Generative AI with Stochastic Differential Equations

An introduction to flow and diffusion models

MIT IAP 2026 | Jan 20, 2025

Peter Holderrieth and Ron Shprints

Sponsor: Tommi Jaakkola

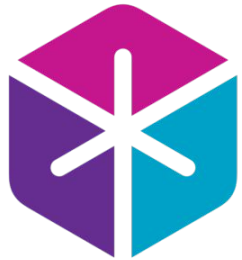

Welcome to class 6.S184!

Course Instructors

Peter

Ron

Sponsor

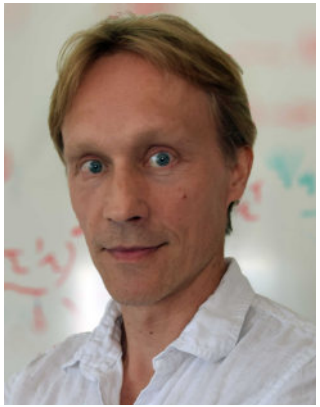

Tommi

Generative AI - A new generation of AI systems

Artistic Images

Realistic Videos

ChatGPT

Draft Texts

These systems are “creative”: they generate new objects.

This class teaches you algorithms to generate objects.

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

Stable Diffusion

DALEE

OpenAI Sora

Meta MovieGen

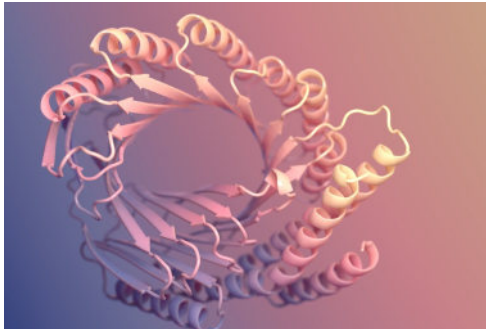

AlphaFold3

Boltz

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

This Class: Theory and Practice of Flow/Diffusion Models

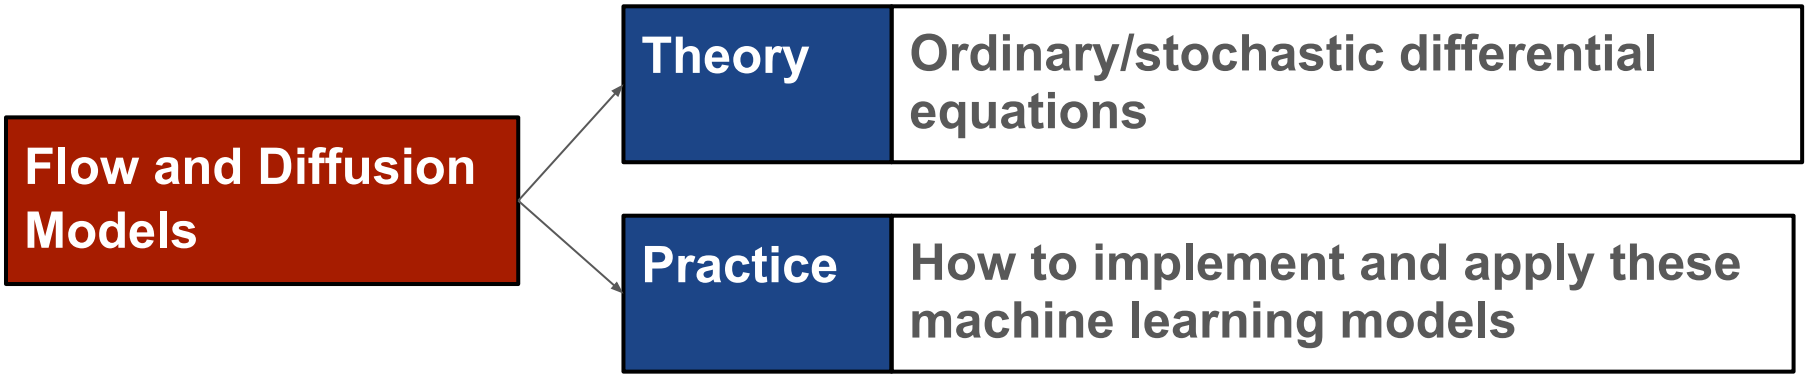

The goal of this class to teach you:

1. Flow and diffusion models from *first principles.*
2. The *minimal* but necessary amount of mathematics for 1.
3. How to implement and apply these algorithms.

Class Overview

- **Lecture 1 (today):**
  - From Generation to Sampling: Formalize “generating an image/etc.”
  - Construct Flow and Diffusion Models
- **Lecture 2 - Flow Matching:** Training algorithm.
- **Lecture 3 - Score Matching, Guidance:** How to condition on a prompt.
- **Lecture 4 - Build Image Generators:** Network architectures + Latent spaces
- **Lecture 5 - Advanced Topics:** Discrete diffusion models, distilled 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 (RBG)

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

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

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

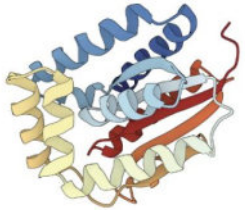

We represent the objects we want to generate as vectors:

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

What does it mean to successfully generate something?

Prompt: “A picture of a dog”

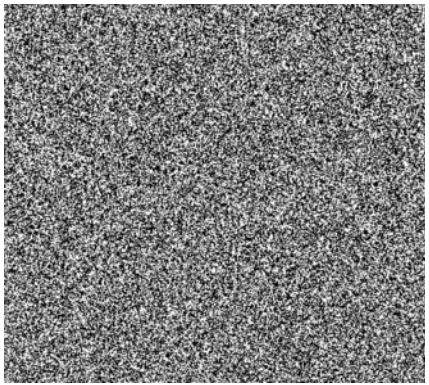

Useless < Bad < Wrong animal < Great!

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”

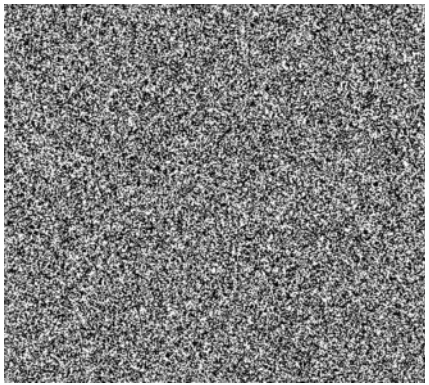

Impossible < **Rare** < **Unlikely** < **Very likely**

How good an image is ~= 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_{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:

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

Conditional Generation allows us to condition on prompts

Data distribution p_{data}

Fixed prompt

"Dog"

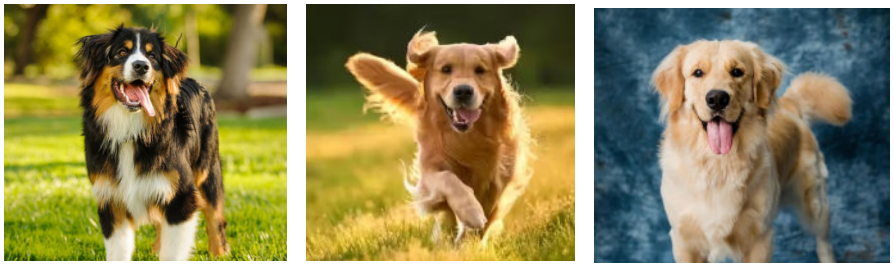

Conditional data distribution $p_{\text{data}}(\cdot|y)$

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_{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 → $z \sim p_{\text{data}}$

Summary

- **Objects to Generate:** We focus on vectors z representing data objects (e.g., images, videos)
- **Data distribution:** Distribution that places higher probability to objects that we consider “good”.
- **Generation as sampling:** generate an object = sampling from the data distribution
- **Dataset**: Finite number of samples from the data distribution used for training
- **Conditional Generation:** Condition on label y and sample from the conditional data distribution
- **Generative Model:** Train a model to transform samples from a simple (e.g., Gaussian) distribution into the data distribution.

Section 2:

Flow and Diffusion Models

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

ODE Trajectory - Example

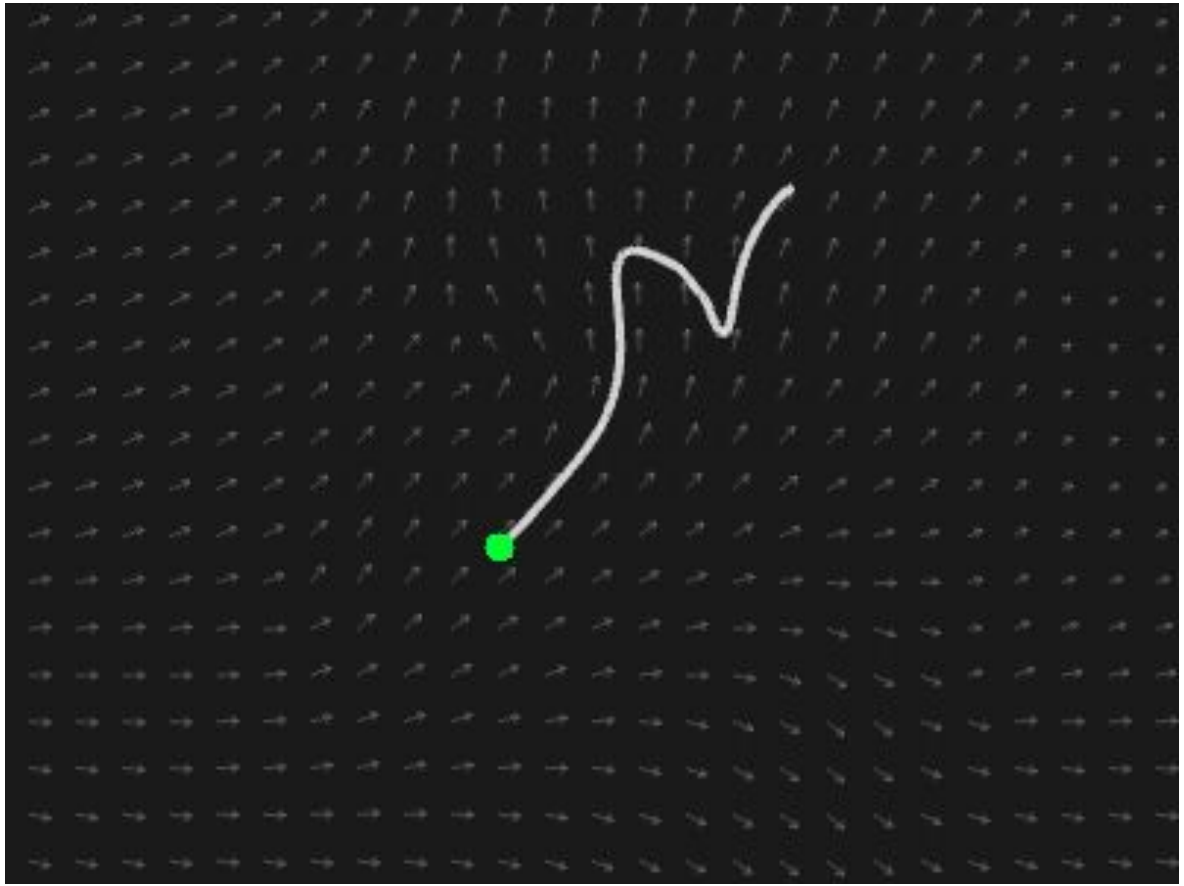

Source: https://mariogemoll.com/flow-matching

Flow - Example

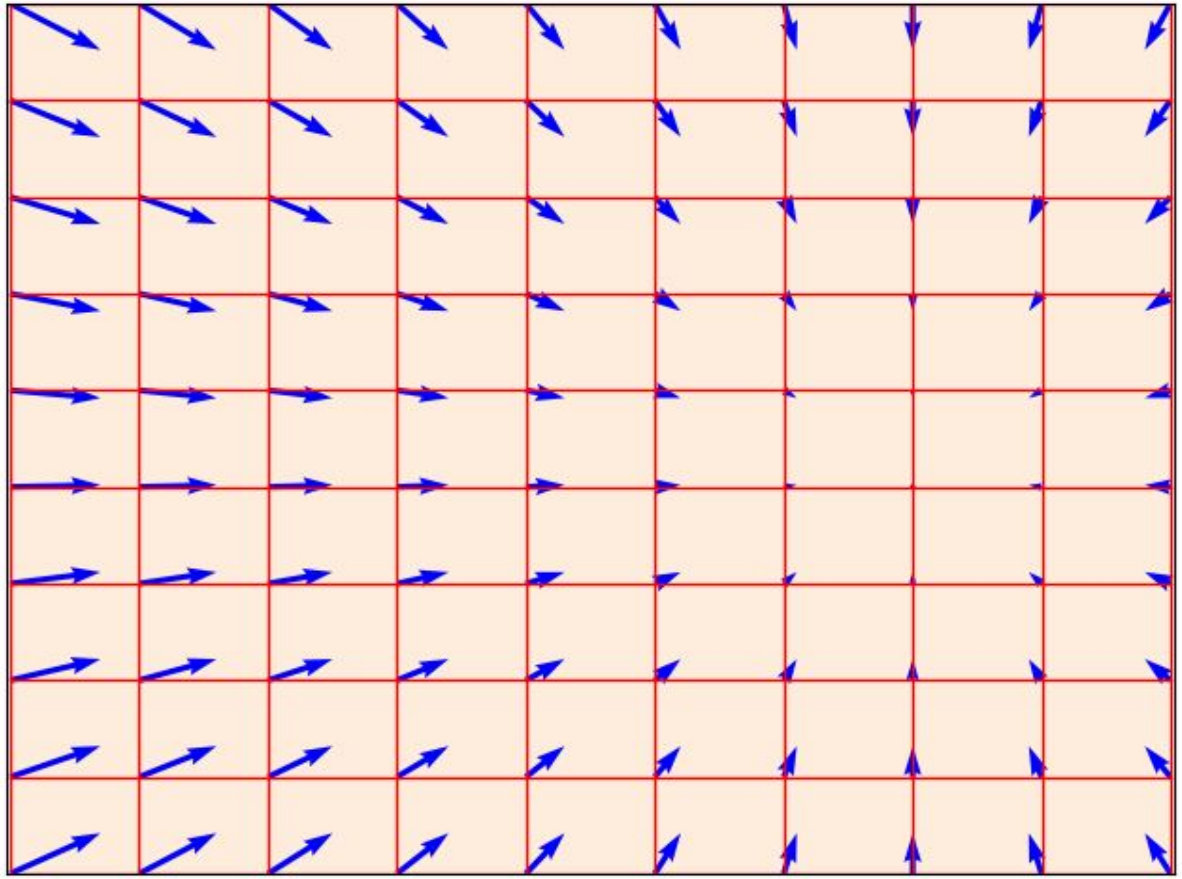

Existence and Uniqueness Theorem ODEs

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

$$X_0 = x_0, \qquad \frac{\mathrm{d}}{\mathrm{d}t} X_t = u_t(X_t)$$

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:

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

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:

$$\frac{\mathrm{d}}{\mathrm{d}t}\psi_t(x_0) = \frac{\mathrm{d}}{\mathrm{d}t}\left(\exp(-\theta t)\, x_0\right) = -\theta \exp(-\theta t)\, x_0 = -\theta \psi_t(x_0) = u_t(\psi_t(x_0))$$

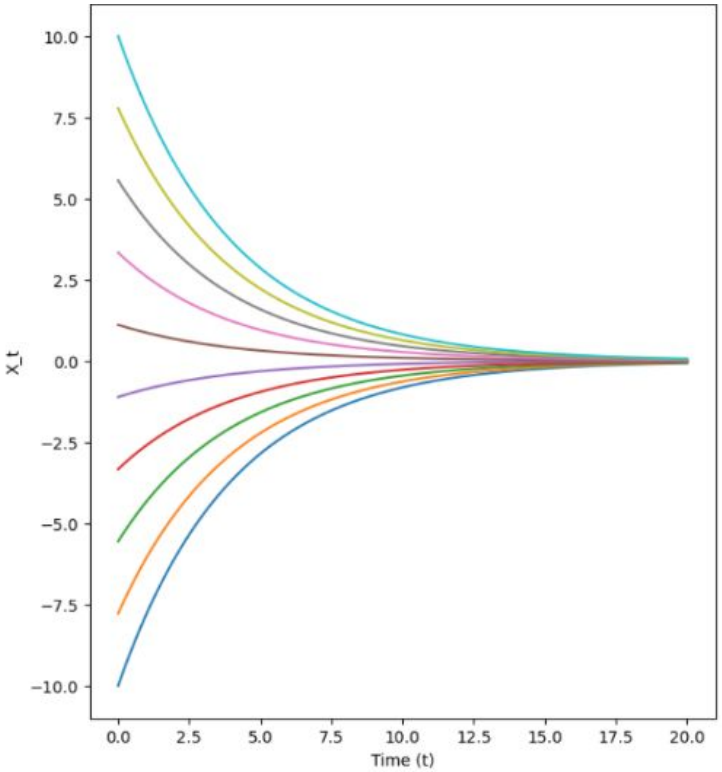

Numerical ODE simulation - Euler method

Algorithm 1 Simulating an ODE with the Euler method

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

Set $t = 0$
Set step size $h = \frac{1}{n}$
Set $X_0 = x_0$
for $i = 1, \ldots, n-1$ **do**
    $X_{t+h} = X_t + hu_t(X_t)$ *Small step into direction of vector field*
    Update $t \leftarrow t + h$
end for
return $X_0, X_h, X_{2h}, \ldots, X_1$ *Return trajectory*

Euler Method - Example

Large step size

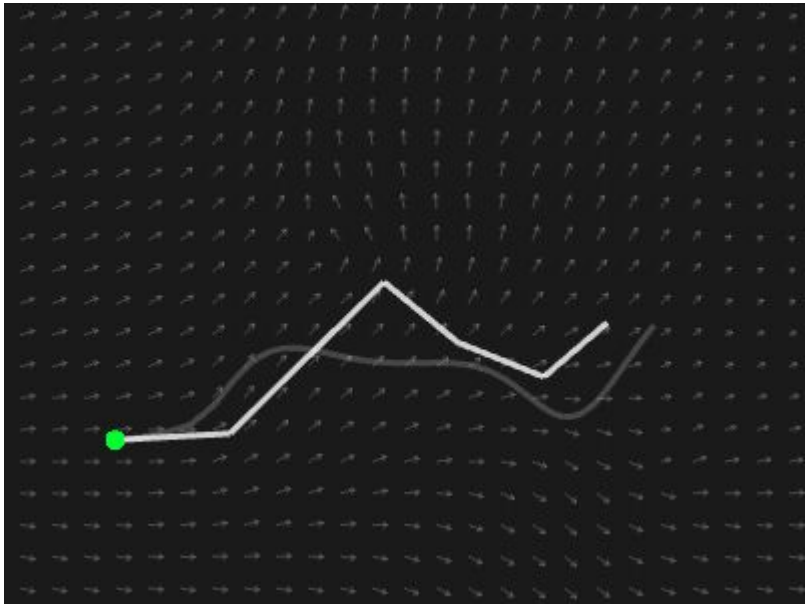

Small step size

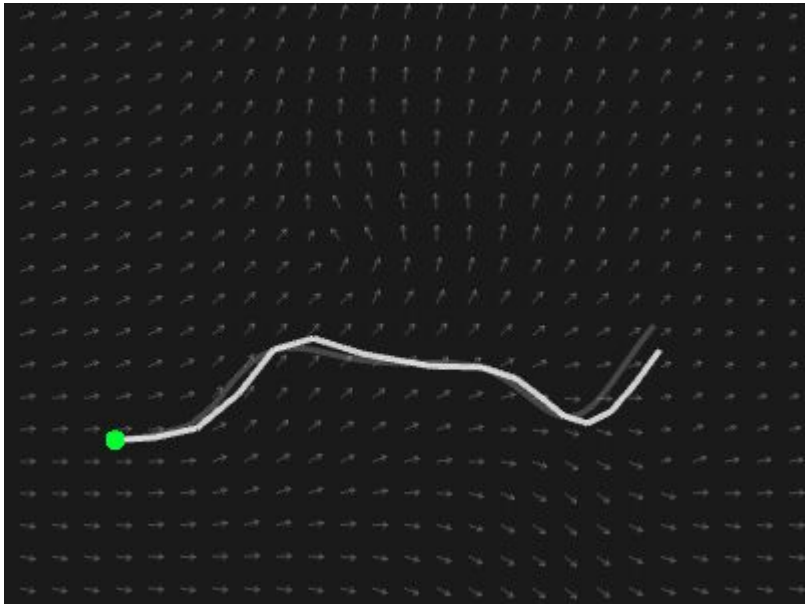

More efficient but *higher error*

Lower error but *less efficient*

Source: https://mariogemoll.com/flow-matching

Toy example

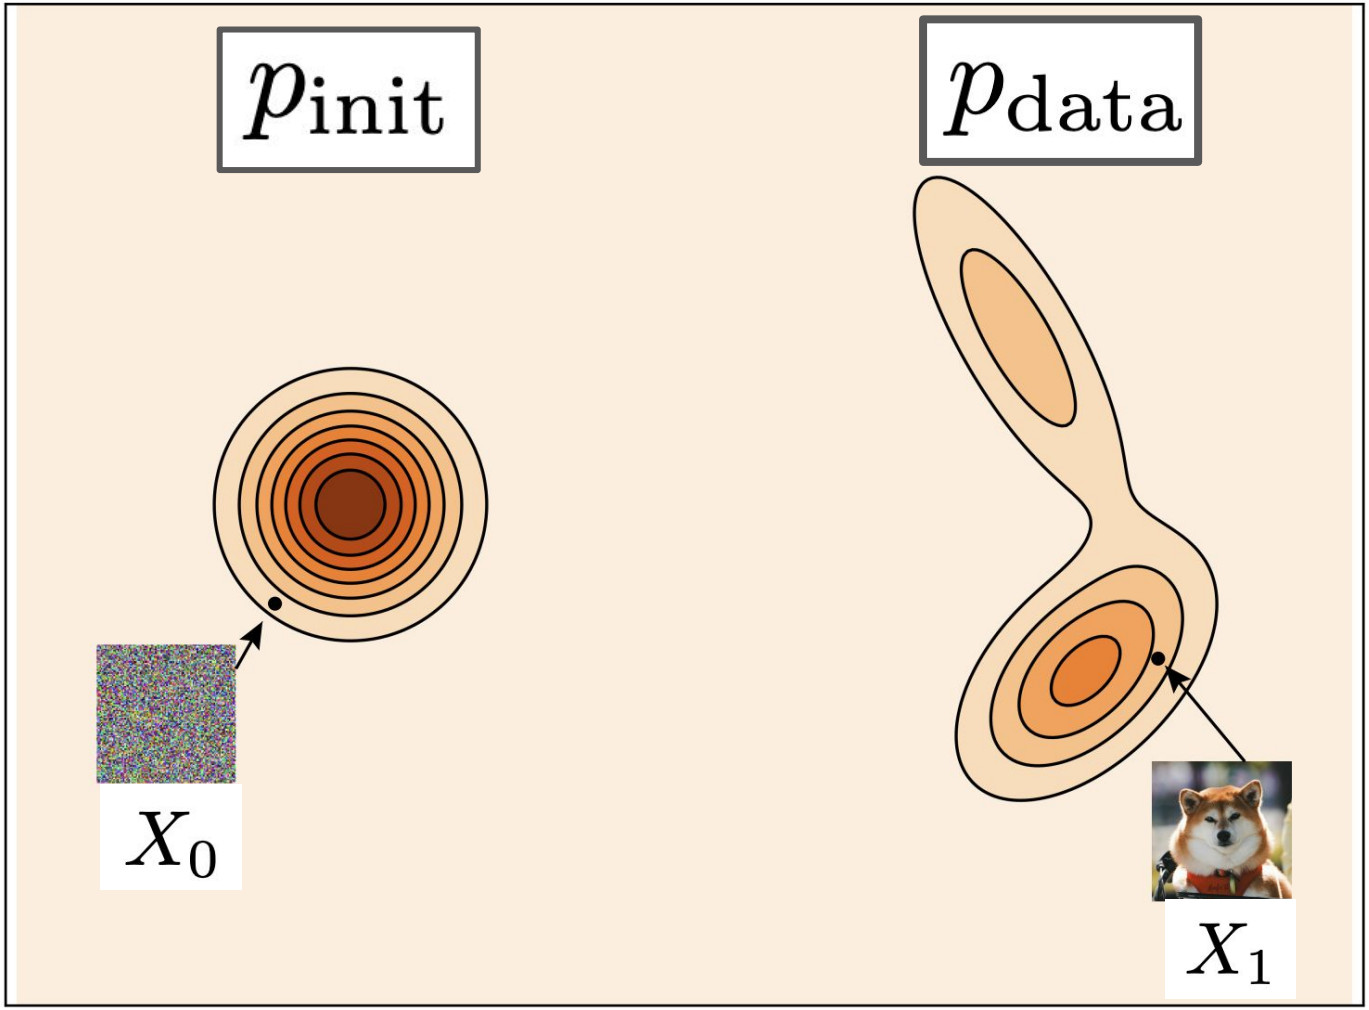

Figure credit: Yaron Lipman

Toy Flow Model

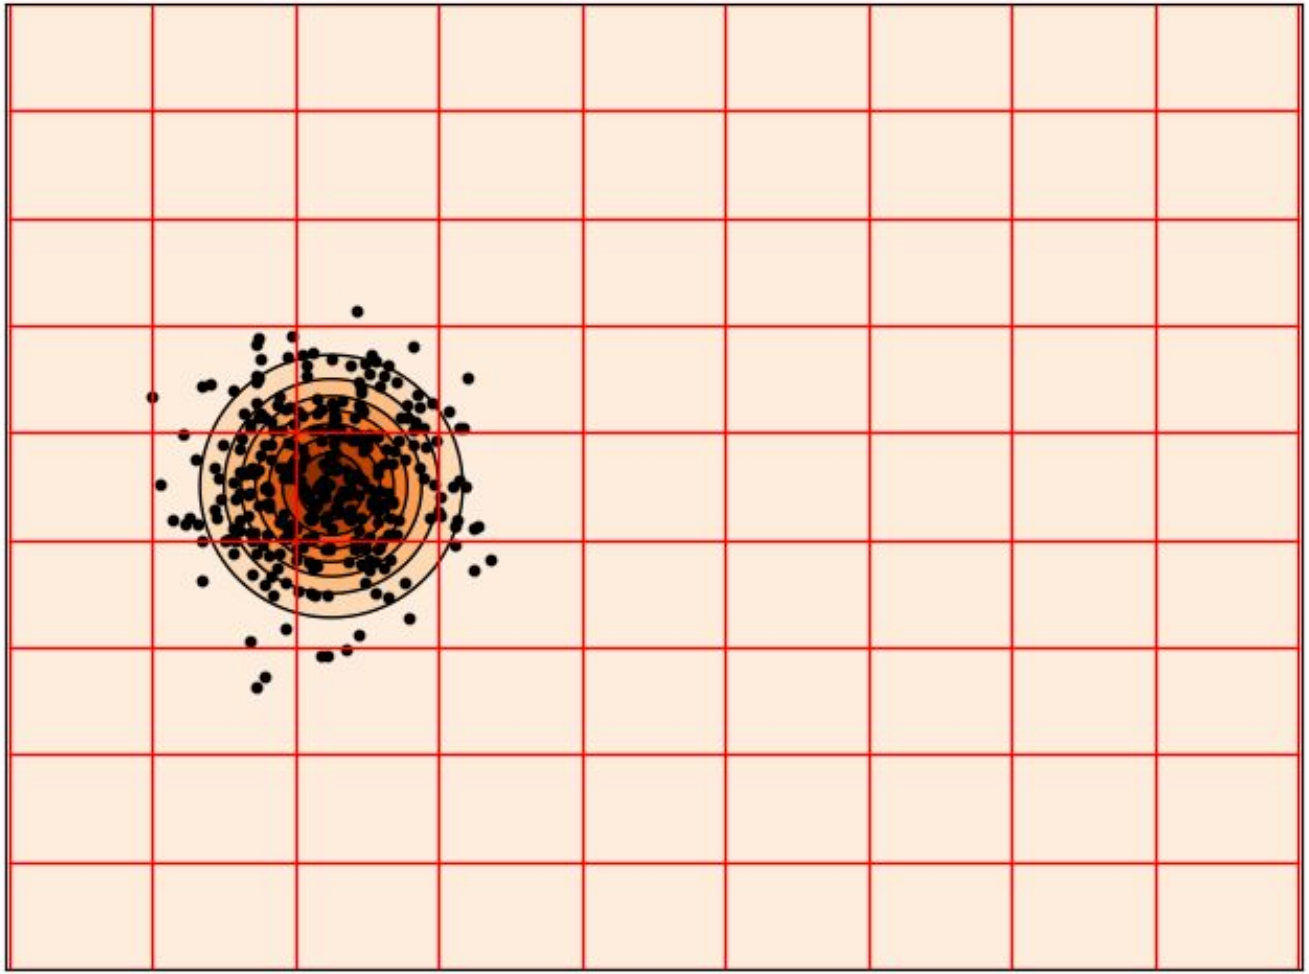

Figure credit: Yaron Lipman

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^θ, 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** $i = 1, \ldots, n-1$ **do**
5: $\quad X_{t+h} = X_t + h u_t^\theta(X_t)$
6: $\quad$ Update $t \leftarrow t + h$
7: **end for**
8: **return** X_1 *Return final point*

Example - Videos Generated via Flow Model Sampling

Stochastic Process: Random Trajectories

Sample 1

Sample 2

Sample 3

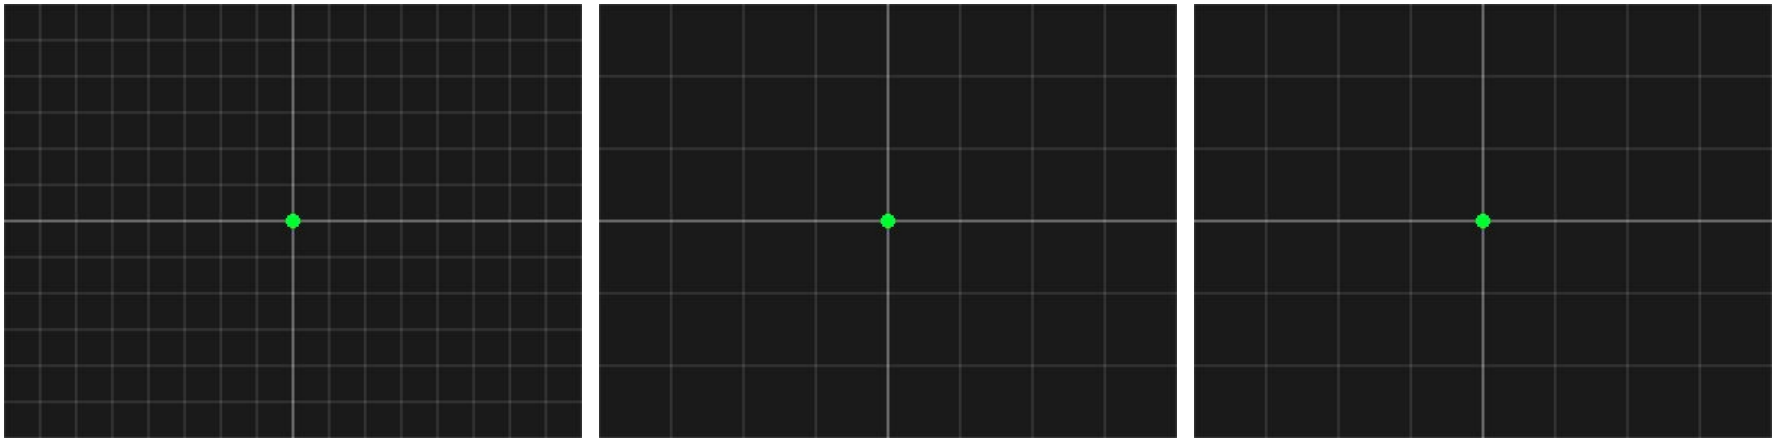

Note: A single stochastic process can give rise to many trajectories as the evolution becomes random.

Source: https://mariogemoll.com/flow-matching

Brownian Motion: Random Trajectories

Sample 1

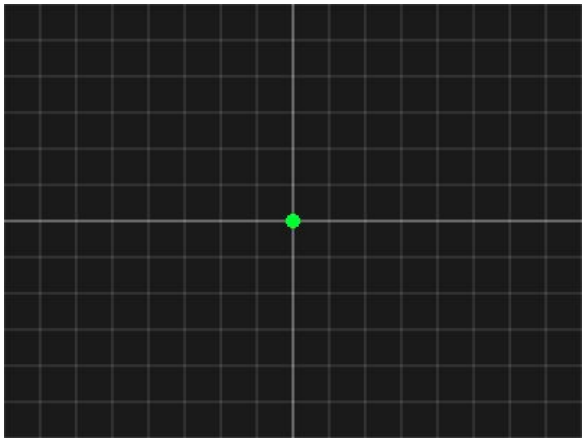

Sample 2

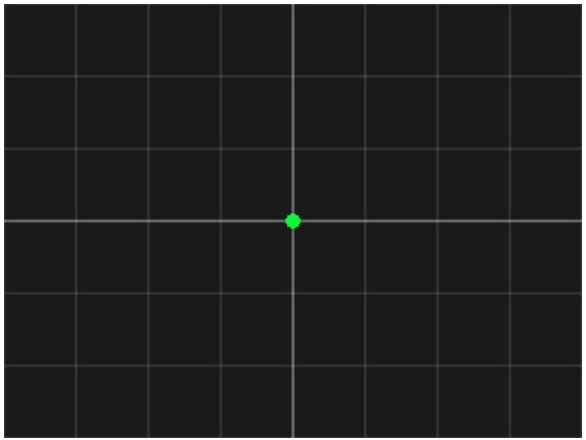

Sample 3

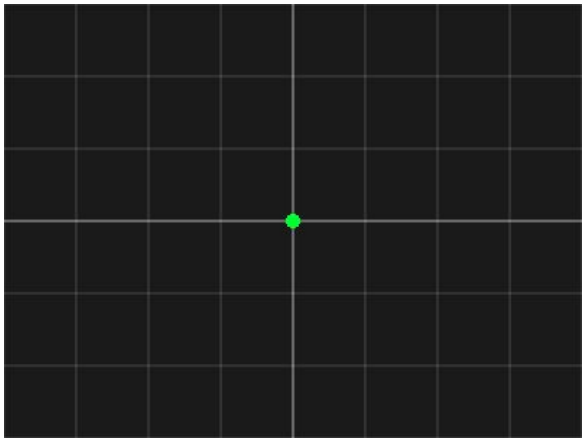

A Brownian motion behaves like a continuous random walk.
Here: 2-dimensional.

Source: https://mariogemoll.com/flow-matching

Brownian Motion

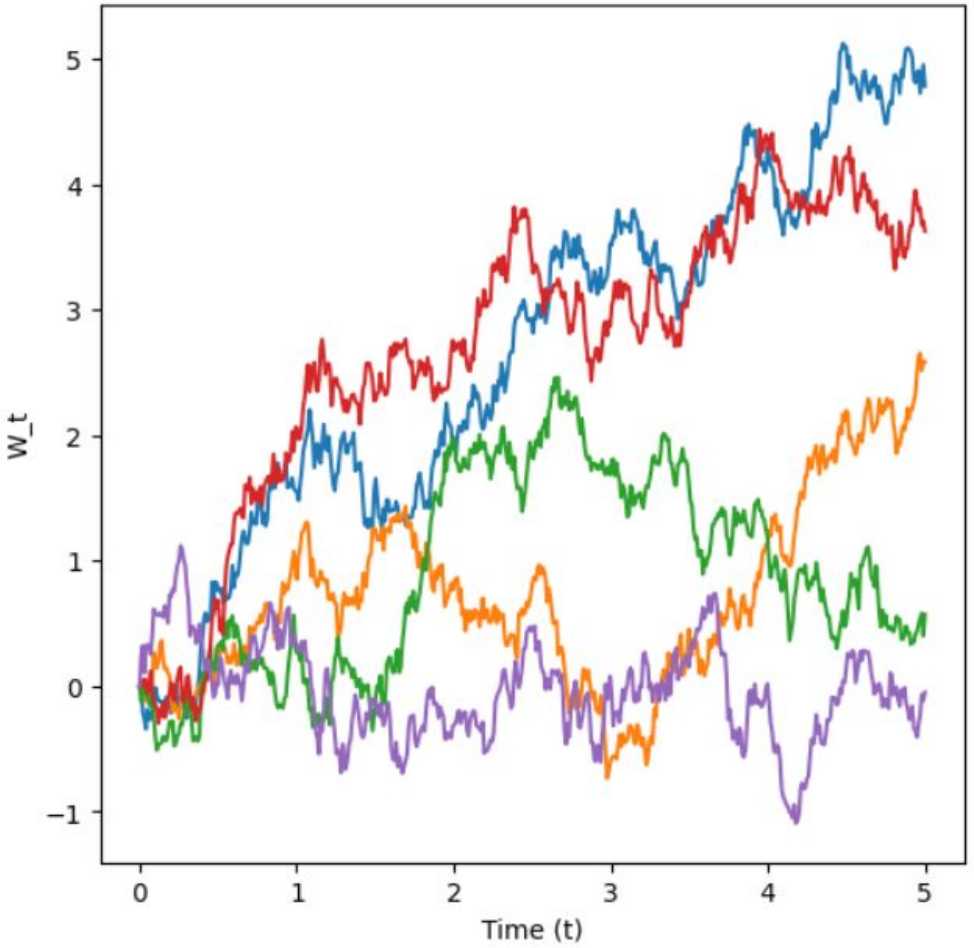

Existence and Uniqueness Theorem SDEs

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

$$X_0 = x_0, \quad \mathrm{d}X_t = u_t(X_t)\mathrm{d}t + \sigma_t \mathrm{d}W_t$$

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 SDEs exist**.

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 σ_t

Set $t = 0$
Set step size $h = \frac{1}{n}$
Set $X_0 = x_0$
for $i = 1, \ldots, n-1$ **do**
  Draw a sample $\epsilon \sim \mathcal{N}(0, I_d)$
  $X_{t+h} = X_t + h u_t(X_t) + \sigma_t \sqrt{h} \epsilon$
  Update $t \leftarrow t + h$
end for
return $X_0, X_h, X_{2h}, X_{3h}, \ldots, X_1$

Add additional noise with var=h scaled by diffusion coefficient σ_t

Ornstein-Uhlenbeck Process

$$\mathrm{d}X_t = -\theta X_t \mathrm{d}t + \sigma dW_t$$

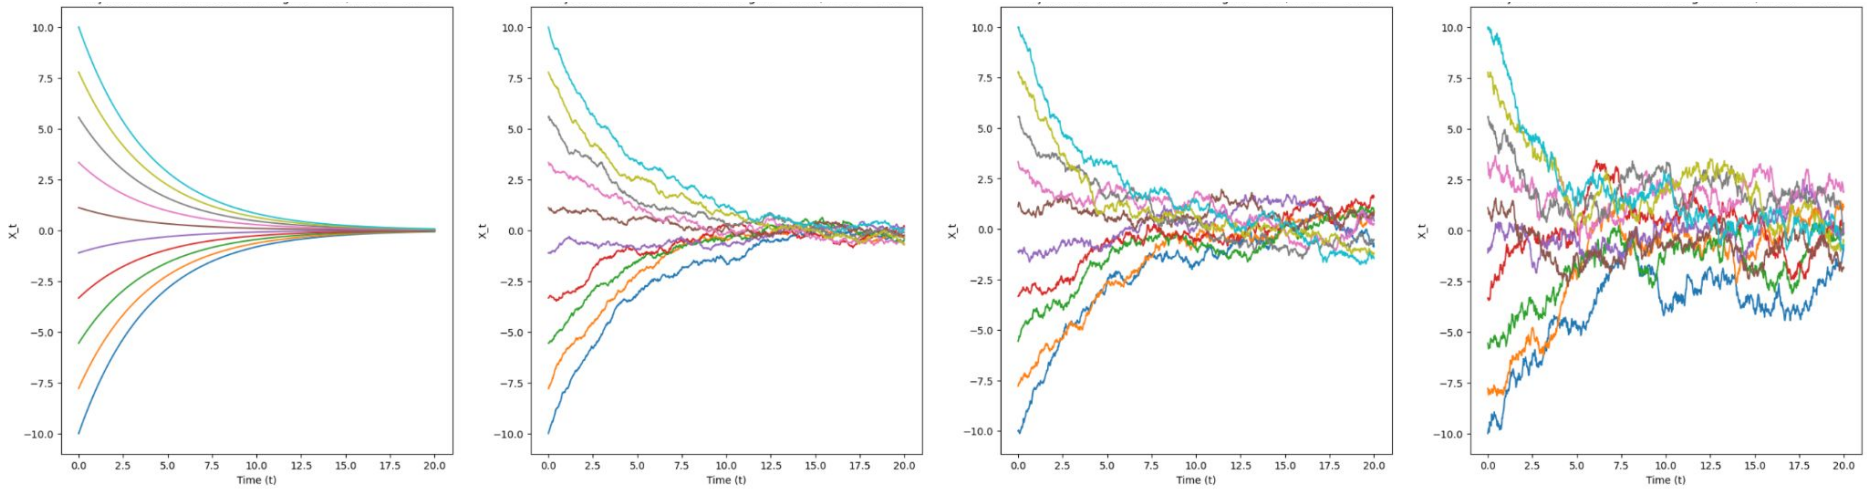

Increasing diffusion coefficient σ

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

Require: Neural network u_t^θ, number of steps n, diffusion coefficient σ_t

Set $t = 0$
Set step size $h = \frac{1}{n}$
Draw a sample $X_0 \sim p_{\text{init}}$
for $i = 1, \ldots, n-1$ **do**
    Draw a sample $\epsilon \sim \mathcal{N}(0, I_d)$
    $X_{t+h} = X_t + h u_t^\theta(X_t) + \sigma_t \sqrt{h} \epsilon$
    Update $t \leftarrow t + h$
end for
return X_1

ODE Sampling

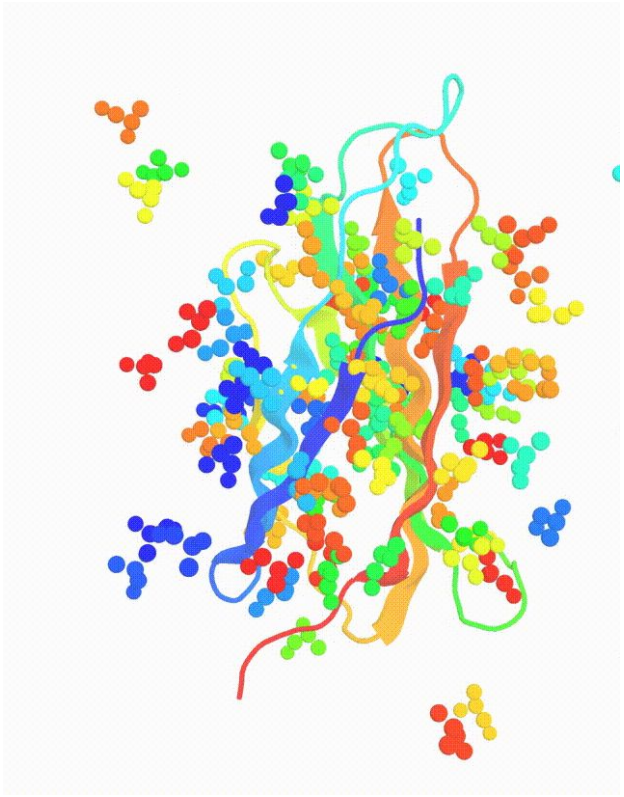

SDE Sampling

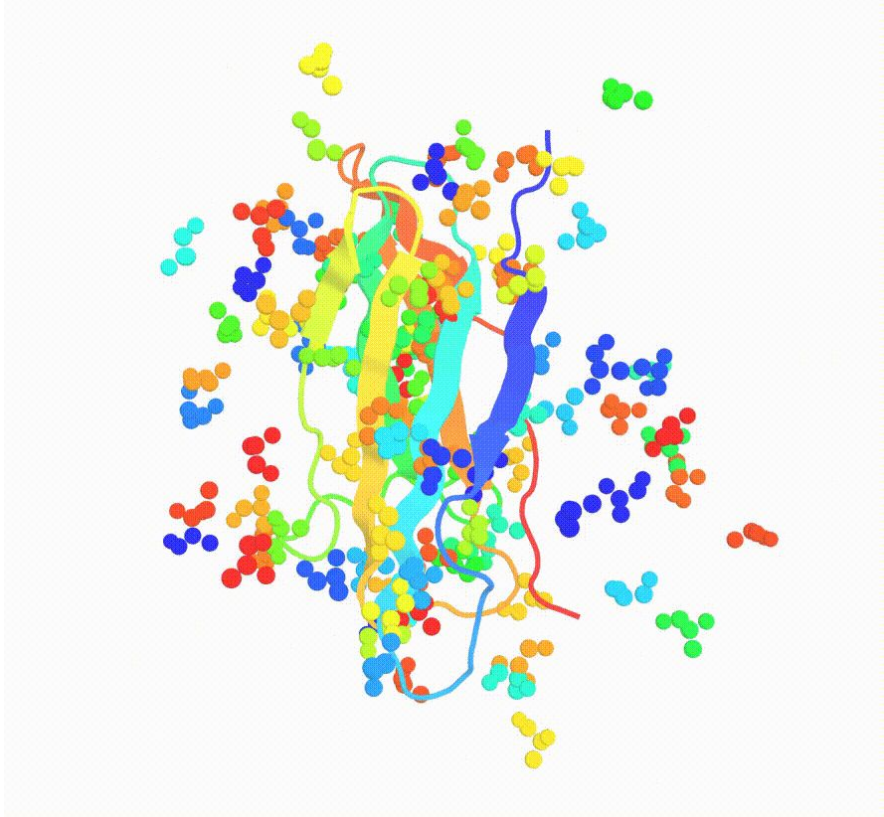

Logistics

Website: https://diffusion.csail.mit.edu/

How to pass this class:

1. Come to lecture
2. **Do the labs (necessary to pass)!**

Support:

1. Use the lecture notes (self-contained)
2. Come to office hours

Lab 1 is out today (see website)! We recommend doing it as soon possible for you!

Next class:

Thursday (Day after tomorrow), 11am-12:30pm

E25-111 (same room)

Office hours: Wednesday (tomorrow), 11am-12:30pm in 36-144