



Carnegie Mellon University

# 10-799 Diffusion & Flow Matching

---

Yutong (Kelly) He  
*Class overview, Jan 13<sup>th</sup>, 2026*

# The year of 2025 is insane



ChatGPT, OpenAI

<https://thunderdungeon.com/2025/03/28/studio-ghibli-memes-ghiblifly-memes/>



Veo, Google

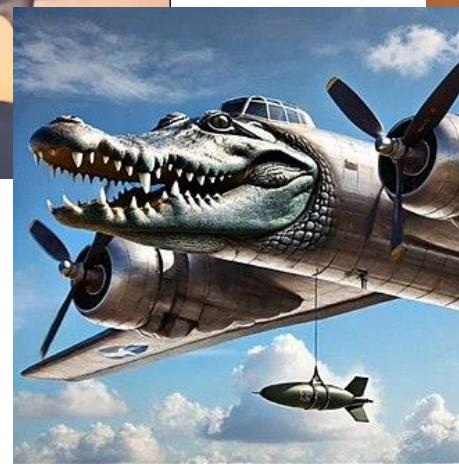
<https://www.youtube.com/shorts/Z8AOj5D70pA>



Sora, OpenAI

<https://www.tiktok.com/@longliveai/video/7563640496497626390>

# Absolutely the best thing that happened to me in 2025



# Besides memes, the industry is blooming in other directions too

[◀ Back to index](#)

SEPTEMBER 25, 2025

FLUX.1 Kontext now  
in Adobe Photoshop:  
Powering Every Pixel



inception

Products ▾ Resources ▾ Company ▾

Try Our API

Announcing our \$50M Seed Round, led by Menlo Ventures

The Fastest LLMs Ever Built

Diffusion LLMs: A Breakthrough for Speed and Quality

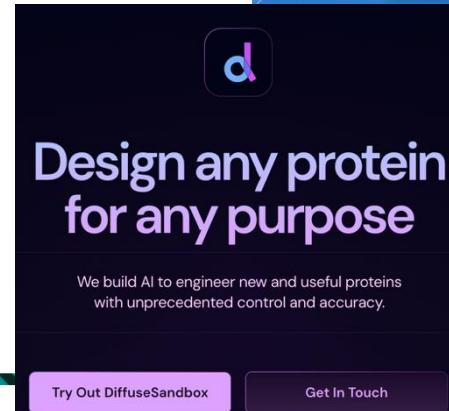
How can I help you code?

Need suggestions?

By using Mercury, you agree to our [Terms of Use](#) and have read our [Privacy Policy](#)



CuspAI is the frontier AI company  
on a mission to solve the  
breakthrough materials needed  
to power human progress.



Design any protein  
for any purpose

We build AI to engineer new and useful proteins  
with unprecedented control and accuracy.

Try Out DiffuseSandbox

Get In Touch

Carnegie  
Mellon  
University

All these advances are built based on the same underlying technology



# This is what we are going to learn in this class



# This is what we are going to learn in this class

In this class, we are going to learn about **diffusion & flow matching** – the technology behind the great AI memes advancements in image generation & beyond in recent years. In particular, we will learn

- The **intuitions and math** behind these algorithms
- How to **implement** basic diffusion & flow matching models in Python
- How to **train a good image generation model** using GPUs **?**
- **Techniques** that **improves** upon the vanilla models
- How to conceptually extend these frameworks to **discrete data**
- How to **demonstrate your thoughts** and **your work**

# What makes a good image generation model



# Fidelity

Do the generated images look “real”?

- Do the people generated have 6 fingers?
- Are the images over/under saturated?
- Do the images have a lot of unnatural artifacts?

...

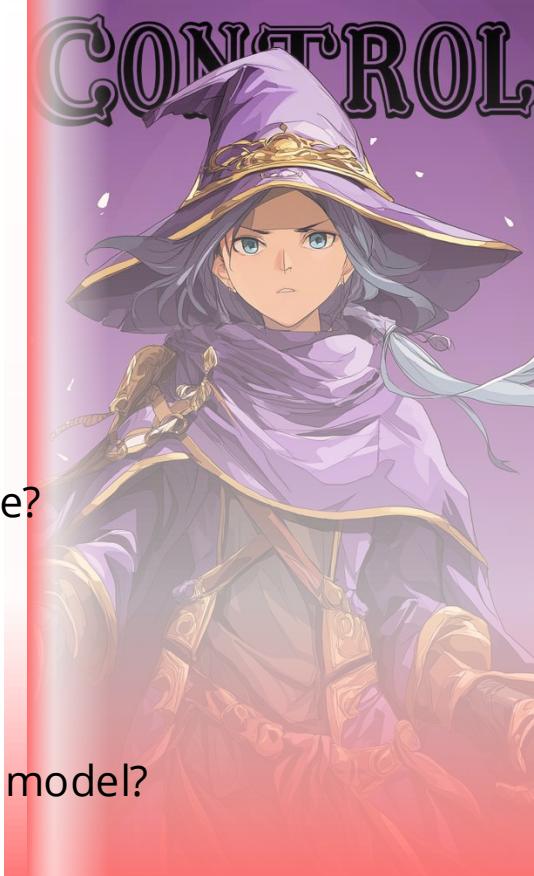


# Controllability

Can I control what I generate in however way I want?

- Can I use text to describe the image I want to generate?
- Can I use other images to specify the image I want to generate?
- Can I generate my cat playing basketball?
- Can I personalize the model?
- Can I use other models to interact with the image generation model?

...

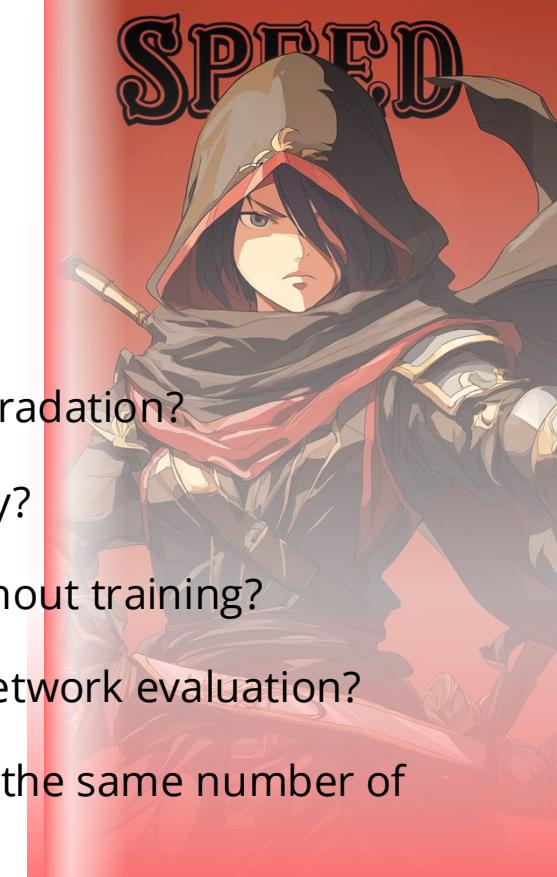


# Speed

How fast is the generation?

- Is it possible to speed up the generation without quality degradation?
- Is it possible to speed up and improve the generation quality?
- Do I need extra training to speed things up or can I do it without training?
- Can I reduce the wall clock time with the same number of network evaluation?
- Can I reduce the wall clock time with the same network and the same number of network evaluation?

...



# Choose your path and level up

I am going to  
be the best  
Pokemon AI  
trainer!

Learning in this class is the same as playing (RPG) video games!

- You start in the beginners' village to learn the basics



# Choose your path and level up

I am going to  
be the best  
Pokemon AI  
trainer!

Learning in this class is the same as playing (RPG) video games!

- You start in the beginners' village to learn the basics
- Before leaving the village, you choose the path you want to focus on and start learning specialized skills



# Choose your path and level up

I am going to  
be the best  
Pokemon AI  
trainer!

Learning in this class is the same as playing (RPG) video games!

- You start in the beginners' village to learn the basics
- Before leaving the village, you choose the path you want to focus on and start learning specialized skills
- During your exploration, you will fight (multiple) small bosses. These fights can help you improve your skills



# Choose your path and level up

Learning in this class is the s

- You start in the beginners
- Before leaving the village learning specialized skills
- During your exploration, help you improve your sk
- At the end, you will defeat the final boss and obtain the holy grail



I am going to  
be the best  
Pokemon AI  
trainer!



d start

hts can

# Homework (70%): Choose your path and level up

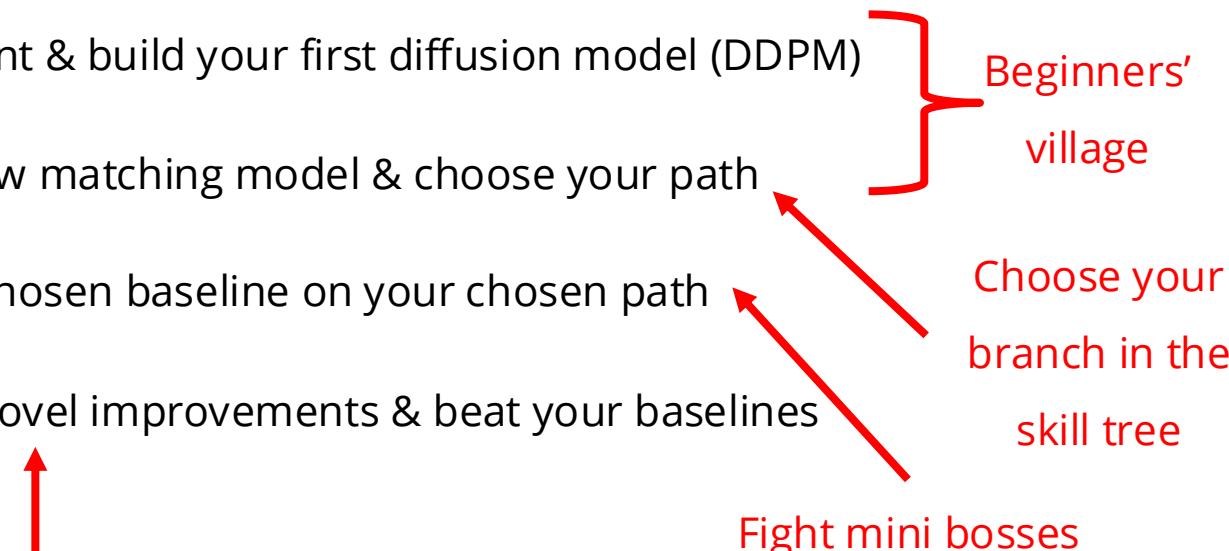
- HW1: Set up environment & build your first diffusion model (DDPM)
- HW2: Build your first flow matching model & choose your path
- HW3: Implement your chosen baseline on your chosen path
- HW4: Implement your novel improvements & beat your baselines

**Fight final boss**

**Fight mini bosses**

**Choose your branch in the skill tree**

**Beginners' village**



# Homework (70%): Choose your path and level up

- HW1: Set up environment & build your first diffusion model (DDPM) <- 15%
- HW2: Build your first flow matching model & choose your path <- 15%
- HW3: Implement your chosen baseline on your chosen path <- 20%
- HW4: Implement your novel improvements & beat your baselines <- 20%

# Final Presentation (15%): Showcase your work

In the last week of class, we will have two poster sessions for everyone to showcase the work they did for their homework. In order to receive full credits, you'll have to

- Submit your poster PDF the day before the first poster session <- 5%
- Actually present your poster at your assigned poster session <- 5%
- Attend the other poster session to check out other people's work <- 5%

We will also conduct a poll at the poster sessions for everyone to vote for their favorite posters. The best poster for each path will receive a small award :)

# AI as your companion on this journey

This class is AI-friendly and open everything!

You can use

- Any AI assistant/tools
- Any open source code
- Any pre-trained model
- Any research paper/book/tutorials etc
- You can also discuss with other human

I am going to  
be the best  
Pokemon AI  
trainer!



Your AI

# AI as your companion on this journey

This class is AI-friendly and open everything!

That being said, you should

- Complete all assignments on your own
- Cite all resources that you use (including AI)
- Do not copy from your classmates
- Do not claim other human or AI's work as your own
- Do not plagiarize any source without attribution

I am going to  
be the best  
Pokemon AI  
trainer!



Your AI

# Human companions that shall help you learn



**Kelly He**  
Instructor



**Russ Salakhutdinov**  
Advisor



**Zico Kolter**  
Advisor



**Nichelle Phillips**  
Education Associate



**Douglas Chen**  
TA



**Krish Vijayan**  
TA

# Office hours starts next week



**Kelly He**  
Instructor



**Douglas Chen**  
TA



**Krish Vijayan**  
TA

Office hours (Kelly + TAs) starts next week

Schedule will be announced on the class website and Discord

TAs are your peers!

# This is what we are going to learn in this class

In this class, we are going to learn about **diffusion & flow matching** – the technology behind the great AI memes advancements in image generation & beyond in recent years. In particular, we will learn

- The **intuitions and math** behind these algorithms
- How to **implement** basic diffusion & flow matching models in Python
- How to **train a good image generation model** using GPUs
- **Techniques** that **improves** upon the vanilla models
- How to conceptually extend these frameworks to **discrete data**
- How to **demonstrate your thoughts** and **your work**

I will provide  
some intuitions  
math, and  
pointers

# This is what we are going to learn in this class

In this class, we are going to learn about **diffusion & flow matching** – the technology behind the great AI memes advancements in image generation & beyond in recent years. In particular, we will learn

- The **intuitions and math** behind these algorithms
- How to **implement** basic diffusion & flow matching models in Python
- How to **train a good image generation model** using GPUs
- **Techniques** that **improves** upon the vanilla models
- How to conceptually extend these frameworks to **discrete data**
- How to **demonstrate your thoughts** and **your work**

You & Your  
pikachu  
figure these  
out through  
experiments

# In-class 10-min Quizzes (15%): Sanity check

The most common concern we got about this fully AI-friendly setting is that, how do we **make sure that we are actually learning stuff** when AI helps us do homework?

In order to provide some sort of reference point, we will have **7 in-class 10-min quizzes** through out the class to benchmark learning in a traditional setting (so **everything closed**).

That being said, you should just treat these quizzes as sanity checks: they will all be **simple questions** with no tricks and they should not heavily impact your final grade.

Doing bad on quizzes simply means that you need to catch up a bit and **you can drop the lowest two quizzes**.

# Extra Credits (1.5~5%): Tell us what you have learned (and what you have not learned)

We want your feedback!

Through out the class, we will release two optional surveys (1.5% each) and conduct one optional interview (2%) to collect feedback on your AI-assisted learning experience.

The extra credits will be applied upon completion in both cases.

# Lecture topics & schedule

- Week 1~2: Basics & foundational algorithms for diffusion & flow matching
- Week 3~4: Advanced algorithms for each path
- Week 5: SOTA models & industrial applications
- Week 6: Discrete variants
- Week 7: Final poster presentation

# Class website & Gradescope

Class website: <https://kellyyutonghe.github.io/10799S26/>

One-stop-shop for all the resources and pointers you will need

Gradescope: <https://www.gradescope.com/courses/1207241>

The place where you will submit your homework and view your grades

(Everyone will be added on Friday)

# Discord server (CMU only)

Our main hub for all communication:

- General question answering
- Announcements
- Forum discussion
- Memes
- Virtual office hours
- Virtual study room



# Computational resources



# Modal

We are super grateful that Modal is sponsoring the compute resources for this class!

- Serviceless cloud compute, everything in code, super easy to use
- All registered students will get \$500 credits for this class (instructions will be sent on Friday 1/16)
- Every month you also get \$30 for free
- This should translate to ~280 GPU hours on L40S in total
- **Modal is giving a guest lecture on Friday 1/16 5 PM SH 105** (same time same location of the lecture) on how to use their service!

# Computational resources



We are super grateful that AWS is also sponsoring this class!

- Classic cloud compute service, many tutorials online
- Details are still being finalized but we are expecting to give each student \$100~\$300 AWS credits as well (will be available later)
- AWS is also sponsoring pizza for our final poster sessions :)

# Computational resources

There are other compute resources available

- CMU clusters: Babel, FLAME, etc
- Your lab's cluster
- Renting other cloud compute service  
(<https://kellyutonghe.github.io/10799S26/self-study/#computational-resources>)

# Waitlist, unofficial auditing & self-study

- If more students drop from this class, we will gradually admit students from the waitlist until Friday 1/16 noon
- If you are a CMU student, please join the discord server regardless
- You can feel free to unofficially audit in any form you'd like
- No need to submit the official audit form
- Guides on auditing and self-studying can be found here:  
<https://kellyutonghe.github.io/10799S26/self-study/>

# Class recording & Zoom

This course will entirely be recorded and streamed on Zoom.

- You can watch the lecture from home via Zoom (although I'd say coming to class shall be fun too ☺), link shared on Discord (CMU only)
- You can access the lecture recordings in the following way
  - Panopto (CMU only for now): link shared on Discord
  - YouTube (Public): link can be found on the class website later (will have some delay)

# Books, tutorials, papers & other helpful resources

Full list here: <https://kellyutonghe.github.io/10799S26/resources/>

## Books

---

### **The Principles of Diffusion Models**

*Chieh-Hsin Lai, Yang Song, Dongjun Kim, Yuki Mitsufuji, Stefano Ermon*

2025

Comprehensive monograph covering diffusion models, flow matching, and transport-based generative modeling from first principles.

[Website](#) • [arXiv](#) • [PDF](#) • [Blog Post](#)

# Books, tutorials, papers & other helpful resources

## Courses

### Stanford CS236: Deep Generative Models

*Stefano Ermon, Aditya Grover*

2023

Stanford course on generative models including VAEs, GANs, EBMs, normalizing flows, diff autoregressive models.

[Course Website](#) • [Lecture Notes](#) • [Video Recordings](#)

### CMU 10-423/10-623: Generative AI

*Matt Gormley, Yuanzhi Li, Henry Chai, Pat Virtue, Aran Nayebi*

2025

CMU course on generative models including LLMs, GANs, and diffusion models.

[Course Website](#) • [Video Recordings](#)

### CMU 18-789: Deep Generative Modeling

*Beidi Chen, Xun Huang*

2025

CMU course on generative models including LLMs, VAEs, and diffusion models.

[Course Website](#)

### MIT 6.S184: Introduction to Flow Matching and Diffusion Models

*Peter Holderrieth, Ezra Erives*

2025

MIT class on diffusion and flow matching from a flow-based theoretical perspective.

[Course Website](#) • [Lecture Notes](#) • [Video Recordings](#)

### CMU 10-708: Probabilistic Graphical Models

*Andrej Risteski, Albert Gu*

2025

CMU course that focuses on probabilistic modeling (including some deep generative models from perspective).

[Course Website](#)

### Stanford CS228: Probabilistic Graphical Models

*Stefano Ermon*

2024

Stanford course that focuses on probabilistic modeling.

[Course Website](#) • [Lecture Notes](#)

# Books, tutorials, papers & other helpful resources

## Tutorials

### **Generative Modeling by Estimating Gradients of the Data Distribution**

*Yang Song*  
2021

Introduction to score-based generative models and their connection to diffusion models.

### **What are Diffusion Models?**

*Lilian Weng*  
2021

Comprehensive introduction to diffusion models with clear explanations and intuitive visualizations.

### **Understanding Diffusion Models: A Unified Perspective**

*Calvin Luo*  
2022

Unifies VAEs, hierarchical VAEs, and diffusion models under a single framework.

[arXiv](#) • [Blog Post](#)

### **Flow Matching Guide and Code**

*Yaron Lipman, Marton Havasi, Peter Holderith, Neta Shaul, Matt Le, Brian Karrer, Ricky T. Q. Chen, David Lopez-Paz, Heli Ben-Hamu, Itai Gat*  
2024

Comprehensive guide to flow matching with code examples and applications.

[arXiv](#) • [PDF](#)

# The only goal in this class is to help you learn & build

This is why

- Lecture attendance is not required, feel free to join on Zoom, watch recordings later or disregard lectures all together if you dislike going to class
  - You still need to take the quizzes in person, lmk ahead of time if you have any emergency/need accommodation
- Your grade will only be curved up and never curved down, and if everyone does well, everyone gets an A/A+ (cutoffs will be announced mid-class)
- Feel free to use any and all resources you can find