

# Learning Protein Structure with a Differentiable Simulator

John Ingraham, Adam Riesselman, Chris Sander, Debora Marks


ICLR 2019

# Protein Folding

Grand challenge in structural biology

Useful for understanding and  
interfacing with proteins

Famously “solved” by AlphaFold2!



# Outline

1. NEMO – Why we're excited
2. Background
3. Model Design
4. Engineering
5. Data & Hyperparameters
6. Results

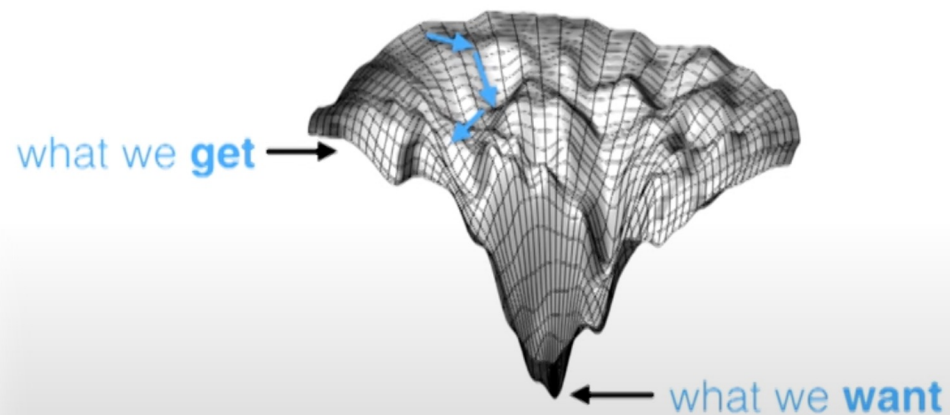
# How does a protein fold?

A good heuristic (from experimental research):

the **energy landscape model** assumes that proteins naturally achieve folds that minimize free energy.

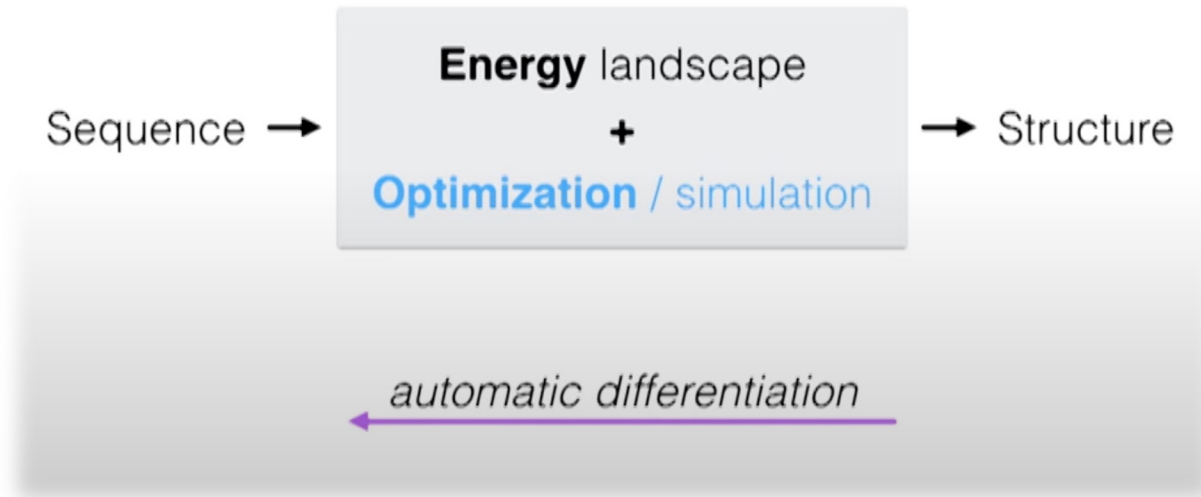
# Just fold the protein!


Even with a correct energy landscape,  
**simulator fails to converge**



# Just learn to simulate the universe!

Could we learn the **energy landscape**  
**and simulator** simultaneously?



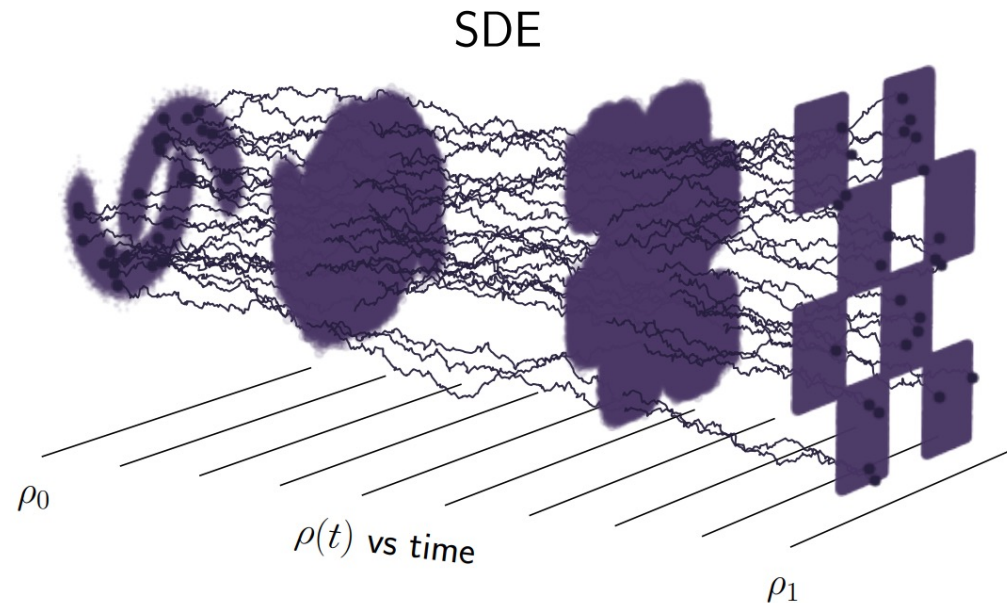


# NEMO

An end-to-end differentiable protein folding simulator.

# Related Work


- Learning distributions and sampling (SPEN, diffusion)
- Protein modeling (like AlphaFold)
- Differentiable computing (neural implicit representations, etc.)
- RNNs and exploding gradients, RL and discount factor





Background

# Boltzmann Distribution & Energy Landscapes



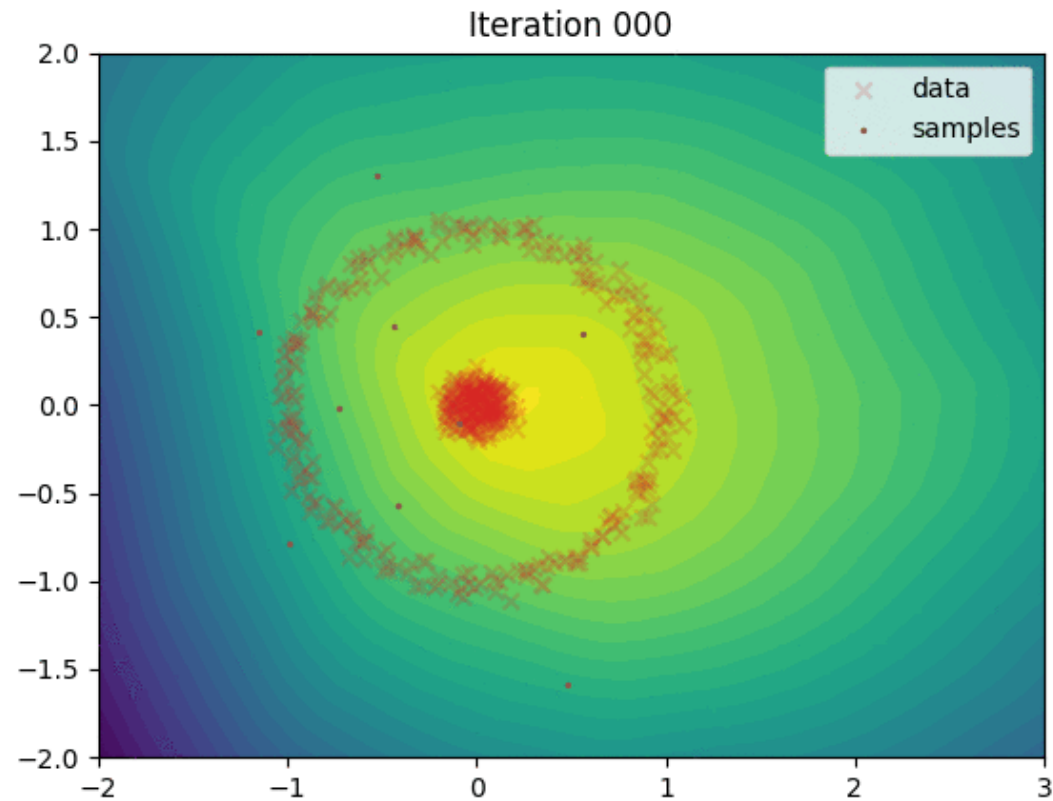
# Sampling low energy states

$$x^{(t+\varepsilon)} \leftarrow x^{(t)} - \frac{\varepsilon}{2} \nabla_x U^{(t)} + \sqrt{\varepsilon} p$$

**(Approximate) Langevin Dynamics**


We update **each atom's position**, by stepping **in the direction of the force**, and perturbing with **Gaussian noise**.

# Langevin Sampling in Action

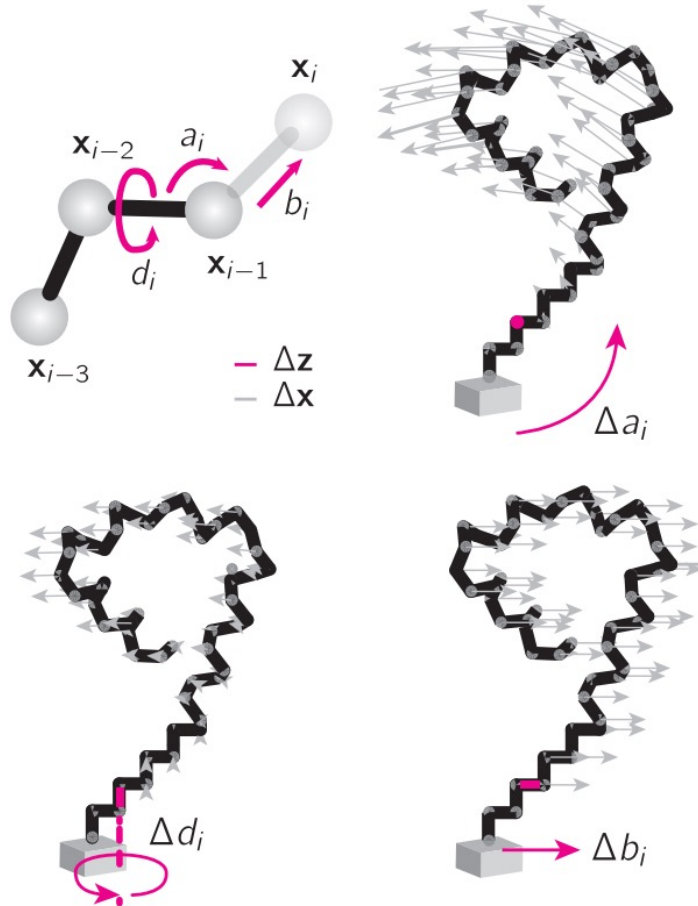


# A Differentiable Simulator

# Building Blocks



# Coordinate Choices



**Internal coordinates** favor coherent movement into new conformations.

**Cartesian coordinates** favor local structure rearrangements.

NEMO uses **both** coordinates, alternating between time steps.

# Sampling Atomic Position

---

## Algorithm 1: Direct integrator

---

**Input** : State  $\mathbf{z}^{(0)}$ , energy  $U(\mathbf{x})$ ,  
step  $\epsilon$ , time  $T$ , scale  $\mathbf{C}$

**Output** : Trajectory  $\mathbf{x}^{(0)}, \dots, \mathbf{x}^{(T)}$

Initialize  $\mathbf{x}^{(0)} \leftarrow \mathcal{F}(\mathbf{z}^{(0)})$ ;

**while**  $t < T$  **do**

    Compute forces  $\mathbf{f}_z = -\frac{\partial \mathbf{x}}{\partial \mathbf{z}}^T \nabla_{\mathbf{x}} U$ ;

    Sample  $\Delta \mathbf{z} \sim \mathcal{N}(\frac{1}{2} \epsilon \mathbf{C} \mathbf{f}_z, \epsilon \mathbf{C})$ ;

    ■  $\mathbf{z}^{(t+\epsilon)} \leftarrow \mathbf{z}^{(t)} + \Delta \mathbf{z}$ ;

    ■  $\mathbf{x}^{(t+\epsilon)} \leftarrow \mathcal{F}(\mathbf{z}^{(t+\epsilon)})$ ;

$t \leftarrow t + \epsilon$ ;

**end**

---

---

## Algorithm 2: Transform integrator

---

**Input** : State  $\mathbf{z}^{(0)}$ , energy  $U(\mathbf{x})$ ,  
step  $\epsilon$ , time  $T$ , scale  $\mathbf{C}$

**Output** : Trajectory  $\mathbf{x}^{(0)}, \dots, \mathbf{x}^{(T)}$

Initialize  $\mathbf{x}^{(0)} \leftarrow \mathcal{F}(\mathbf{z}^{(0)})$ ;

**while**  $t < T$  **do**

    Compute forces  $\mathbf{f}_z = -\frac{\partial \mathbf{x}}{\partial \mathbf{z}}^T \nabla_{\mathbf{x}} U$ ;

    Sample  $\Delta \mathbf{z} \sim \mathcal{N}(\frac{1}{2} \epsilon \mathbf{C} \mathbf{f}_z, \epsilon \mathbf{C})$ ;

    ■  $\tilde{\mathbf{x}} \leftarrow \mathbf{x}^{(t)} + \frac{\partial \mathbf{x}}{\partial \mathbf{z}}^{(t)} \Delta \mathbf{z}^{(t)}$ ;

    ■  $\mathbf{x}^{(t+\epsilon)} \leftarrow \mathbf{x}^{(t)} + \frac{1}{2} \left( \frac{\partial \mathbf{x}}{\partial \mathbf{z}}^{(t)} + \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{z}} \right) \Delta \mathbf{z}^{(t)}$ ;

$t \leftarrow t + \epsilon$ ;

Heun's Method

**end**

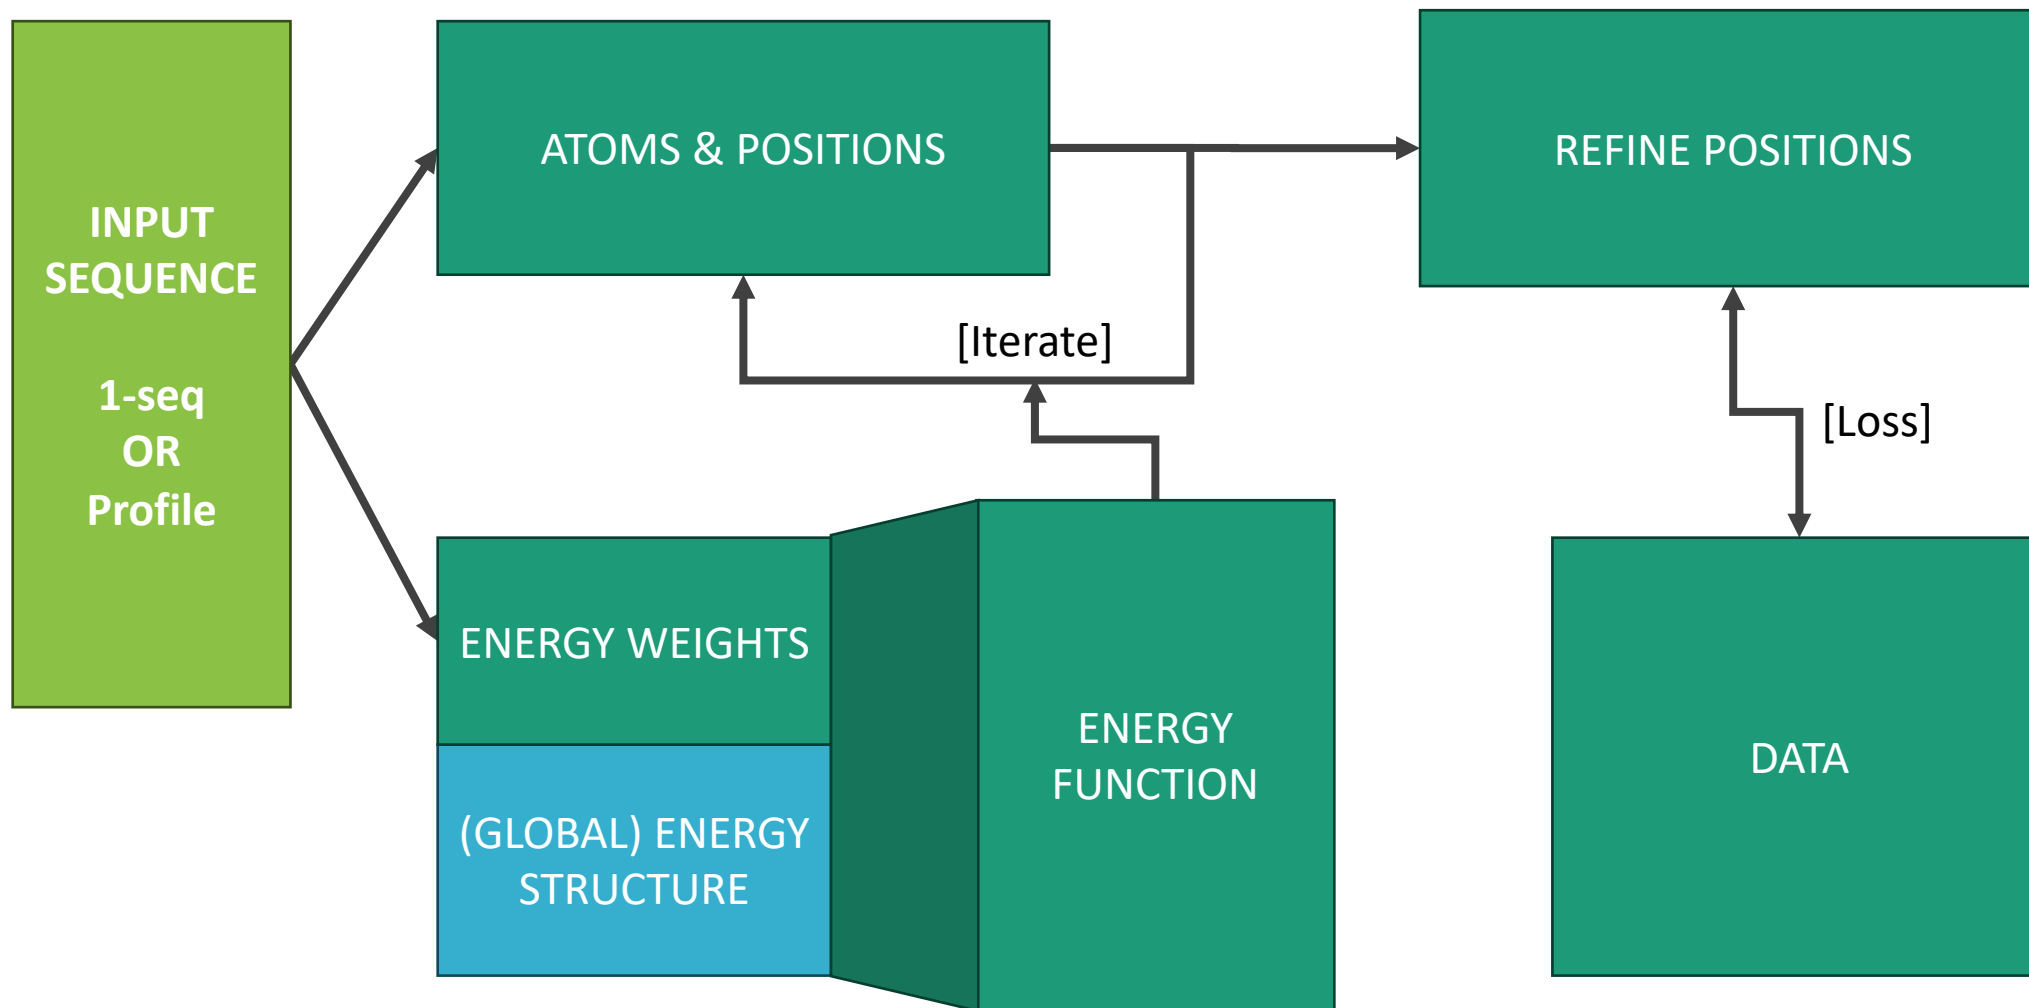
---

Parallelizable!

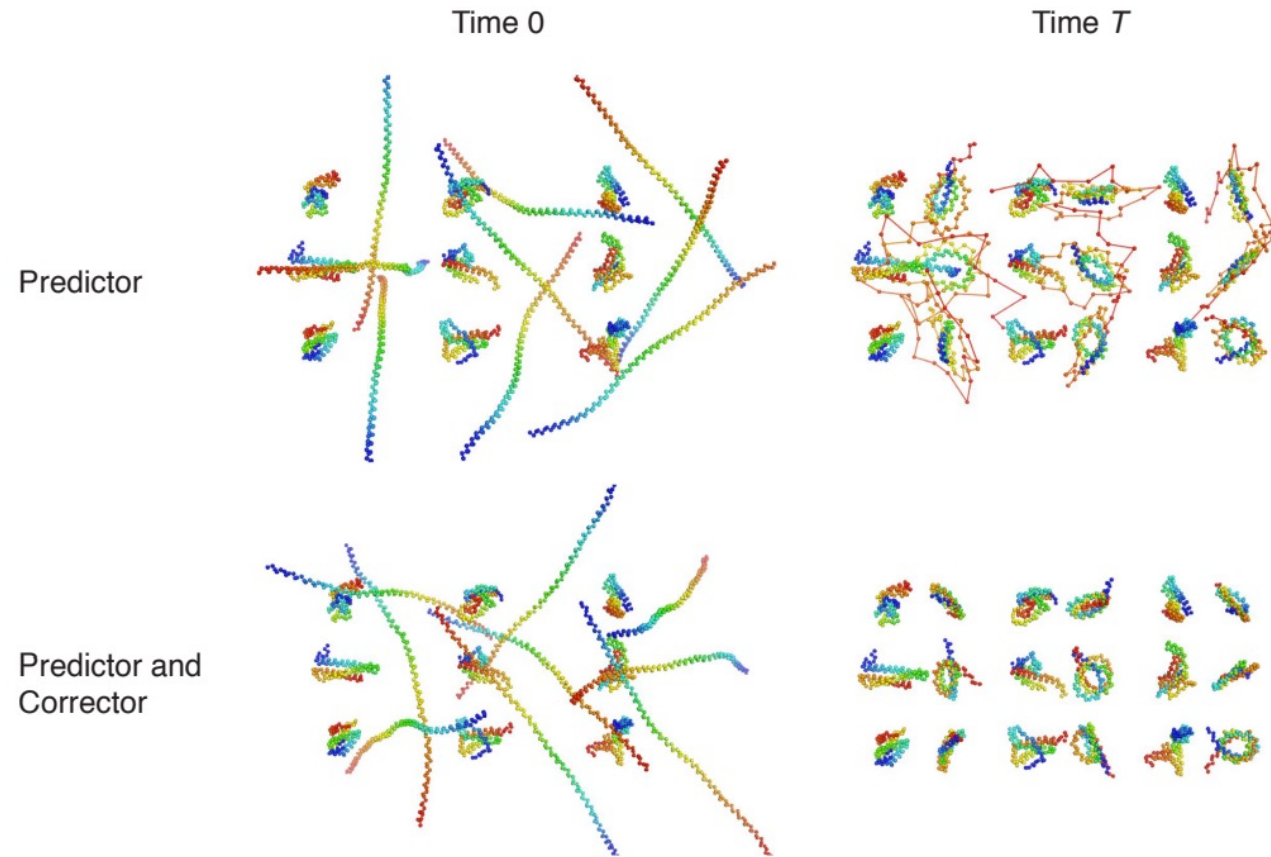




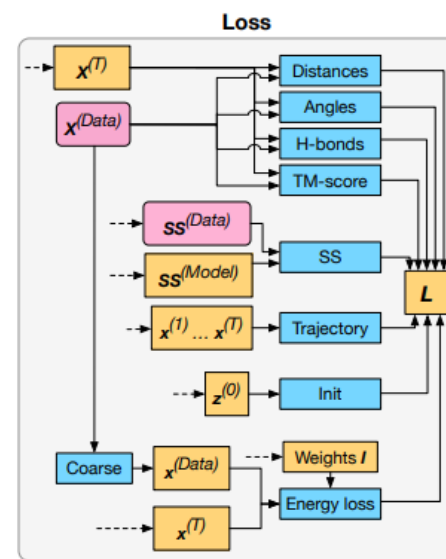
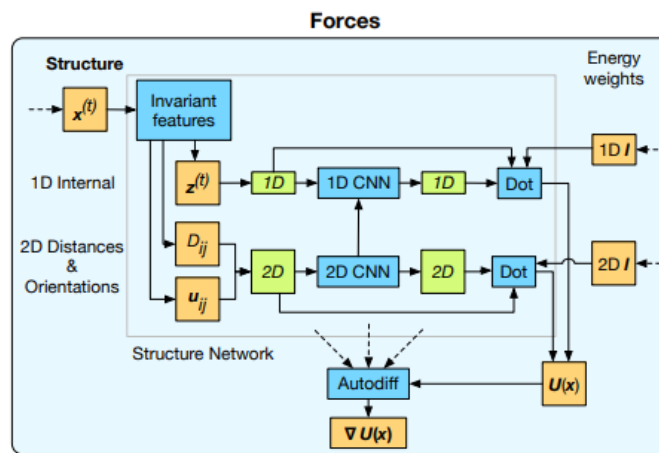


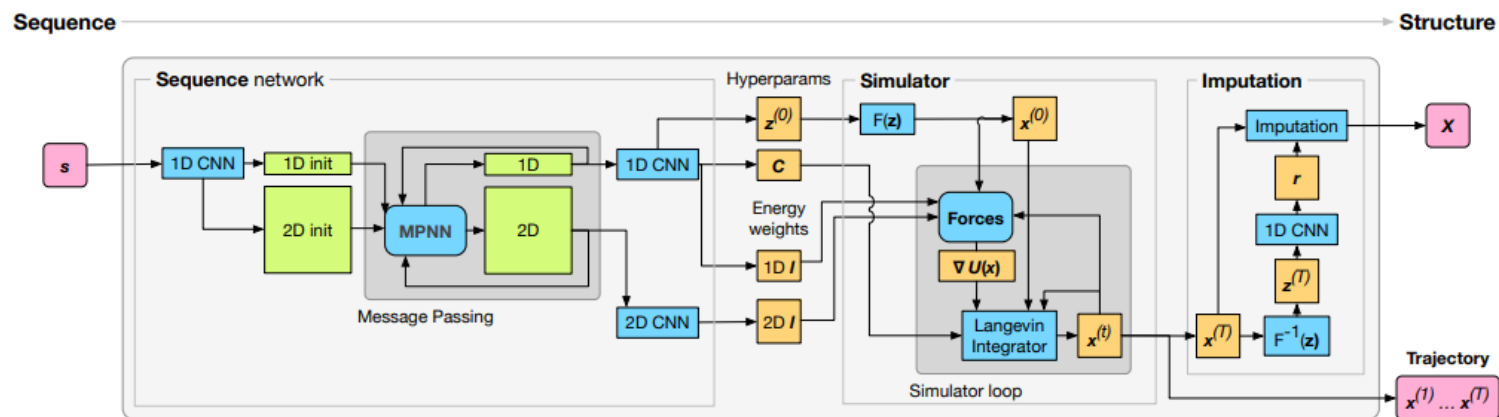
# Learning to simulate



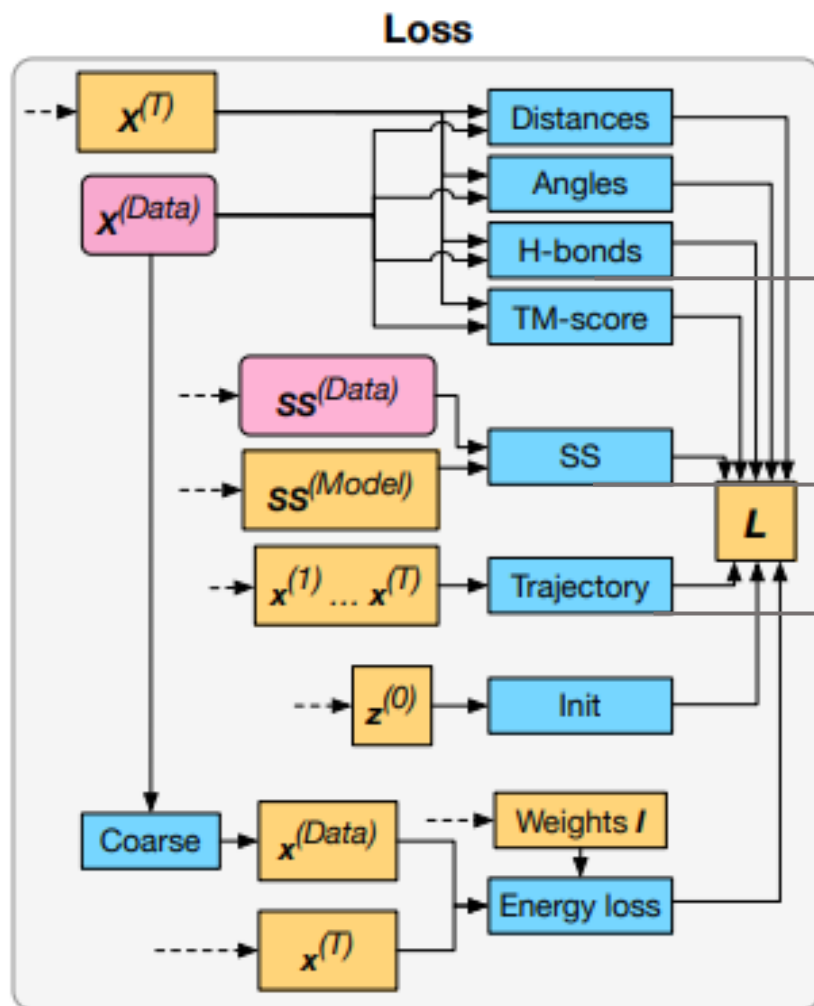
# Corrector & Imputation



# Full model architecture



# Focusing on Loss



$$\mathcal{L}_{ML} = U_{\theta}(\perp(\mathbf{x}^{(D)}); \mathbf{s}) - U_{\theta}(\perp(\mathbf{x}^{(M)}); \mathbf{s})$$

Probabilistic model of H-bonds + cross-entropy loss with data

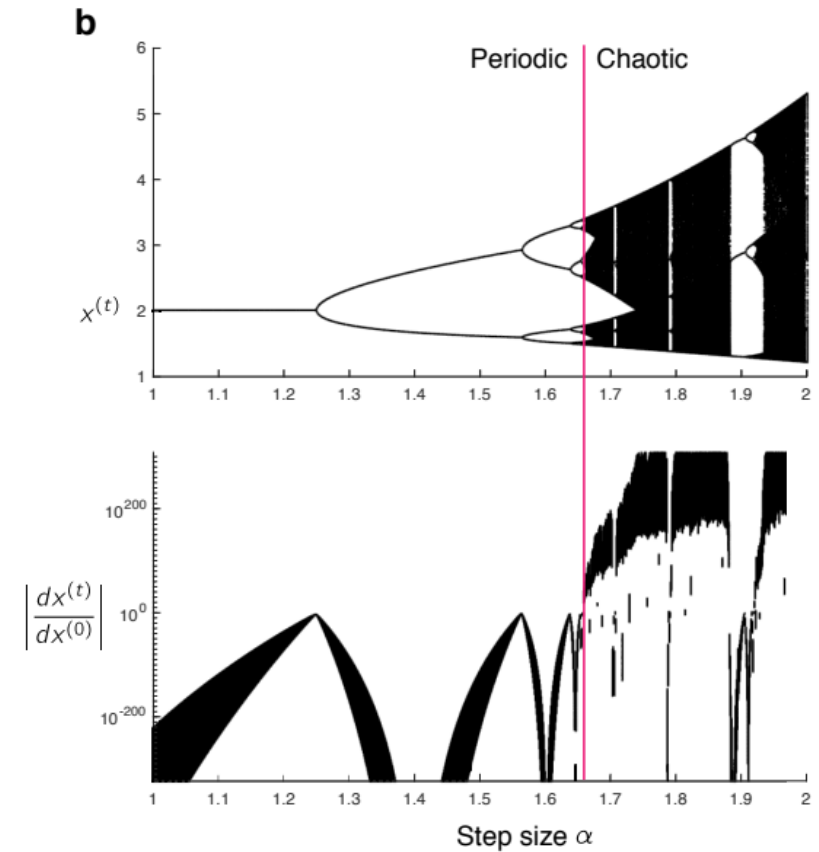

Predicting secondary structure (8 classes)

Required because of discount factor

# Training is *very* unstable!

Some problems:

- Gradients accumulate **exponentially** over time
- High forces explode simulations!
- Sensitivity to initial conditions → instability through time steps



# Damped Gradients

Gradient clipping? Doesn't work...

...so multiply each backwards iteration by  $\gamma$ ,  
like in reinforcement learning!

$$\frac{\partial \hat{\mathbf{x}}^{(t)}}{\partial \mathbf{x}^{(t-k)}} = \left( \left( \left( \left( \gamma \frac{\partial \mathbf{x}^{(t)}}{\partial \mathbf{x}^{(t-1)}} \right) \gamma \frac{\partial \mathbf{x}^{(t-1)}}{\partial \mathbf{x}^{(t-2)}} \right) \cdots \gamma \frac{\partial \mathbf{x}^{(t-k+1)}}{\partial \mathbf{x}^{(t-k)}} \right) = \gamma^k \frac{\partial \mathbf{x}^{(t)}}{\partial \mathbf{x}^{(t-k)}}.$$

# Speed Clipping

---

## Algorithm 3: Mixed Integrator

---

**Input** : Initial state  $\mathbf{z}^{(0)}$ , energy  $U(\mathbf{x})$ , time steps  $\epsilon_x, \epsilon_z$ , total time  $T$ , preconditioners  $\mathbf{C}_x, \mathbf{C}_z$ ,

**Output** : Trajectory  $\mathbf{x}^{(0)}, \dots, \mathbf{x}^{(T)}$

Initialize  $\mathbf{x}^{(0)} \leftarrow \mathcal{F}(\mathbf{z}^{(0)})$ ;

**while**  $t < T$  **do**

$\mathbf{f}_x \leftarrow \nabla_x U$ ;

$\Delta \mathbf{x}^{(Cart)} \leftarrow \text{CartesianStep}(\mathbf{x}^{(t)}, \mathbf{f}_x, \epsilon_x, \mathbf{C}_x)$ ;

$\Delta \mathbf{x}^{(Int)} \leftarrow \text{ClippedInternalStep}(\mathbf{x} + \Delta \mathbf{x}^{(Cart)}, \mathbf{f}_x, \epsilon_z, \mathbf{C}_z)$ ;


$\mathbf{x} \leftarrow \mathbf{x} + \text{Detrend}(\Delta \mathbf{x}^{(Cart)} + \Delta \mathbf{x}^{(Int)})$ ;

$t \leftarrow t + \epsilon$ ;

**end**

---

# Limiting Sensitivity to Initial Conditions



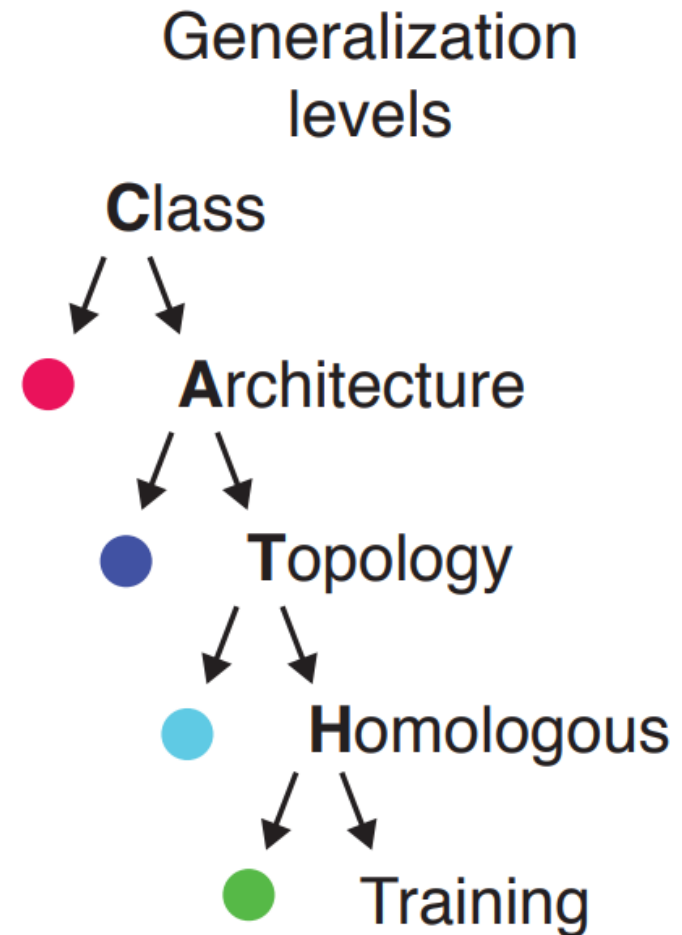

$$|F^{(t)}(\mathbf{x}, \theta) - F^{(t)}(\mathbf{x} + \Delta\mathbf{x}, \theta)| \leq |\Delta\mathbf{x}|$$

$$\mathcal{L}_{Lyapunov} = \max\left(0, \log \frac{|F(\mathbf{x}^{(t)}) - F(\mathbf{x}^{(t)} + \boldsymbol{\delta})|}{|\boldsymbol{\delta}|}\right)$$



# Data & Metadata

# Dataset: CATH -- hierarchy of protein structure



# Data Split

- Train: 35k folds
- Validation: 21k folds
- Test: 10k folds
  
- Testing generalization:
  - **H** validation: superfamilies excluded from train
  - **T** validation: topologies excluded from train
  - **A** validation: secondary structures excluded from train

# Guesswork

figure [Figure 12](#). We note that the validation set was not explicitly used to tune hyperparameters due to the large cost of training (2 months on 2 M40 GPUs), but we did keep track of validation statistics during training.

!!

# So... how did we do?

A scary question

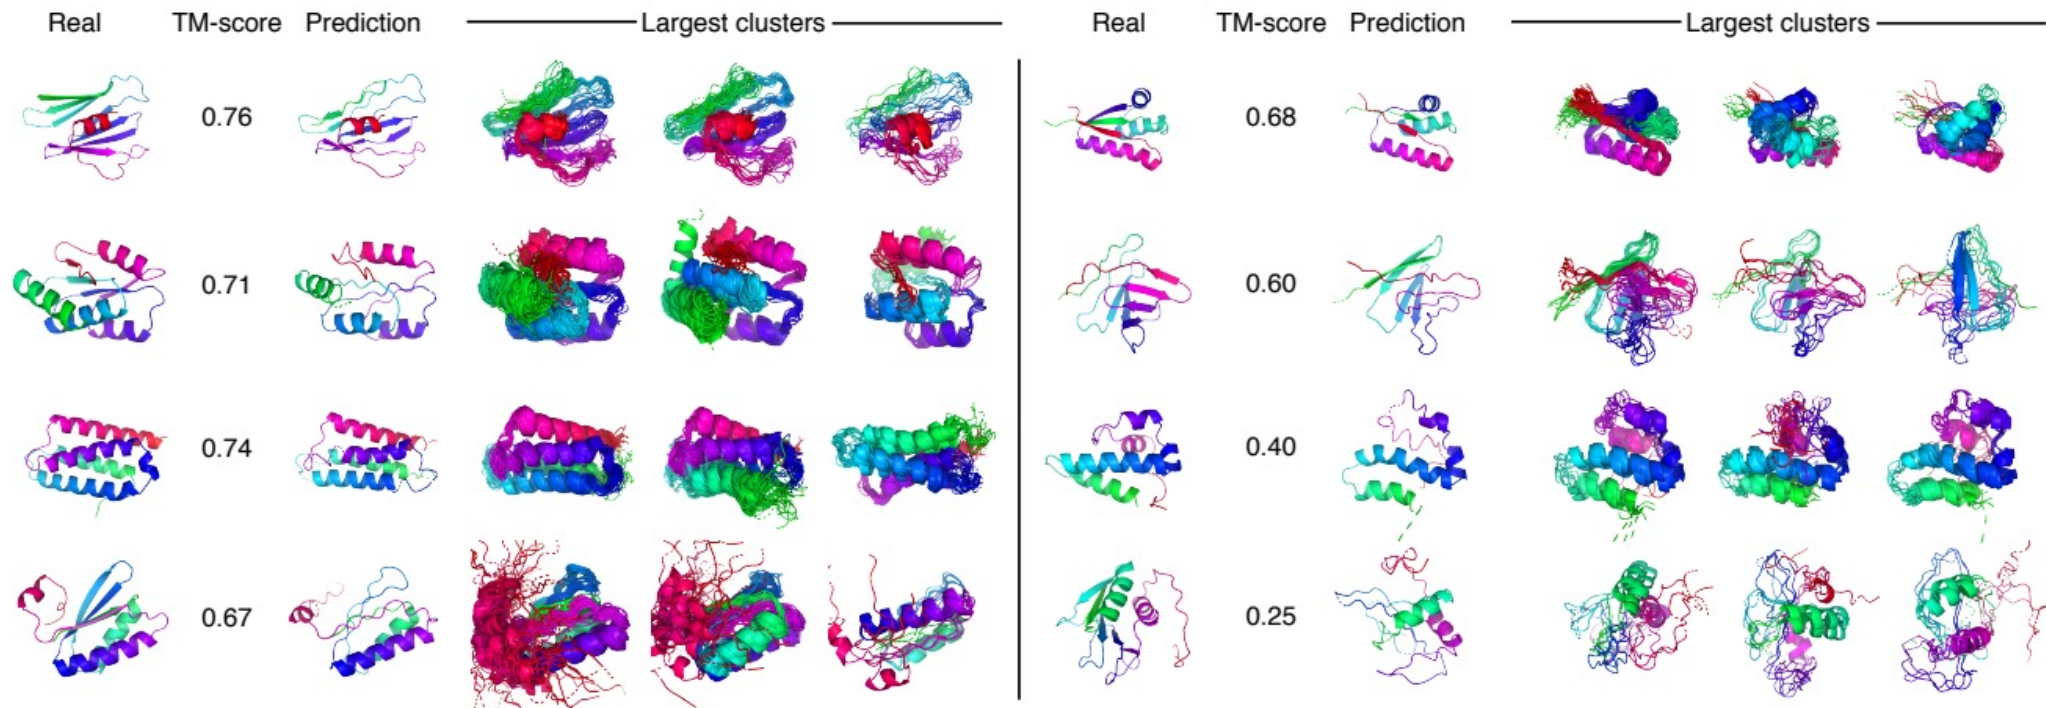
# One answer

Table 1: Test set performance across different levels of generalization

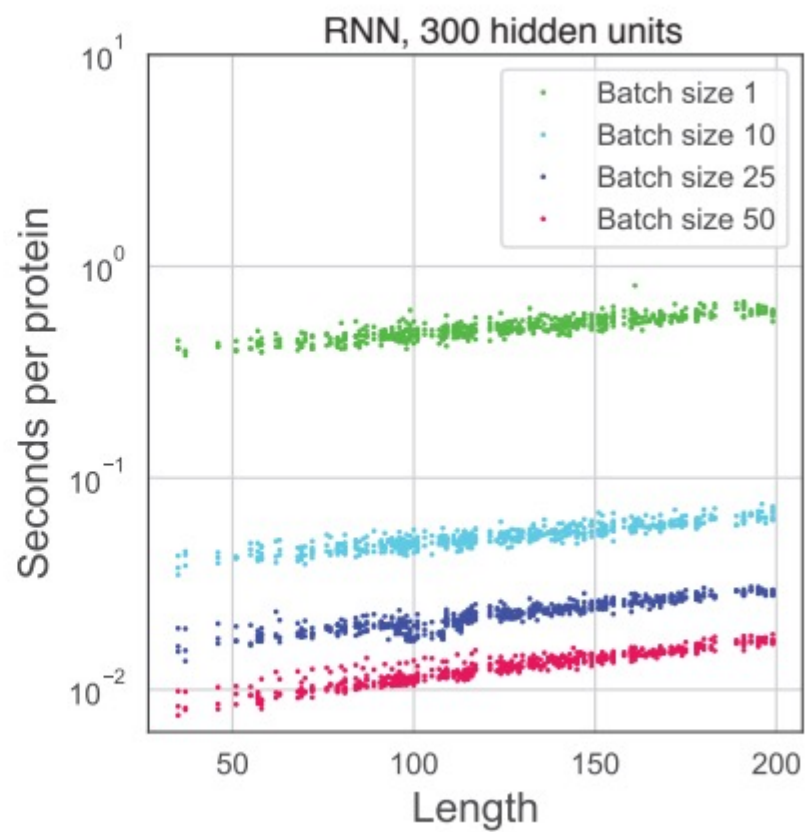
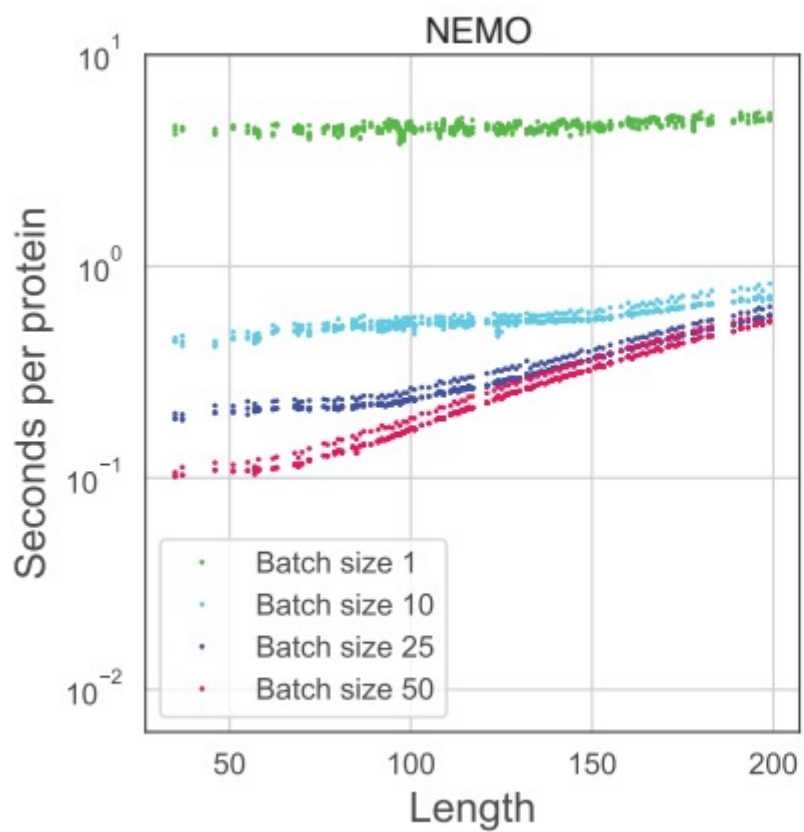
Model	# params	Total	C	A	T	H
NEMO (ours, profile)	21.3m	<b>0.366</b>	<b>0.274</b>	<b>0.361</b>	<b>0.331</b>	0.431
NEMO (ours, sequence-only)	19.1m	0.248	0.198	0.245	0.254	0.263
RNN baseline model (profile)						
2x100	5.9m	0.293	0.213	0.230	0.247	0.388
2x300 (avg. of 3)	8.8m	0.335	0.229	0.282	0.278	0.446
2x500	13.7m	0.347	0.222	0.272	0.286	<b>0.477</b>
2x700	21.4m	0.309	0.223	0.259	0.261	0.403
Number of structures		10381	1537	1705	3198	3941

# Confidence

Uncertainty measured by “clusters” – similar structures sampled in distribution



# Time



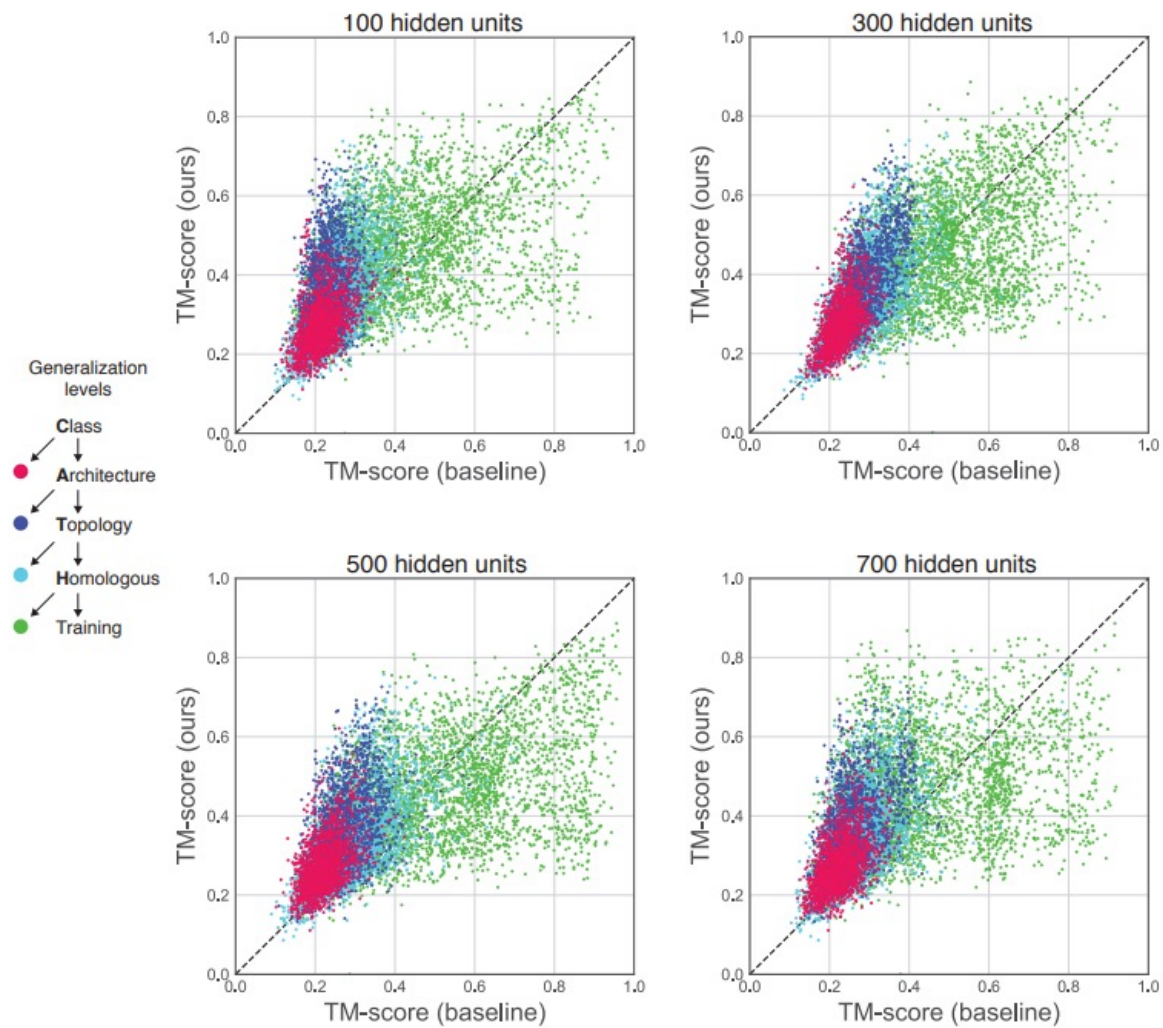


# One measurement metric

- TM-Score
  - Scored from 0 to 1
  - >0.5 is considered “good”
  - Best possible score across all possible positions:
    - Found by optimizing (with autograd throughout)

$$\sum_i \frac{1}{1 + \left(\frac{D_i}{D_0}\right)^2},$$

# Baselines



# loss.backward()

Any questions?