


# Generative AI with Stochastic Differential Equations

*An introduction to flow and diffusion models*

**MIT IAP 2025 | Jan 21, 2025**

Peter Holderrieth and Ezra Erives

*Sponsor: Tommi Jaakkola*



# Flow and Diffusion Models: state-of-the-art models for generating images, videos, proteins!




*Stable Diffusion*

*DALEE*



*OpenAI Sora*

*Meta MovieGen*



*AlphaFold3*

*RFDiffusion*

**Most SOTA generative AI models for  
images/videos/proteins/robotics: Diffusion and Flow Models**

## Section 1:

# **From Generation to Sampling**

*Goal:* Formalize what it means to “generate” something.

# We represent images/videos/protein as vectors

## Images:

- Height  $H$  and Width  $W$
- 3 color channels (RGB)

$$z \in \mathbb{R}^{H \times W \times 3}$$



## Videos:

- $T$  time frames
- Each frame is image


$$z \in \mathbb{R}^{T \times H \times W \times 3}$$



## Molecular structures:

- $N$  atoms
- Each atom has 3 coordinate

$$z \in \mathbb{R}^{N \times 3}$$




**We represent the objects we want to generate as vectors:**

$$z \in \mathbb{R}^d$$



# What does it mean to successfully generate something?

Prompt: “A picture of a dog”




Useless   <   Bad   <   Wrong animal   <   Great!

***These are subjective statements - Can we formalize this?***

# Data Distribution: How “likely” are we to find this picture in the internet?

Prompt: “A picture of a dog”



Impossible   <   Rare   <   Unlikely   <   Very likely

***How good an image is  $\sim$  How likely it is under the data distribution***

# Generation as sampling from the data distribution

**Data distribution:** Distribution of objects that we want to generate:

$p_{\text{data}}$


**Probability density:**

$$p_{\text{data}} : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0},$$
$$z \mapsto p_{\text{data}}(z)$$


*Note: We don't know the probability density!*

**Generation means sampling the data distribution:**

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



$$z =$$



# A Dataset consists of samples from the data distribution

**Training requires datasets:** To train our algorithms, we need a **dataset**.

**Examples:**

- Images: Publicly available images from the internet
- Videos: YouTube
- Protein structures: Scientific data (e.g. Protein Data Bank)

**A dataset consists of a finite number of samples from the data distribution:**

$$z_1, \dots, z_N \sim p_{\text{data}}$$





# Conditional Generation allows us to condition on prompts

Data distribution  $p_{\text{data}}$

Fixed prompt


“Dog”



Conditional data distribution  $p_{\text{data}}(\cdot|y)$

Condition variable:  $y$

$y$  = “Dog”



$y$  = “Cat”



$y$  = “Landscape”



**Conditional generation means sampling the conditional data distribution:**

$$z \sim p_{\text{data}}(\cdot|y)$$

*We will first focus on unconditional generation and then learn how to translate an unconditional model to a conditional one.*

Generative Models generate samples from data distribution

Initial distribution:


$p_{\text{init}}$

**Default:**


$$p_{\text{init}} = \mathcal{N}(0, I_d)$$

**A generative model converts samples from a initial distribution (e.g. Gaussian) into samples from the data distribution:**


$$x \sim p_{\text{init}}$$



Generative  
Model



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




## Section 2:

# **Flow and Diffusion Models**

*Goal:* Understand differential equations and how we can build generative models with them.

Flow - Example





# Existence and Uniqueness Theorem ODEs

**Theorem (Picard–Lindelöf theorem):** If the vector field  $u_t(x)$  is continuously differentiable with bounded derivatives, then a *unique* solution to the ODE

$$X_0 = x_0, \quad \frac{d}{dt}X_t = u_t(X_t)$$

exists. In other words, a flow map exists. More generally, this is true if the vector field is **Lipschitz**.

**Key takeaway:** In the cases of practical interest for machine learning, **unique solutions to ODE/flows exist.**

*Math class: Construct solutions via Picard-Iteration*

# Example: Linear ODE

**Simple vector field:**

$$u_t(x) = -\theta x \quad (\theta > 0)$$

**Claim:** Flow is given by

$$\psi_t(x_0) = \exp(-\theta t) x_0$$


**Proof:**

1. Initial condition:

$$\psi_t(x_0) = \exp(0)x_0 = x_0$$

2. ODE:

$$\frac{d}{dt}\psi_t(x_0) = \frac{d}{dt}(\exp(-\theta t)x_0) = -\theta \exp(-\theta t)x_0 = -\theta\psi_t(x_0) = u_t(\psi_t(x_0))$$



# Numerical ODE simulation - Euler method

---

**Algorithm 1** Simulating an ODE with the Euler method

---

**Require:** Vector field  $u_t$ , initial condition  $x_0$ , number of steps  $n$

1: Set  $t = 0$

2: Set step size  $h = \frac{1}{n}$

3: Set  $X_0 = x_0$

4: **for**  $i = 1, \dots, n - 1$  **do**

5:      $X_{t+h} = X_t + hu_t(X_t)$      *Small step into direction of vector field*

6:     Update  $t \leftarrow t + h$

7: **end for**

8: **return**  $X_0, X_h, X_{2h}, \dots, X_1$      *Return trajectory*

---

## Toy example




Figure credit:  
Yaron Lipman

# How to generate objects with a Flow Model

---


## Algorithm 1 Sampling from a Flow Model with Euler method

---

**Require:** Neural network vector field  $u_t^\theta$ , number of steps  $n$

- 1: Set  $t = 0$
  - 2: Set step size  $h = \frac{1}{n}$
  - 3: Draw a sample  $X_0 \sim p_{\text{init}}$  *Random initialization!*
  - 4: **for**  $i = 1, \dots, n - 1$  **do**
  - 5:      $X_{t+h} = X_t + hu_t^\theta(X_t)$
  - 6:     Update  $t \leftarrow t + h$
  - 7: **end for**
  - 8: **return**  $X_1$  *Return final point*
-

# Brownian Motion



# Existence and Uniqueness Theorem SDEs

**Theorem:** If the vector field  $u_t(x)$  is continuously differentiable with bounded derivatives and the diffusion coeff. is continuous, then a *unique* solution to the SDE

$$X_0 = x_0, \quad dX_t = u_t(X_t)dt + \sigma_t dW_t$$

exists. More generally, this is true if the vector field is **Lipschitz**.

**Key takeaway:** In the cases of practical interest for machine learning, **unique solutions to SDE.**

*Stochastic calculus class: Construct solutions via stochastic integrals and Ito-Riemann sums*

# Numerical SDE simulation (Euler-Maruyama method)

---

**Algorithm 2** Sampling from a SDE (Euler-Maruyama method)

---


**Require:** Vector field  $u_t$ , number of steps  $n$ , diffusion coefficient  $\sigma_t$

- 1: Set  $t = 0$
  - 2: Set step size  $h = \frac{1}{n}$
  - 3: Set  $X_0 = x_0$
  - 4: **for**  $i = 1, \dots, n - 1$  **do**
  - 5:     Draw a sample  $\epsilon \sim \mathcal{N}(0, I_d)$
  - 6:      $X_{t+h} = X_t + hu_t(X_t) + \sigma_t\sqrt{h}\epsilon$  *Add additional noise with var=h scaled by diffusion coefficient  $\sigma_t$*
  - 7:     Update  $t \leftarrow t + h$
  - 8: **end for**
  - 9: **return**  $X_0, X_h, X_{2h}, X_{3h}, \dots, X_1$
-



# Ornstein-Uhlenbeck Process

$$dX_t = -\theta X_t dt + \sigma dW_t$$



*Increasing diffusion coefficient*  $\sigma$

---

**Algorithm 2** Sampling from a Diffusion Model (Euler-Maruyama method)

---

**Require:** Neural network  $u_t^\theta$ , number of steps  $n$ , diffusion coefficient  $\sigma_t$

- 1: Set  $t = 0$
  - 2: Set step size  $h = \frac{1}{n}$
  - 3: Draw a sample  $X_0 \sim p_{\text{init}}$
  - 4: **for**  $i = 1, \dots, n - 1$  **do**
  - 5:     Draw a sample  $\epsilon \sim \mathcal{N}(0, I_d)$
  - 6:      $X_{t+h} = X_t + hu_t^\theta(X_t) + \sigma_t\sqrt{h}\epsilon$
  - 7:     Update  $t \leftarrow t + h$
  - 8: **end for**
  - 9: **return**  $X_1$
-

# Reminder: Flow and Diffusion Models

**Flow  
Model**

Initialize:

$$X_0 \sim p_{\text{init}},$$

*E.g. Gaussian*

ODE:

$$dX_t = u_t^\theta(X_t)dt$$

*Neural network  
vector field*

*Diffusion coeff.*

**Diffusion  
Model**

Initialize:

$$X_0 \sim p_{\text{init}},$$

SDE:

$$dX_t = u_t^\theta(X_t)dt + \sigma_t dW_t$$

**To get samples, simulate ODE/SDE from  $t=0$  to  $t=1$  and return  $X_1$**

# Next step: Training the model

Without training, the model produces “non-sense” → **We need to train**  $u_t^\theta$

---

**Training = Finding parameters  $\theta$  such that**

$$X_0 \sim p_{\text{init}}, \quad dX_t = u_t^\theta(X_t)dt \quad \xrightarrow{\text{Implies}} \quad X_1 \sim p_{\text{data}}$$

*Start with initial  
distribution*

*Follow along  
the vector field*

*The distribution of  
the final point = data  
distribution*

**Goal of lecture 2 (today) and lecture 3 (tomorrow):**

**Derive training algorithm**

# Today's goal: Derive a Training Target

- Typically, we train the model by minimizing a **mean squared error**:

$$L(\theta) = \|u_t^\theta(x) - u_t^{\text{target}}(x)\|^2$$

**Training target**

- In regression or classification, the training target is the label.
- Here: No label :(  $\rightarrow$  We have to **derive a training target**

**Today: Derive a formula for the training target:**  $u_t^{\text{target}}(x)$

**Tomorrow: Training algorithm using**  $u_t^{\text{target}}(x)$

## Section 2:

# Constructing a Training Target



*Goal:* Derive a formula for a training target for training our models

Key terminology:

“Conditional” = “Per single data point”

“Marginal” = “Across distribution of data points”

# Probability Paths: The Path from Noise to Data





# Conditional Prob. Path

Notation

Key property

Gaussian example

Conditional  
Probability Path

$$p_t(\cdot|z)$$

Interpolates  $p_{\text{init}}$   
and a data point  $z$

$$\mathcal{N}(\alpha_t z, \beta_t^2 I_d)$$

Conditional  
Vector Field

Conditional  
Score  
Function

# Marginal Prob. Path

	Notation	Key property	Formula
Marginal Probability Path	$p_t$	Interpolates $p_{\text{init}}$ and $p_{\text{data}}$	$\int p_t(x z)p_{\text{data}}(z)\mathrm{d}z$

Marginal  
Vector  
Field


;

Marginal  
Score  
Function

;


$p_{\text{init}}$



Conditional Probability Path  $p_t(\cdot|z)$



$p_{\text{init}}$


Conditional Probability Path  $p_t(\cdot|z)$




$p_{\text{init}}$ Conditional Probability Path  $p_t(\cdot|z)$  $p_{\text{init}}$ Marginal Probability Path  $p_t$  $p_{\text{data}}$

$p_t(\cdot|z)$


Ground truth






ODE samples



ODE Trajectories



$p_t$



# Conditional Prob. Path, Vector Field, and Score

	Notation	Key property	Gaussian example
Conditional Probability Path	$p_t(\cdot z)$	Interpolates $p_{\text{init}}$ and a data point $z$	$\mathcal{N}(\alpha_t z, \beta_t^2 I_d)$
Conditional Vector Field	$u_t^{\text{target}}(x z)$	ODE follows conditional path	$\left(\dot{\alpha}_t - \frac{\dot{\beta}_t}{\beta_t} \alpha_t\right) z + \frac{\dot{\beta}_t}{\beta_t} x$
Conditional Score Function			

# Gaussian Conditional Probability Path And Conditional Vector Field

*Figure credit:  
Yaron Lipman*

