translation using cycle-consistent adversarial networks](#), In: ICCV, 2017.

<sup>6</sup> I. Gulrajani et al., [Improved training of Wasserstein gans](#), In: NeurIPS, 2017.

<sup>7</sup> A. Van Den Oord et al., [Pixel recurrent neural networks](#), In: International conference on machine learning, 2016.

<sup>8</sup> A. Van Den Oord et al., [Wavenet: A generative model for raw audio](#), In: arXiv preprint arXiv:1609.03499, 2016.

<sup>9</sup> J. Sohl-Dickstein et al., [Deep unsupervised learning using nonequilibrium thermodynamics](#), In: ICML, 2015.

<sup>10</sup> J. Ho et al., [Denosing diffusion probabilistic models](#), In: Advances in neural information processing systems, 2020.

<sup>11</sup> R. Gao et al., [Diffusion Meets Flow Matching: Two Sides of the Same Coin](#), In: , 2024.

<sup>12</sup> Y. Lipman et al., [Flow matching for generative modeling](#), In: ICLR, 2023.

<sup>13</sup> X. Liu et al., [Flow straight and fast: Learning to generate and transfer data with rectified flow](#), In: ICLR, 2023.

<sup>14</sup> A. Gagneux et al., [A Visual Dive into Conditional Flow Matching](#), In: ICLR Blogposts, 2025.

# Zoo of Generative Models <sup>1</sup>

- Variational Autoencoders (2014) <sup>2 3</sup>
- Generative Adversarial Networks (GANs) <sup>4</sup> (2014)
  - Cycle GANs <sup>5</sup>
  - Wasserstein GANs <sup>6</sup>
- Autoregressive Models (2016)
  - Pixel RNNs <sup>7</sup>
  - Wavenets <sup>8</sup>

---

<sup>1</sup> F. Caetano, [GenerativeZoo: A Model Zoo for Generative Models](#), In: GitHub repository, 2024.

<sup>2</sup> D. Kingma et al., [Auto-encoding variational bayes](#), In: ICLR, 2014.

<sup>3</sup> C. Doersch, [Tutorial on variational autoencoders](#), In: arXiv preprint arXiv:1606.05908, 2016.

<sup>4</sup> I. Goodfellow et al., [Generative adversarial nets](#), In: NeurIPS, 2014.

<sup>5</sup> J. Zhu et al., [Unpaired image-to-image translation using cycle-consistent adversarial networks](#), In: ICCV, 2017.

<sup>6</sup> I. Gulrajani et al., [Improved training of Wasserstein gans](#), In: NeurIPS, 2017.

<sup>7</sup> A. Van Den Oord et al., [Pixel recurrent neural networks](#), In: International conference on machine learning, 2016.

<sup>8</sup> A. Van Den Oord et al., [Wavenet: A generative model for raw audio](#), In: arXiv preprint arXiv:1609.03499, 2016.

<sup>9</sup> J. Sohl-Dickstein et al., [Deep unsupervised learning using nonequilibrium thermodynamics](#), In: ICML, 2015.

<sup>10</sup> J. Ho et al., [Denosing diffusion probabilistic models](#), In: Advances in neural information processing systems, 2020.

<sup>11</sup> R. Gao et al., [Diffusion Meets Flow Matching: Two Sides of the Same Coin](#), In: , 2024.

<sup>12</sup> Y. Lipman et al., [Flow matching for generative modeling](#), In: ICLR, 2023.

<sup>13</sup> X. Liu et al., [Flow straight and fast: Learning to generate and transfer data with rectified flow](#), In: ICLR, 2023.

<sup>14</sup> A. Gagneux et al., [A Visual Dive into Conditional Flow Matching](#), In: ICLR Blogposts, 2025.

# Zoo of Generative Models <sup>1</sup>

- Variational Autoencoders (2014) <sup>2 3</sup>
- Generative Adversarial Networks (GANs) <sup>4</sup> (2014)
  - Cycle GANs <sup>5</sup>
  - Wasserstein GANs <sup>6</sup>
- Autoregressive Models (2016)
  - Pixel RNNs <sup>7</sup>
  - Wavenets <sup>8</sup>
- Diffusion (2020) <sup>9 10 11</sup> / Flow Matching Models (2023) <sup>12 13 14</sup>

---

<sup>1</sup> F. Caetano, [GenerativeZoo: A Model Zoo for Generative Models](#), In: GitHub repository, 2024.

<sup>2</sup> D. Kingma et al., [Auto-encoding variational bayes](#), In: ICLR, 2014.

<sup>3</sup> C. Doersch, [Tutorial on variational autoencoders](#), In: arXiv preprint arXiv:1606.05908, 2016.

<sup>4</sup> I. Goodfellow et al., [Generative adversarial nets](#), In: NeurIPS, 2014.

<sup>5</sup> J. Zhu et al., [Unpaired image-to-image translation using cycle-consistent adversarial networks](#), In: ICCV, 2017.

<sup>6</sup> I. Gulrajani et al., [Improved training of Wasserstein gans](#), In: NeurIPS, 2017.

<sup>7</sup> A. Van Den Oord et al., [Pixel recurrent neural networks](#), In: International conference on machine learning, 2016.

<sup>8</sup> A. Van Den Oord et al., [Wavenet: A generative model for raw audio](#), In: arXiv preprint arXiv:1609.03499, 2016.

<sup>9</sup> J. Sohl-Dickstein et al., [Deep unsupervised learning using nonequilibrium thermodynamics](#), In: ICML, 2015.


<sup>10</sup> J. Ho et al., [Denosing diffusion probabilistic models](#), In: Advances in neural information processing systems, 2020.

<sup>11</sup> R. Gao et al., [Diffusion Meets Flow Matching: Two Sides of the Same Coin](#), In: , 2024.

<sup>12</sup> Y. Lipman et al., [Flow matching for generative modeling](#), In: ICLR, 2023.


<sup>13</sup> X. Liu et al., [Flow straight and fast: Learning to generate and transfer data with rectified flow](#), In: ICLR, 2023.

<sup>14</sup> A. Gagneux et al., [A Visual Dive into Conditional Flow Matching](#), In: ICLR Blogposts, 2025.



# Normalizing Flows a.k.a., Iterative Refinement

# Normalizing Flows




Definition (push-forward): if  $x \sim p_0$  then  $T(x) \sim T\#p_0$

Normalizing flow (intuition):

- denoting  $p_{gen} = T\#p_0$
- locally, if  $T$  compresses the space by a factor 42, then  $p_{gen}(T(x)) = 42 \cdot p_0(x)$
- formally, change of variable,  $p_{gen}(T(x)) = |\det(J_{T^{-1}}(x))| \cdot p_0(x)$  (determinant of the jacobian of  $T^{-1}$ )

Principle: parametrize and learn  $T$  ... so that its inverse exists (and has an easy jacobian det).

# Normalizing Flows, with Composed Functions



Learn a deep  $T$ , i.e.,


$$T = \phi_1 \circ \phi_2 \circ \dots \circ \phi_K$$

Chain rule of change of variable,

$$|\det(J_{T^{-1}}(x))| = \prod_k |\det(J_{\phi_k^{-1}}(x))|$$

Principle: compose invertible blocks (with easy jacobian det)

# Normalizing Flows, with Composed Functions



Learn a deep  $T$ , i.e.,


$$T = \phi_1 \circ \phi_2 \circ \dots \circ \phi_K$$

Chain rule of change of variable,

$$|\det(J_{T^{-1}}(x))| = \prod_k |\det(J_{\phi_k^{-1}}(x))|$$

Principle: compose invertible blocks (with easy jacobian det)

# Normalizing Flows, with Composed Functions



Learn a deep  $T$ , i.e.,


$$T = \phi_1 \circ \phi_2 \circ \dots \circ \phi_K$$

Chain rule of change of variable,

$$|\det(J_{T^{-1}}(x))| = \prod_k |\det(J_{\phi_k^{-1}}(x))|$$

Principle: compose invertible blocks (with easy jacobian det)

# Normalizing Flows, with Composed Functions



Learn a deep  $T$ , i.e.,

$$T = \phi_1 \circ \phi_2 \circ \dots \circ \phi_K$$


Chain rule of change of variable,

$$|\det(J_{T^{-1}}(x))| = \prod_k |\det(J_{\phi_k^{-1}}(x))|$$


Principle: compose invertible blocks (with easy jacobian det)

# Continuous Normalizing Flows


# Continuous Normalizing Flows




# Continuous Normalizing Flows



# Continuous Normalizing Flows




# Continuous Normalizing Flows



Pushing to the limit

- Infinitely many infinitely-small steps
- Making depth continuous  $k \mapsto t$
- Replacing  $\phi_k(x)$  by  $u(x, t)$ , or  $u_t(x)$

# Continuous Normalizing Flows




Pushing to the limit

Continuous Normalizing Flow


- Infinitely many infinitely-small steps
- Making depth continuous  $k \mapsto t$
- Replacing  $\phi_k(x)$  by  $u(x, t)$ , or  $u_t(x)$
- easier: less constraints on  $u$  than  $\phi$

$$\min_{\theta} \text{Dist}(T_{\theta} \# p_0, p_{\text{data}})$$

# Continuous Normalizing Flows: Visual Summary



# Continuous Normalizing Flows: "limitation"




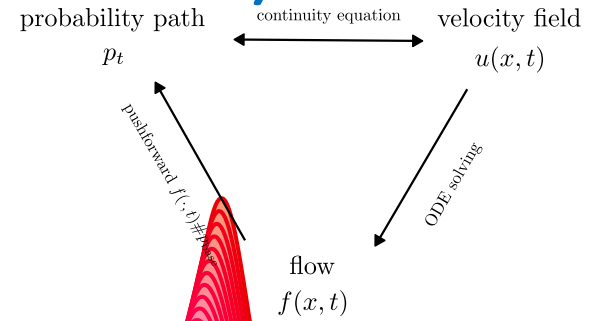
The flow is unspecified!  
(there is an infinity of equally good solutions)

# Probability paths, velocity fields (and flows)






Continuity Equation / Transport Equation  

$$\partial_t p_t + \nabla \cdot u_t p_t = 0$$





# Conditional Flow Matching (CFM)



probability path

continuity equation

velocity field

$p_t$

$u(x, t)$


pushforward  $f(\cdot, t) \# p_{\text{base}}$

ODE solving


flow  
 $f(x, t)$

# How to Learn a Good Velocity $u_\theta$ ?


# How to Learn a Good Velocity $u_\theta$ ?



t=0



t=0.25



t=0.5




t=0.75




t=1


# How to Learn a Good Velocity $u_\theta$ ?



t=0



t=0.25



t=0.5



t=0.75




t=1


**Idea** : learn  $u_\theta$

$$u_\theta(tx_1 + (1-t)x_0, t) \approx x_1 - x_0$$


# How to Learn a Good Velocity $u_\theta$ ?




t=0



t=0.25



t=0.5



t=0.75



t=1

Idea : learn  $u_\theta$

$$u_\theta(tx_1 + (1-t)x_0, t) \approx x_1 - x_0$$

Diffusion/Conditional Flow Matching<sup>15,16,17</sup>

$$\mathbb{E}_{\substack{x_0 \sim p_0 \\ x_1 \sim p_{\text{data}} \\ t \sim \mathcal{U}([0,1])}} \left\| u_\theta(tx_1 + (1-t)x_0, t) - (x_1 - x_0) \right\|^2$$

<sup>15</sup> Y. Lipman et al., **Flow Matching for Generative Modeling**, In: ICLR, 2023.

<sup>16</sup> M. Albergo et al., **Building Normalizing Flows with Stochastic Interpolants**, In: ICLR, 2023.

<sup>17</sup> X. Liu et al., **Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow**, In: ICLR, 2023.

# Conditional Flow Matching Algorithm

# Conditional Flow Matching Algorithm

$$x_0 \sim \mathcal{N}(0, I)$$

$$x_1 \sim p_{\text{data}}$$

$$t \sim \text{Uniform}([0, 1])$$

$$x_t = t \cdot x_1 + (1 - t) \cdot x_0$$

SGD step on  $\theta$  with loss:  $\|u_\theta(x, t) - (x_1 - x_0)\|_2^2$

# Conditional Flow Matching Algorithm

$$x_0 \sim \mathcal{N}(0, I)$$

$$x_1 \sim p_{\text{data}}$$

$$t \sim \text{Uniform}([0, 1])$$

$$x_t = t \cdot x_1 + (1 - t) \cdot x_0$$

SGD step on  $\theta$  with loss:  $\|u_\theta(x, t) - (x_1 - x_0)\|_2^2$

**That's it!**

# Conditional Flow Matching Algorithm

$$x_0 \sim \mathcal{N}(0, I)$$

$$x_1 \sim p_{\text{data}}$$

$$t \sim \text{Uniform}([0, 1])$$


$$x_t = t \cdot x_1 + (1 - t) \cdot x_0$$

SGD step on  $\theta$  with loss:  $\|u_\theta(x, t) - (x_1 - x_0)\|_2^2$

**That's it!**

**(up to practical hacks and a few days of training)**

# CFM: Does it works? the "inversion", path un-mixing



# CFM: Design choices

## CFM: Design choices

Decide on  $p_0$ , typically  $\mathcal{N}(0, I)$

## CFM: Design choices

Decide on  $p_0$ , typically  $\mathcal{N}(0, I)$

Decide on the conditioning variable (and its distribution), e.g.

# CFM: Design choices

Decide on  $p_0$ , typically  $\mathcal{N}(0, I)$

Decide on the conditioning variable (and its distribution), e.g.

- $z$  is a pair  $(x_0, x_1)$
- $z$  is a target point  $x_1$
- $z$  is a minibatch of source and target
- $z$  is a pair, constrained by some clusters

# CFM: Design choices

Decide on  $p_0$ , typically  $\mathcal{N}(0, I)$

Decide on the conditioning variable (and its distribution), e.g.

- $z$  is a pair  $(x_0, x_1)$
- $z$  is a target point  $x_1$
- $z$  is a minibatch of source and target
- $z$  is a pair, constrained by some clusters

Decide on the conditional "flow"

# CFM: Design choices

Decide on  $p_0$ , typically  $\mathcal{N}(0, I)$

Decide on the conditioning variable (and its distribution), e.g.

- $z$  is a pair  $(x_0, x_1)$
- $z$  is a target point  $x_1$
- $z$  is a minibatch of source and target
- $z$  is a pair, constrained by some clusters

Decide on the conditional "flow"

- conditional probability path  $p_t(x|z)$  (or  $p(x, t|z)$ )
- and associated velocity field  $u^{\text{cond}}(x, t)$

# CFM: Design choices

Decide on  $p_0$ , typically  $\mathcal{N}(0, I)$

Decide on the conditioning variable (and its distribution), e.g.

- $z$  is a pair  $(x_0, x_1)$
- $z$  is a target point  $x_1$
- $z$  is a minibatch of source and target
- $z$  is a pair, constrained by some clusters

Decide on the conditional "flow"

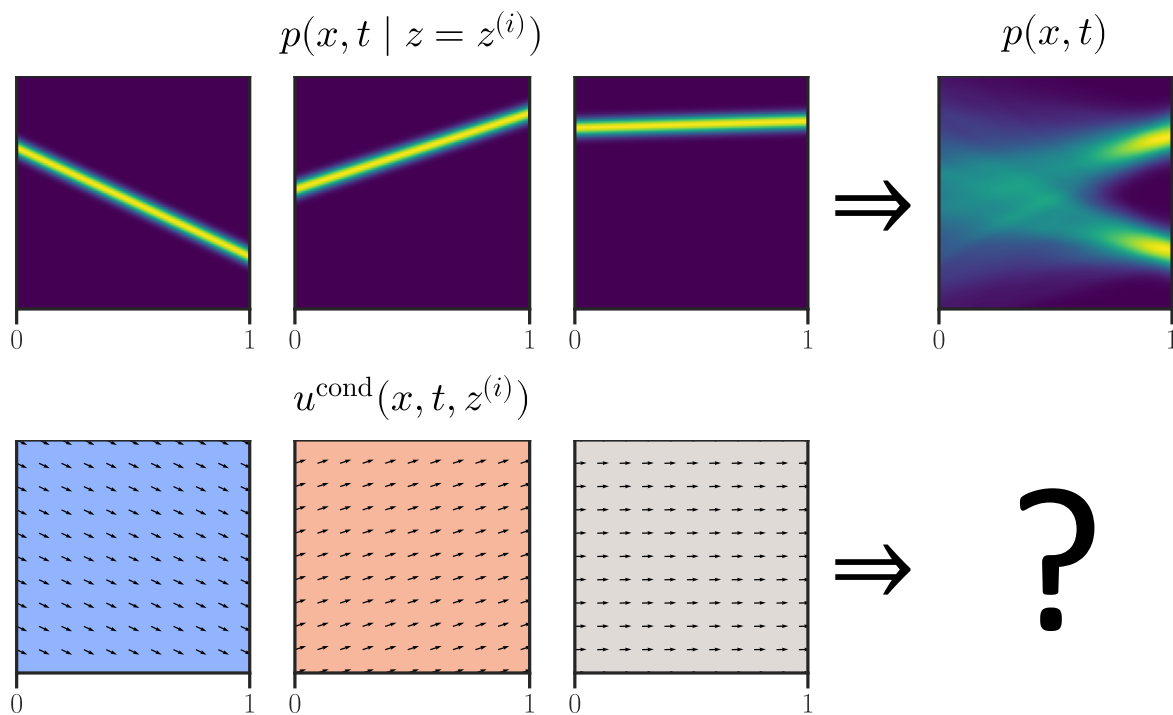
- conditional probability path  $p_t(x|z)$  (or  $p(x, t|z)$ )
- and associated velocity field  $u^{\text{cond}}(x, t)$

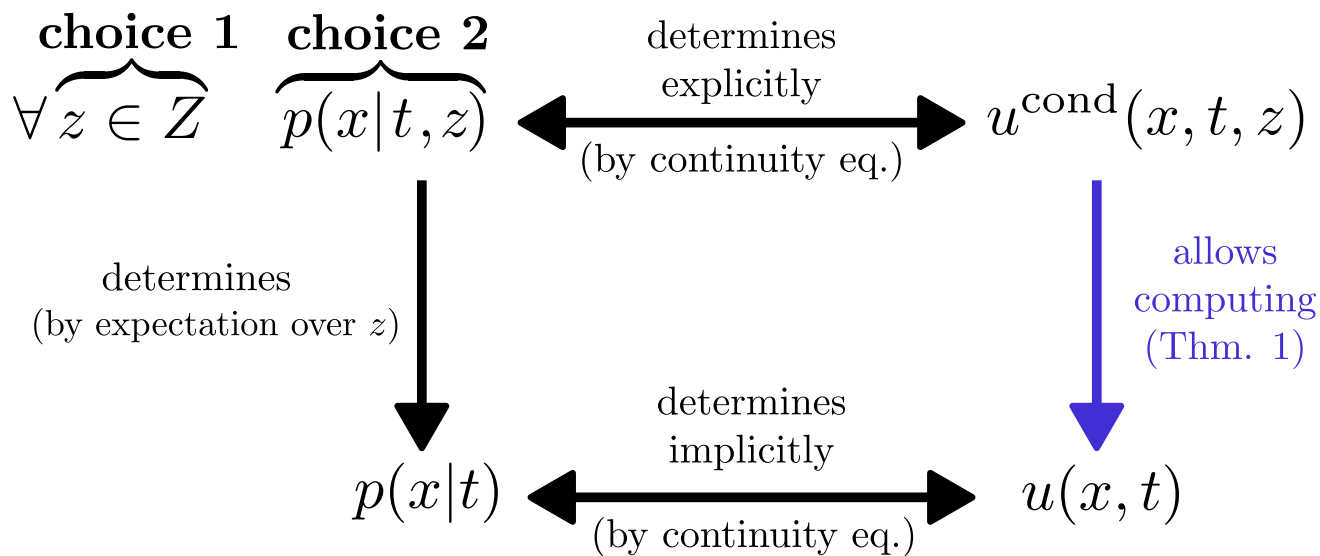
(under marginal constraints, on  $p(x, t)$ )

CFM:  $p(x, t|z)$  (conditional) to  $p(x, t)$  is easy

$$p(x, t) = \int p(x, t|z)p(z)dz = E_z [p(x, t|z)]$$

CFM:  $u^{cond}(x, t, z)$  to  $u(x, t)$  is less easy





# CFM: Closed form expression (Theorem 1)

$$\forall t, \forall \mathbf{x},$$

$$u(\mathbf{x}, t) = E_{z|x,t}[u^{\text{cond}}(\mathbf{x}, t, z)]$$

(also written as)

$$\forall t, \forall \mathbf{x},$$

$$u(\mathbf{x}, t) = \int_z u^{\text{cond}}(\mathbf{x}, t, z)p(z|x, t)$$

(or Bayes)

$$\forall t, \forall \mathbf{x},$$

$$u(\mathbf{x}, t) = \int_z u^{\text{cond}}(\mathbf{x}, t, z) \frac{p(\mathbf{x}, t|z)p(z)}{p(\mathbf{x}, t)} = E_z \left[ \frac{u^{\text{cond}}(\mathbf{x}, t, z)p(\mathbf{x}, t|z)}{p(\mathbf{x}, t)} \right] = E_z \left[ \frac{u^{\text{cond}}(\mathbf{x}, t, z)p(\mathbf{x}, t|z)}{\sum_{z'} p(\mathbf{x}, t|z')p(z')} \right]$$

# CFM: Closed form expression (Theorem 1)

$$\forall t, \forall \mathbf{x},$$

$$u(\mathbf{x}, t) = E_{z|x,t}[u^{\text{cond}}(\mathbf{x}, t, z)]$$

(also written as)

$$\forall t, \forall \mathbf{x},$$

$$u(\mathbf{x}, t) = \int_z u^{\text{cond}}(\mathbf{x}, t, z)p(z|x, t)$$

(or Bayes)

$$\forall t, \forall \mathbf{x},$$

$$u(\mathbf{x}, t) = \int_z u^{\text{cond}}(\mathbf{x}, t, z) \frac{p(\mathbf{x}, t|z)p(z)}{p(\mathbf{x}, t)} = E_z \left[ \frac{u^{\text{cond}}(\mathbf{x}, t, z)p(\mathbf{x}, t|z)}{p(\mathbf{x}, t)} \right] = E_z \left[ \frac{u^{\text{cond}}(\mathbf{x}, t, z)p(\mathbf{x}, t|z)}{\sum_{z'} p(\mathbf{x}, t|z')p(z')} \right]$$

# CFM: Closed form expression (Theorem 1)

$$\forall t, \forall x,$$

$$u(x, t) = E_{z|x,t}[u^{\text{cond}}(x, t, z)]$$

(also written as)

$$\forall t, \forall x,$$

$$u(x, t) = \int_z u^{\text{cond}}(x, t, z)p(z|x, t)$$

(or Bayes)

$$\forall t, \forall x,$$

$$u(x, t) = \int_z u^{\text{cond}}(x, t, z) \frac{p(x, t|z)p(z)}{p(x, t)} = E_z \left[ \frac{u^{\text{cond}}(x, t, z)p(x, t|z)}{p(x, t)} \right] = E_z \left[ \frac{u^{\text{cond}}(x, t, z)p(x, t|z)}{\sum_{z'} p(x, t|z')p(z')} \right]$$

# CFM playground

R

$p_0$

G  $\rightarrow$  .

U  $\rightarrow$  .

G2  $\rightarrow$  .

U2  $\rightarrow$  .

$p_1$

.  $\rightarrow$  G

.  $\rightarrow$  U

.  $\rightarrow$  G2

.  $\rightarrow$  U2

$z$

-

$\nabla$

$\approx$

$\text{☹}$


$\text{☺}$

$\equiv$

$\otimes$

$\subset$

$\hookrightarrow$



z

1

3

\*

$\approx$

$\approx$

$\text{↺}$


$\times$

1s

.5s

.2s

# Diffusion: denoising diffusion probabilistic models (DDPM)



## Principle


Figure 2: The directed graphical model considered in this work.

- progressively noise you data
- use that data to learn an infinitesimal denoiser

---

1. ... the latent space dimension is the same as the data space dimension (like CNF, contrary to PCA, GAN, VAE) ⇔

# Diffusion: denoising diffusion probabilistic models (DDPM)



## Principle

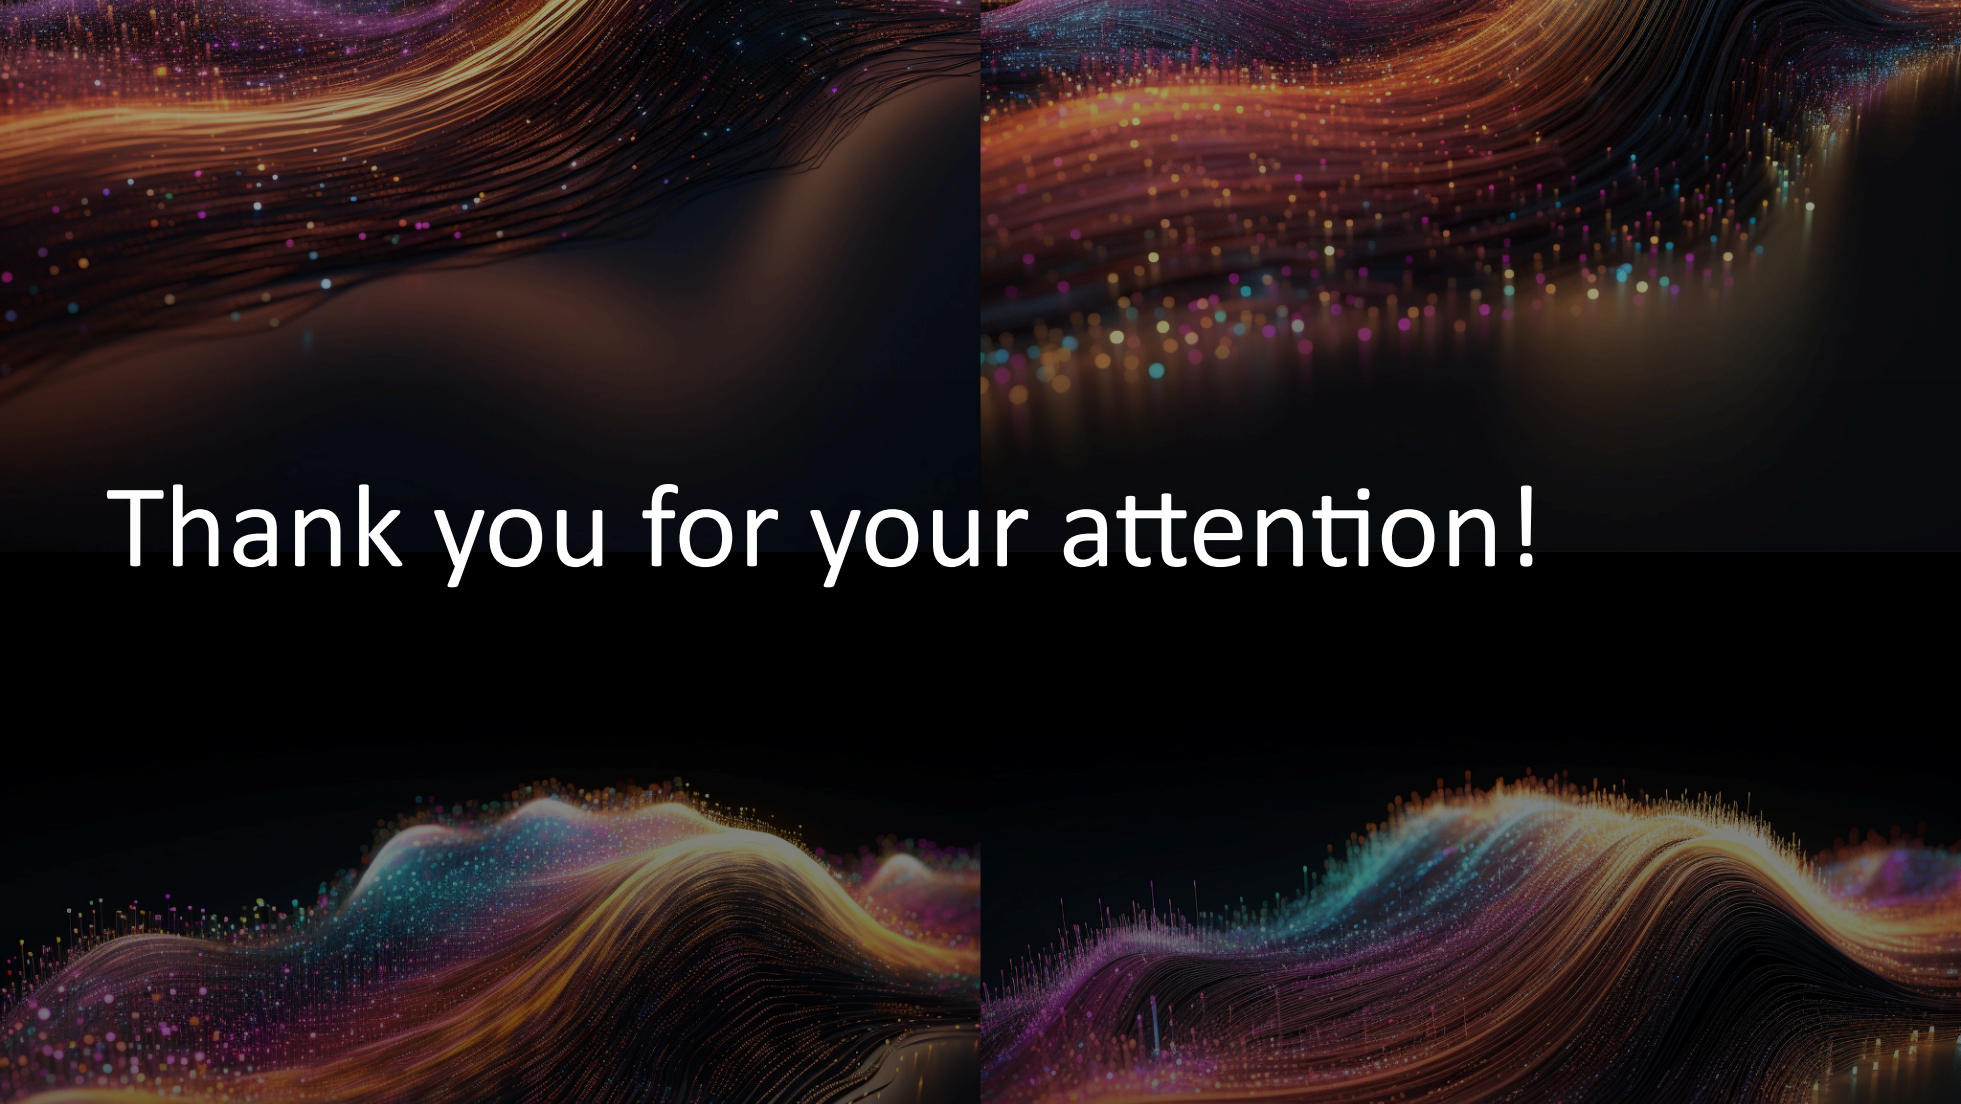
Figure 2: The directed graphical model considered in this work.

- progressively noise you data
- use that data to learn an infinitesimal denoiser

## Actually

- a VAE with successive latent representations<sup>[1]</sup>
- learning a velocity field (notation trap:  $t \in \llbracket T, 0 \rrbracket$  instead of  $t \in [0, 1]$  )
- specifying a unique probability path (but stochastic flow)
- effectively supervising at every step (vs CNF)

1. ... the latent space dimension is the same as the data space dimension (like CNF, contrary to PCA, GAN, VAE) ↔



Thank you for your attention!