

BME 6938 · ADVANCED BIOMEDICAL AI · SPRING 2027

LECTURE 11-1

Generative Models and AI for Science

Diffusion models for protein design and structure



Parisa Rashidi · Herbert Wertheim College of Engineering, University of Florida



Where We Are

From models that predict and act to models that generate.

Week 9

**Retrieval, tools, and
agents**

Models that look things up and
act

Week 11: today

**Generative models and
AI for science**

Diffusion for protein design and
structure

Week 12

**Evaluation, robustness,
fairness**

Benchmarks, intervals,
shortcuts

Predict a Structure, or Design One

Thursday's two papers use diffusion for opposite directions.

Prediction: AlphaFold 3

given: sequences, ligands, ions

output: 3D atom coordinates

diffusion over raw atoms

Design: RFdiffusion

given: a motif or a target

output: a new backbone

then ProteinMPNN writes sequences


diffusion over residue frames




Adding Noise

Learn to Generate by Learning to Undo Noise


Two directions: a fixed one that destroys, a learned one that rebuilds.




x_0 real backbone



x_{100}



x_{500}



x_{999} pure noise

forward: add noise (fixed, no learning)

reverse: remove noise (learned network)

Jump to Any Step in One Line

x_t is one weighted mix of the structure and a single noise draw.

$$x_t = \sqrt{\text{signal}_t} \cdot x_0 + \sqrt{(1 - \text{signal}_t)} \cdot \text{noise}$$

x_0

the clean backbone

noise

one draw from a standard Gaussian


signal_t

product of $(1 - \beta)$ over the first t steps; falls from 1 toward 0

The Schedule

β rises from 0.0001 to 0.02 over 1,000 steps (the standard DDPM schedule).

Step t	Structure weight $\sqrt{\text{signal}_t}$	Noise weight $\sqrt{1 - \text{signal}_t}$
100	0.95	0.32
250	0.72	0.69
500	0.28	0.96
999	0.006	1.00



Training Data for Free

One real structure gives a thousand training examples.

1. Pick

a real structure x_0 and a random step t

2. Noise it

$$x_t = \sqrt{\text{signal}_t} \cdot x_0 + \sqrt{(1 - \text{signal}_t)} \cdot \text{noise}$$

3. Predict

the network sees x_t and t , outputs predicted noise

4. Loss

mean squared error between true and predicted noise

Animation: Diffusion, Part 1: Adding Noise

Animation clip: week11-clip1.mp4

Silent version embedded (click to play). The narrated version is in the course repository.



Learning the Way Back

One Reverse Step

Move a little toward the predicted clean structure, then add a little fresh noise.

Input

the noisy structure x_t and the step number t

Output

the predicted noise, equivalently an estimate of x_0

Step

move part of the way toward that estimate

Then

add a little fresh noise and repeat, $t = 1000$ down to 1

Start from different noise and you get a different structure: generation, not retrieval.


Two Papers, Two Representations

What the diffusion model works on decides what the model can cover.

	RFdiffusion	AlphaFold 3
Diffuses over	residue frames	raw atom coordinates
Protein-specific machinery	frames and orientations	none
Covers	protein backbones	proteins, nucleic acids, ligands, ions
Built from	a fine-tuned structure predictor	a pairformer trunk plus diffusion

A Backbone Is a Shape, Not a Protein

The design pipeline has a second stage.



Animation: Diffusion, Part 2: Learning the Way Back

Animation clip: week11-clip2.mp4

Silent version embedded (click to play). The narrated version is in the course repository.

When the Textbook Step Is Not What You Want

RFdiffusion reports better results without the fresh noise.

Textbook DDPM

add fresh noise each reverse step

keeps samples diverse

standard for images

RFdiffusion for design

often adds no fresh noise

higher in silico success in 17 of 23

less diversity, more that fold



Steering the Sampler

Unconditional Sampling Is Not Enough

Science asks for a specific structure, not any valid one.

Unconditional

start from noise
get some valid structure

no control over what

Conditional

give the denoiser a requirement
it is respected at every step

many scaffolds for one motif

Conditioning versus Guidance

Two different levers; RFdiffusion uses the first.

Conditioning (RFdiffusion)

extra input to the denoiser

held fixed at every step

e.g. motif coordinates

the motif never moves

Classifier-free guidance (images)

mix two predictions after the fact

$\text{uncond} + w \cdot (\text{cond} - \text{uncond})$

$0.2 + 3 \times (0.5 - 0.2) = 1.1$

larger w : follows harder, less diverse

Guidance numbers illustrative. RFdiffusion does not use classifier-free guidance.

Animation: Diffusion, Part 3: Steering the Sampler


Animation clip: week11-clip3.mp4

Silent version embedded (click to play). The narrated version is in the course repository.

Results, and Which Ones Count

In silico benchmarks, then the wet lab.

Motif scaffolding benchmark (25 problems)



scored in silico

Tested in the lab

five targets

19% experimental success

fewer than 100 designs per target

best flu binder: 28 nM

cryo-EM matched the design to 0.63 Å

The In Silico Judge Shares the Designer's Blind Spots

A benchmark scored by a model can agree with the model.

Circular

designs scored by a predictor trained on the same structures

Independent

expression, binding affinity, a solved structure

In the notebook

the toy is scored by a classifier trained on the same data, on purpose

In the Notebook: What the Toy Does Not Reproduce

A 2D diffusion model from scratch, and three honest negative results.

What was measured	Result	Reading
Class of generated chains	100% correct	the shapes are learned
Bond length (helix)	0.46 against 0.52	slightly short
Motif held in place	RMSE 0.026	conditioning works
Scaffold diversity for one motif	0.65 against 5.66 free	the toy is too small
Guidance trade-off	best realism at $w = 2$	not the textbook curve

Recap: Diffusion for Structure and Design

Forward

Fixed noising; x_t is one weighted mix; one structure gives 1,000 examples.

Reverse

A network predicts the noise; sampling walks from noise to a structure.

Representation

Residue frames (RFdiffusion) or raw atoms (AlphaFold 3).

Evidence

23 of 25 in silico; 19% in the lab. Prefer the number a model did not score.

Before Thursday

Read

Watson et al., RFdiffusion (Nature 2023) and Abramson et al., AlphaFold 3 (Nature 2024).

Write

Reading responses on Canvas before 1:00 PM Thursday.

Submit

Project progress report due Fri Mar 26, 11:59 PM: baseline results and the remaining plan.

Explore

Week 11 reference notebook (ungraded): a 2D diffusion model from scratch, conditioning, and a guidance sweep.