SE(3) Denoising Score Matching with Neural Euler's Rotation Equation

Towards unsupervised models for protein-ligand binding

Wengong Jin*, Sisi Sarkizova, Xun Chen, Nir Hacohen, Caroline Uhler

*Eric and Wendy Schmidt Center, Broad Institute

Protein-ligand binding affinity prediction

Applications of equivariant neural networks to drug discovery

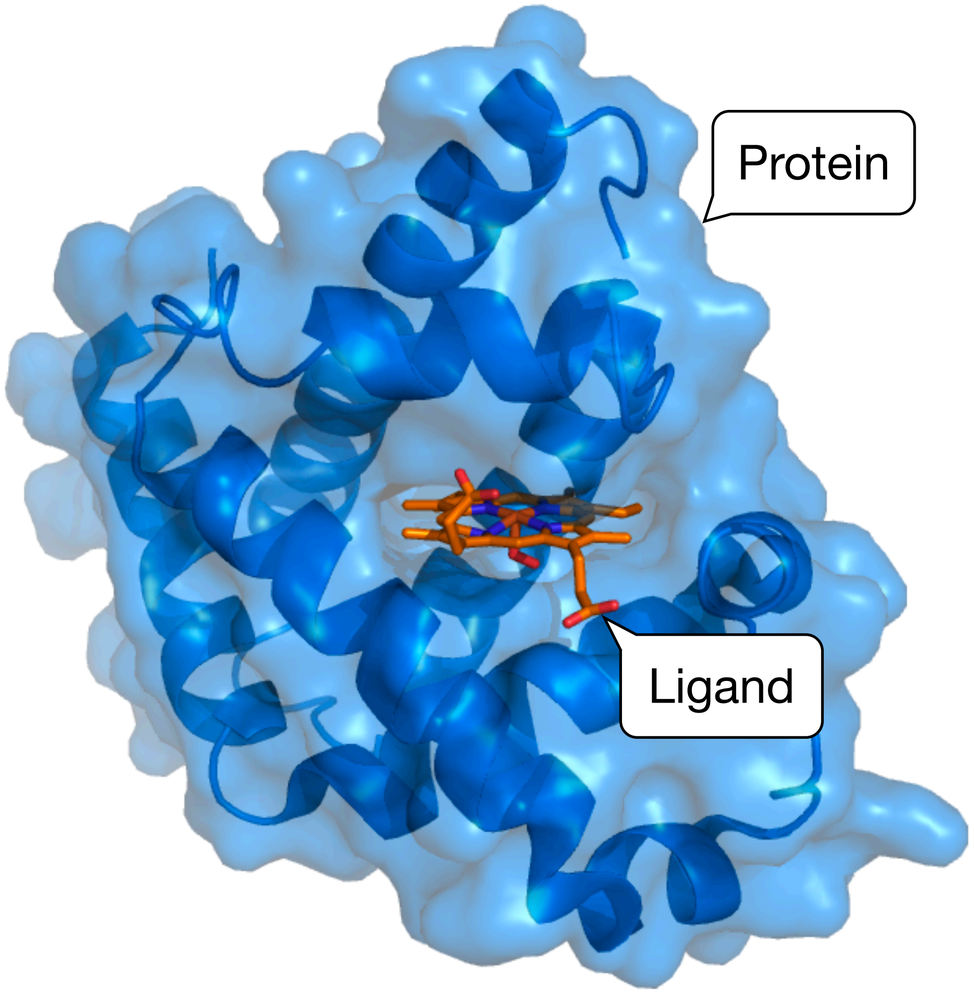

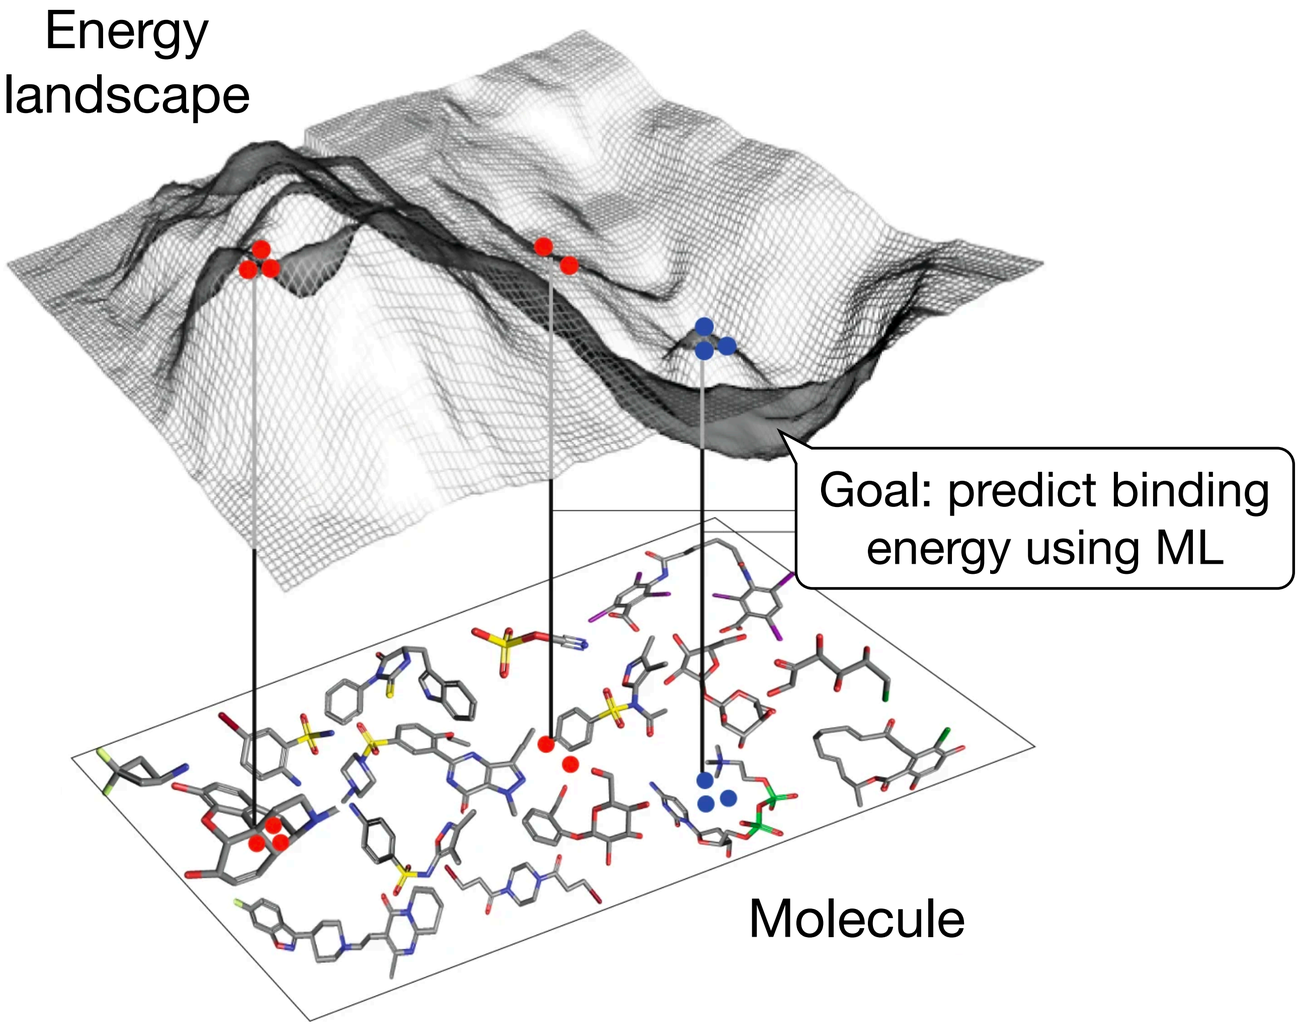

Figure: Gromski et al. 2019 & Wikipedia

Protein-ligand binding affinity prediction

A typical workflow: docking + scoring

Protein structure

Ligand structure

Docking

(e.g., AlphaFold, DiffDock)

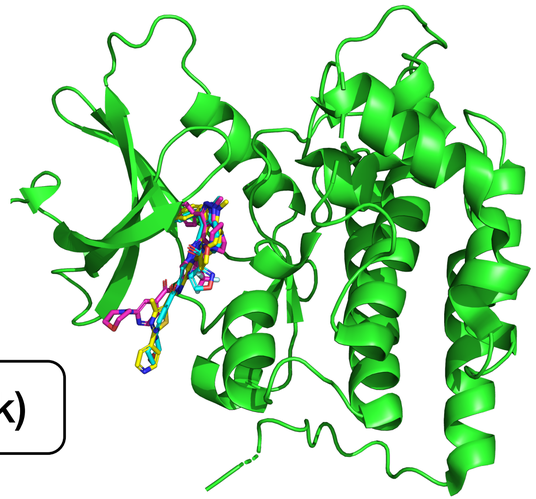

Scoring function

Today's focus

Binding affinity

Outline of this talk

1. Formulate binding affinity prediction as a generative modeling problem
2. Train the generative model using SE(3) denoising score matching (DSM)
3. Propose a **simple** & **equivariant** rotation prediction module for SE(3) DSM

Protein-ligand binding affinity prediction

Background: supervised & unsupervised models

Supervised models

Protein-ligand complex

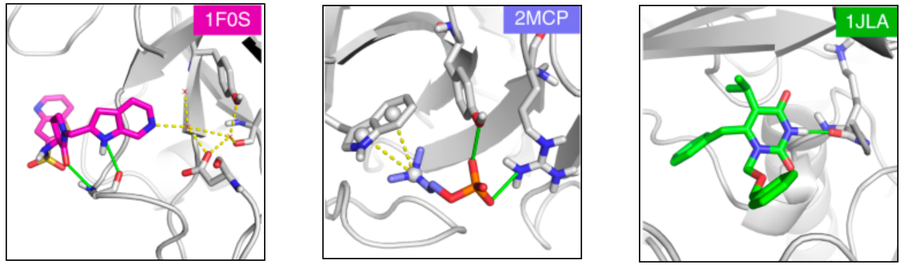

Affinity: Ki=18nM Ki=160nM Ki=1500nM

Training criteria: regression

Limitation: requires lots of binding affinity labels

Unsupervised models

Protein-ligand complex

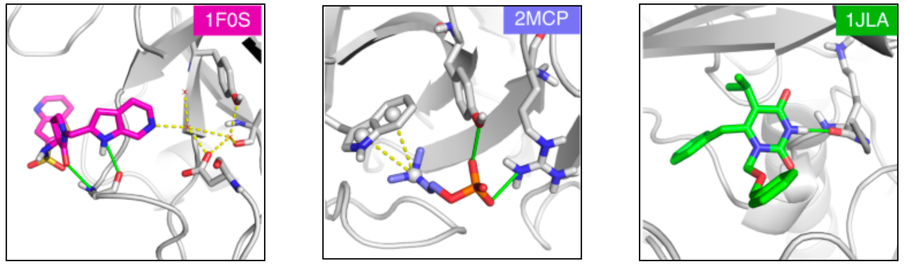

Training criteria: maximum likelihood

Benefit: works when binding affinity data is limited (e.g., antibodies)

Protein-ligand binding affinity prediction

Background: connection between binding energy and log-likelihood

- **Intuition**: if a protein-ligand complex has a strong binding affinity, it will appear more **often**
- The likelihood of a complex $p(x) \propto \exp(-E_\theta(x))$, where $E_\theta(x)$ is the energy of a complex
- Previous work (e.g., DrugScore2018 [1]) showed log-likelihood is correlated with binding energy
  - $\log p(x) = \sum_{i,j} \log p(D_{ij} = d_{ij})$, where d_{ij} is the distance between atom pair (i, j)
  - Model is not expressive due to factorization

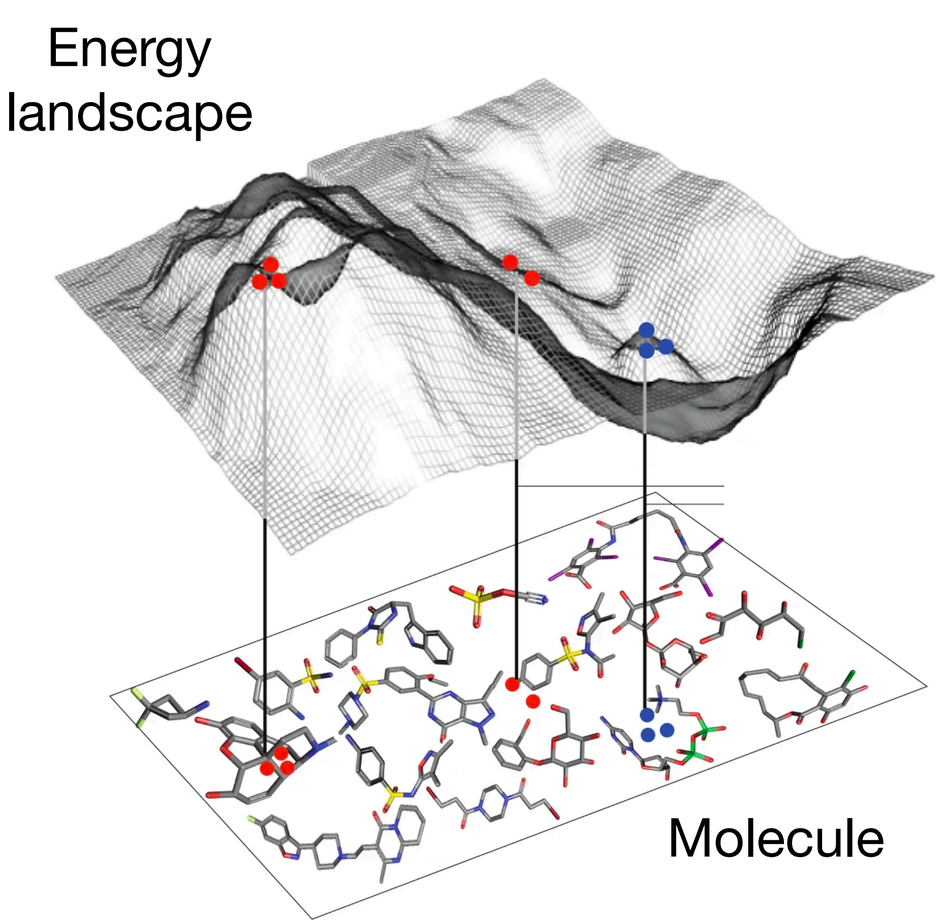

[1] Dittrich et al., "Converging a Knowledge-Based Scoring Function: DrugScore2018" *J. Chem. Inf. Model.* 2019, 59, 1, 509–521

Our approach: neural network energy models

Learn binding energy from crystal structures (data-driven)

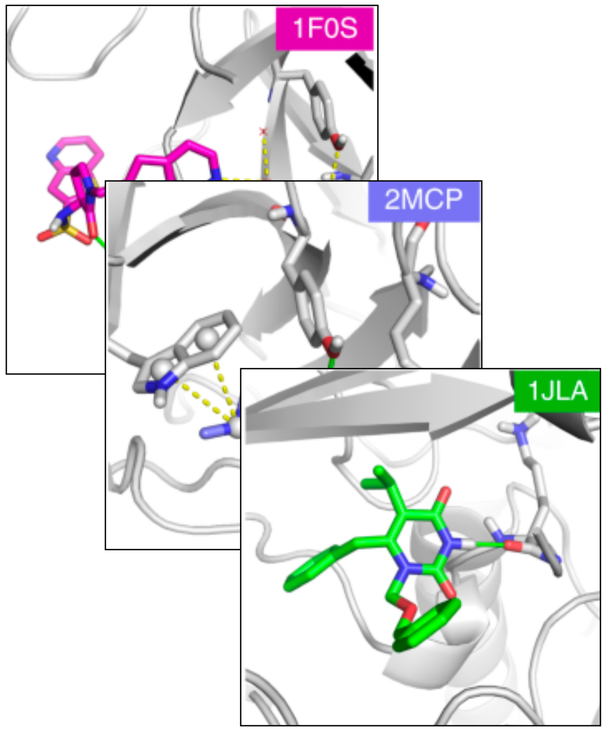

Protein-ligand 3D structures

Energy-based Model (EBM)

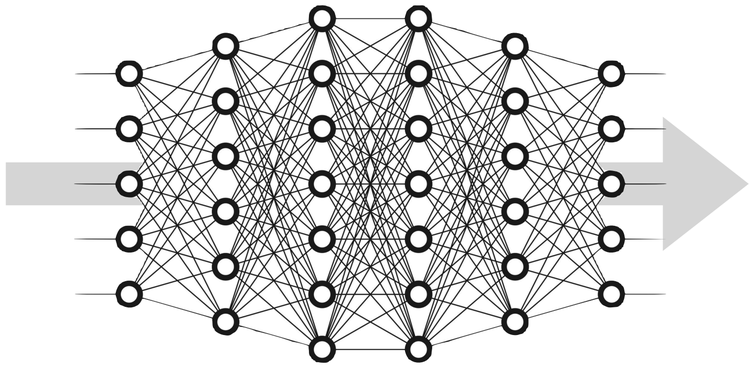

Energy landscape

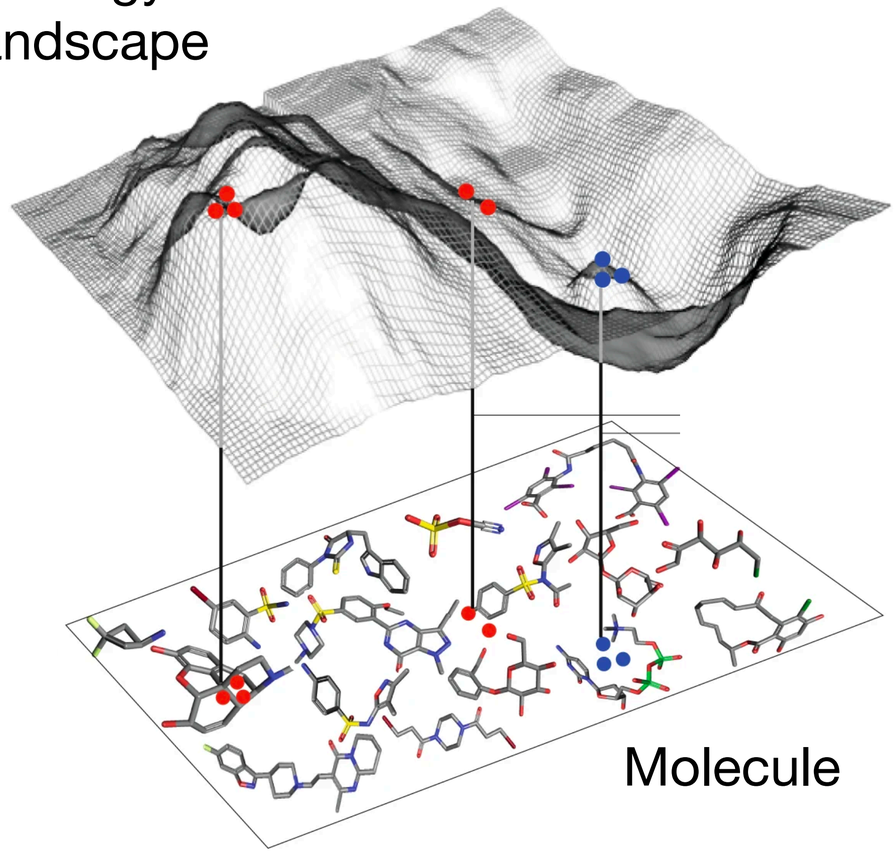

Molecule

EBM architecture for protein-ligand binding

Requirement: $E(X)$ is SE(3)-invariant and differentiable w.r.t. X

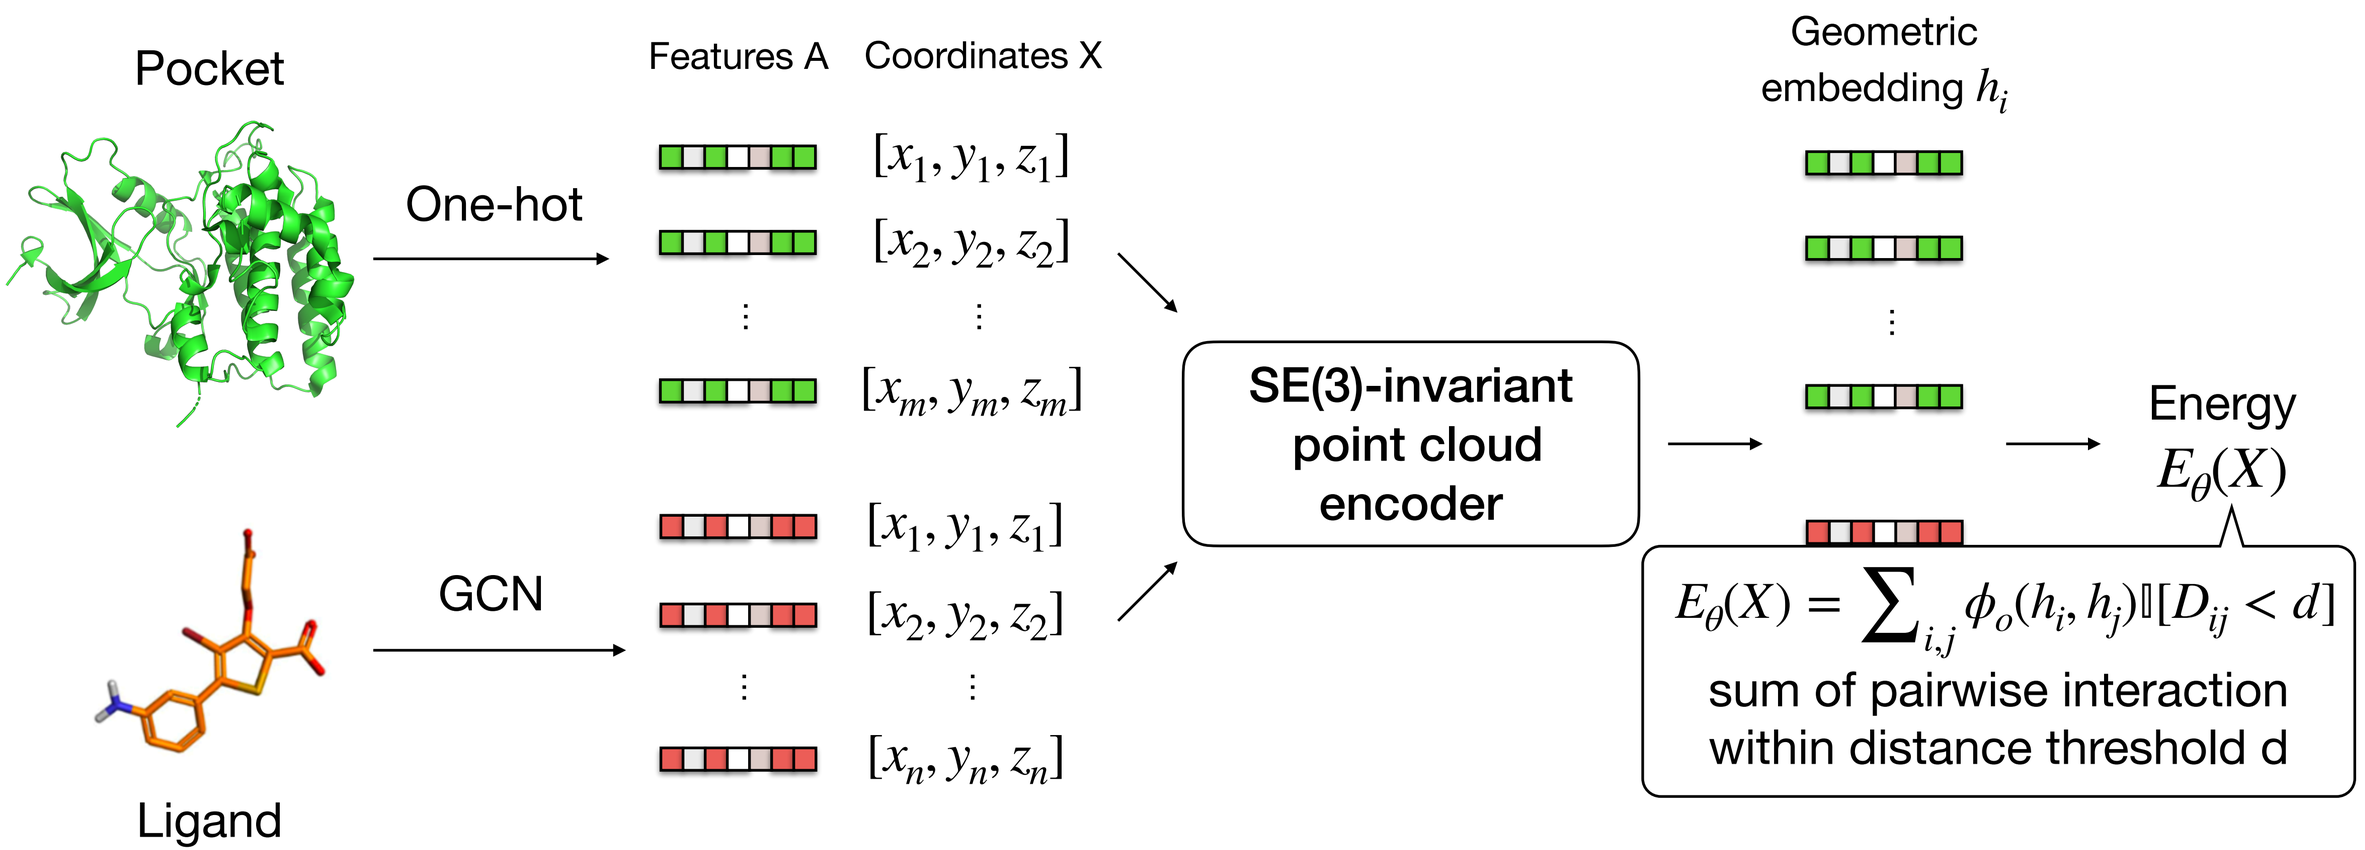

Training EBMs with denoising score matching

Standard approach: Gaussian noise

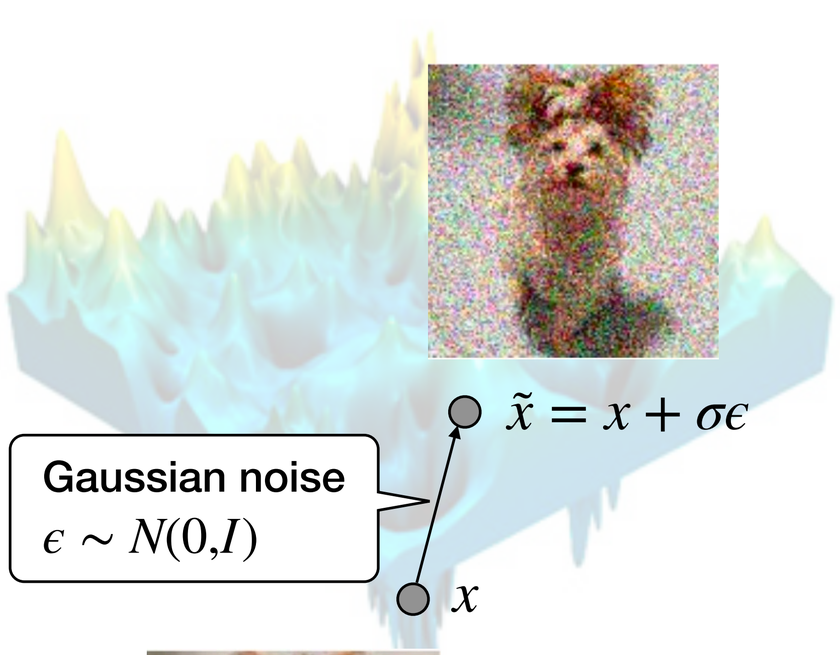

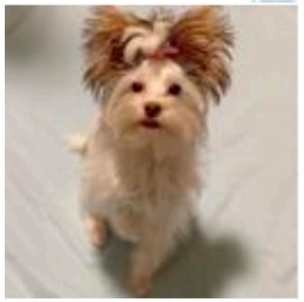

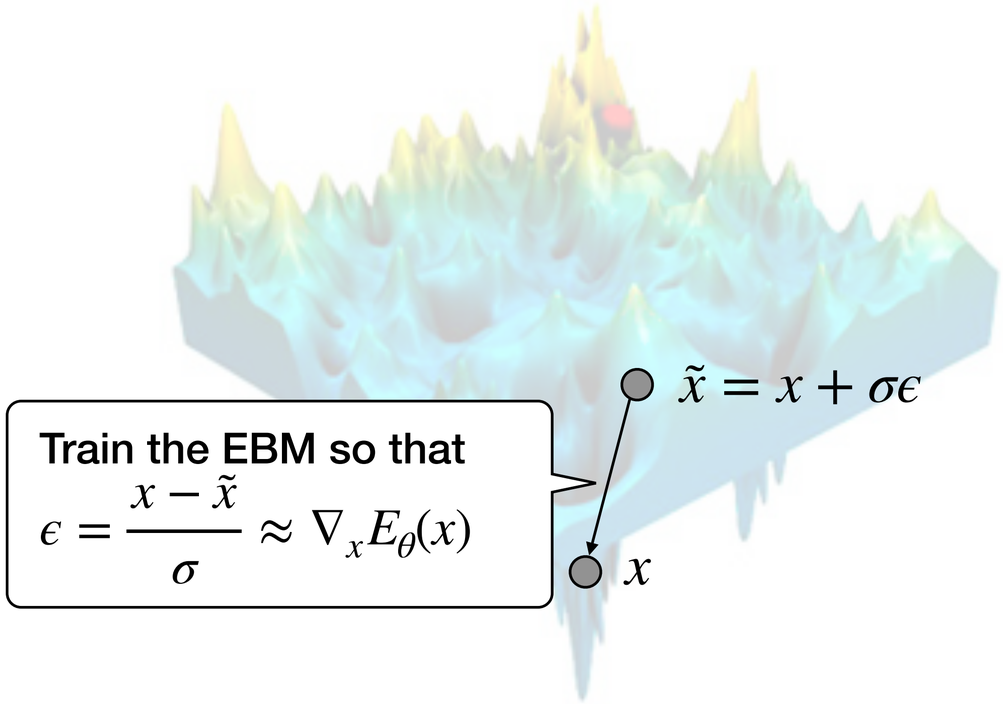

- Train EBMs with denoising score matching: $\|\nabla_x E_\theta(x) - \epsilon\|$
- By shaping the gradient $\nabla_x E_\theta(x)$, we can recover the true energy function (up to affine transformation)

Training EBMs with denoising score matching

For molecules, we should use rigid transformation noises

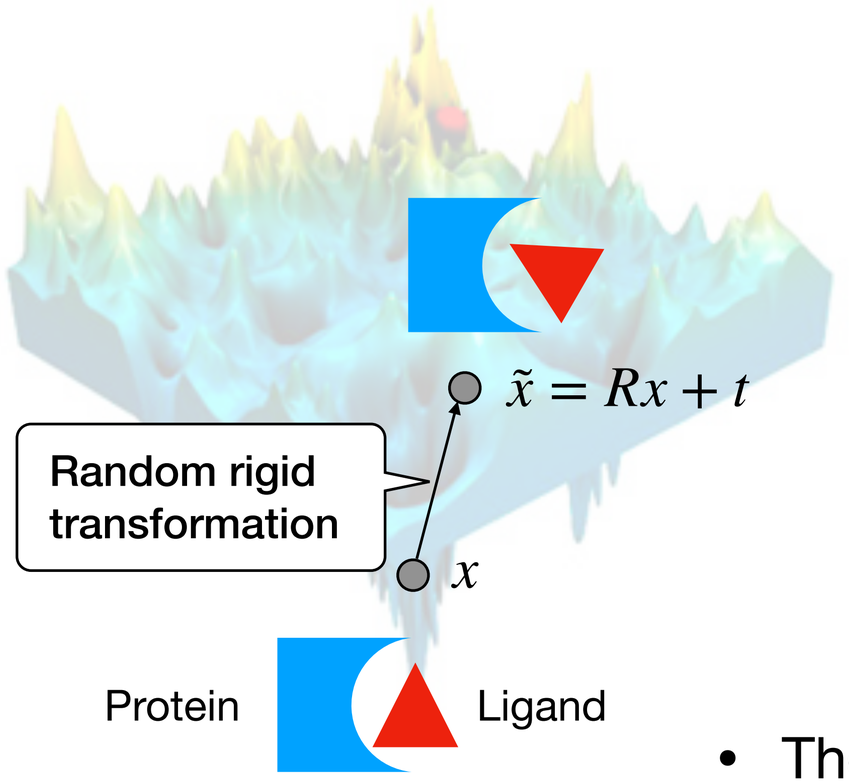

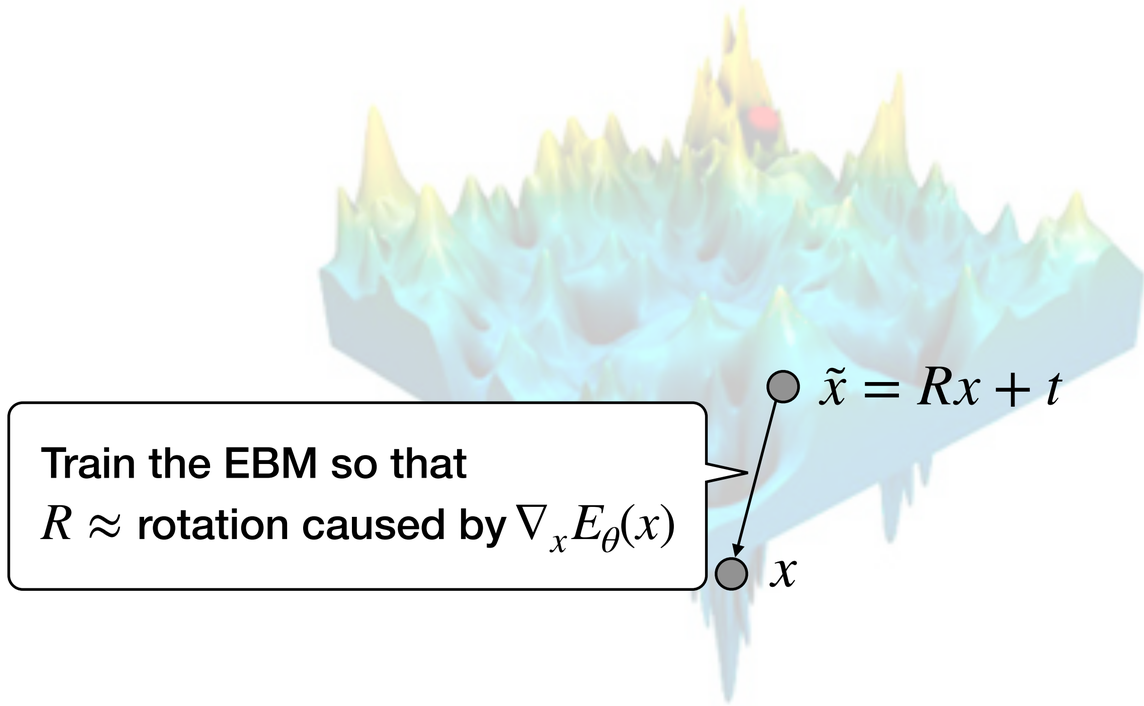

- The key step in SE(3) denoising score matching is to infer the rotation induced by the score $\nabla_x E_\theta(x)$

Euler's rotation equation

Infer rotation R from gradient $\nabla_x E_\theta(x)$ (force)

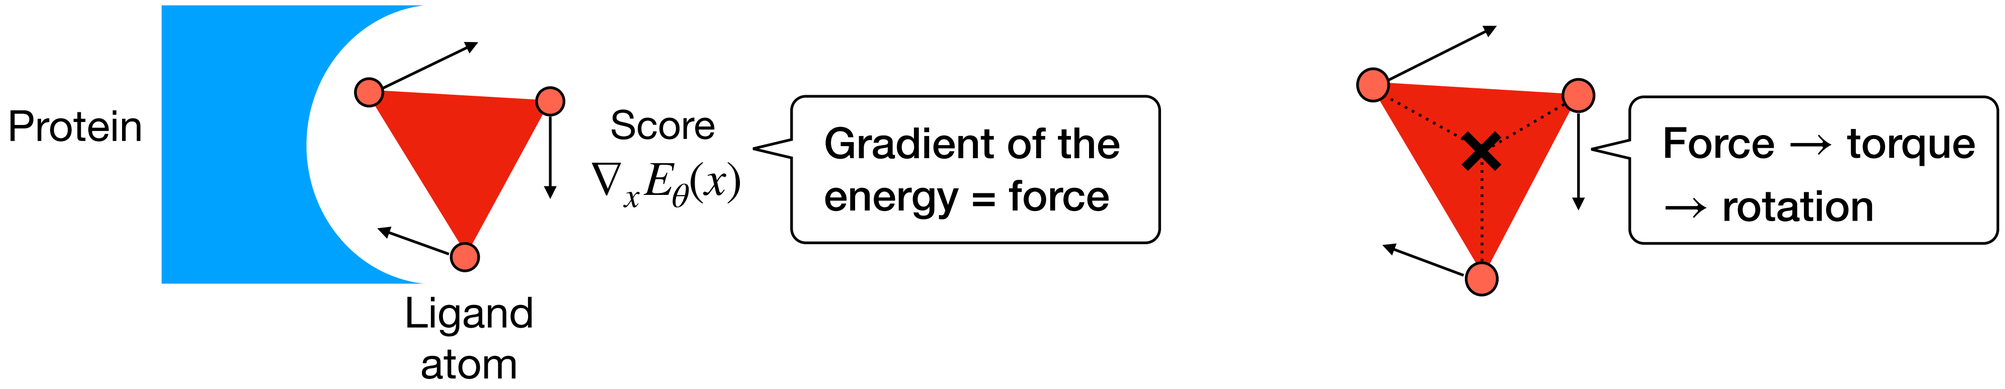

- The torque applied to the ligand $\tau = \sum_i (x_i - \mu) \times \nabla_{x_i} E_\theta(x)$
- (**Euler's rotation equation**) Angular acceleration of the ligand $\alpha = I^{-1}\tau$, where I is the inertia matrix
- Angular velocity $\omega = I^{-1}\tau\Delta t$ for an infinitesimal time Δt
- Rotation matrix R is the exponential of the following matrix $W(\omega) = \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix}$

Euler's rotation equation

Infer rotation R from gradient $\nabla_x E_\theta(x)$ (force)

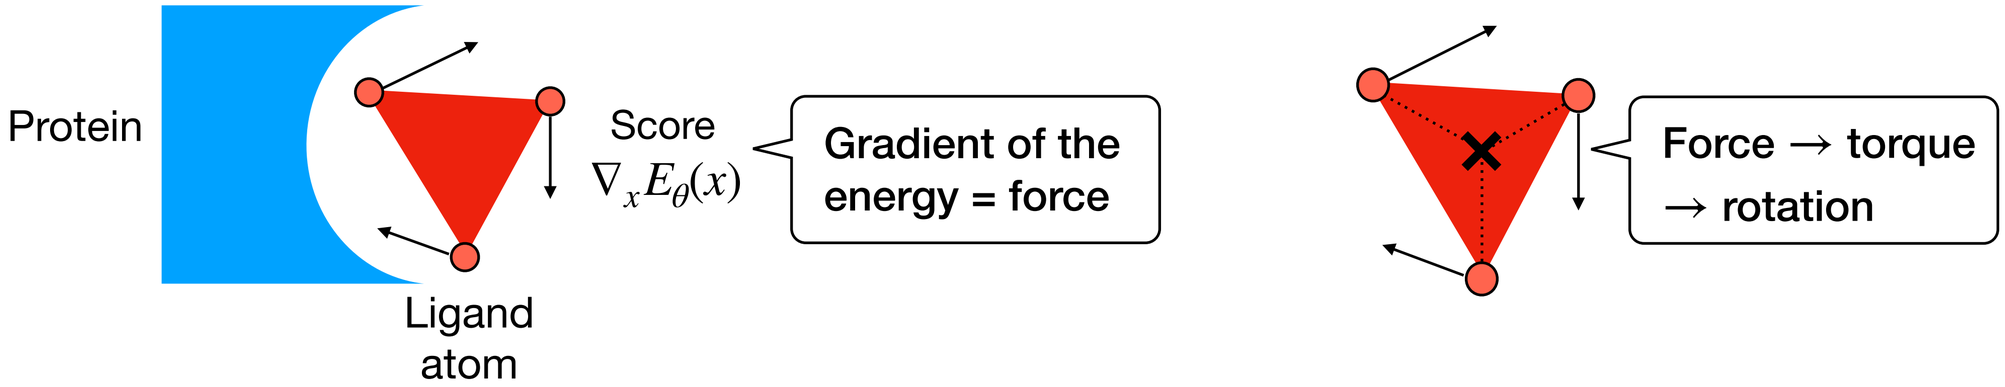

- The good news is that $W(\omega)$ is a skew symmetric matrix and its exponential has a closed form
- The rotation operation can be further simplified to $Rx = e^{W(\omega)}x = x + c_1\omega \times x + c_2\omega \times (\omega \times x)$, where $c_1 = \sin \|\omega\|/\|\omega\|$ and $c_2 = (1 - \cos \|\omega\|)/\|\omega\|^2$
- The above rotation formula only requires vector cross product, which is very efficient.

Euler's rotation equation as a rotation layer

A simple, yet equivariant way to predict rotations

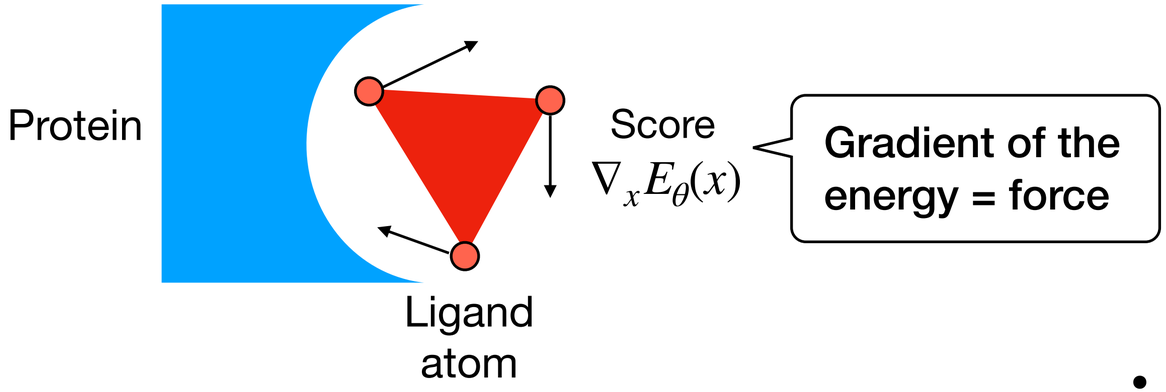

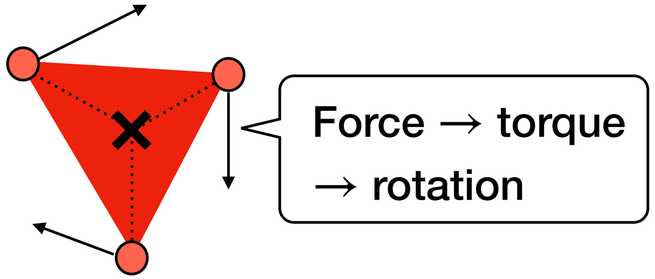

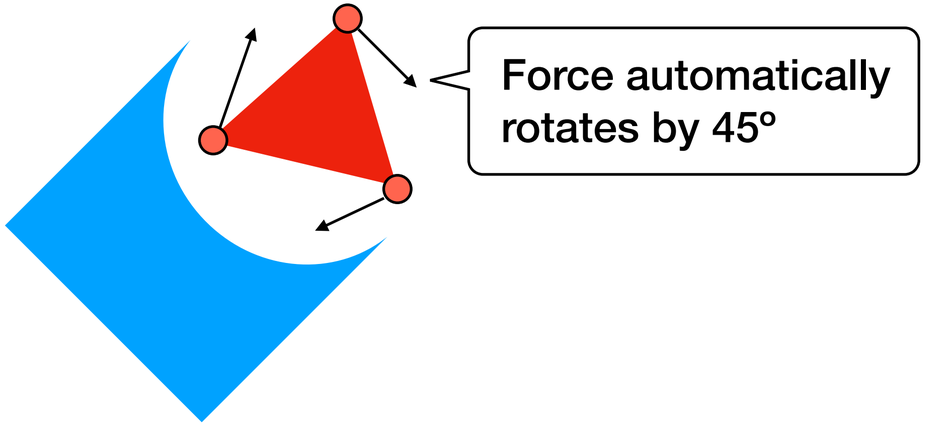

- If we rotate a complex by Q, we have:
- Torque $\tau_Q = \sum_i (Qx_i - Q\mu) \times Q \nabla_{x_i} E_\theta(x) = Q\tau$
- Inertia matrix $I_Q = QIQ^{-1}$
- Angular velocity $\omega_Q = I_Q^{-1} \tau_Q \Delta t = Q\omega$
- Rotation matrix $R_Q = QR$

SE(3) denoising score matching

Training procedure

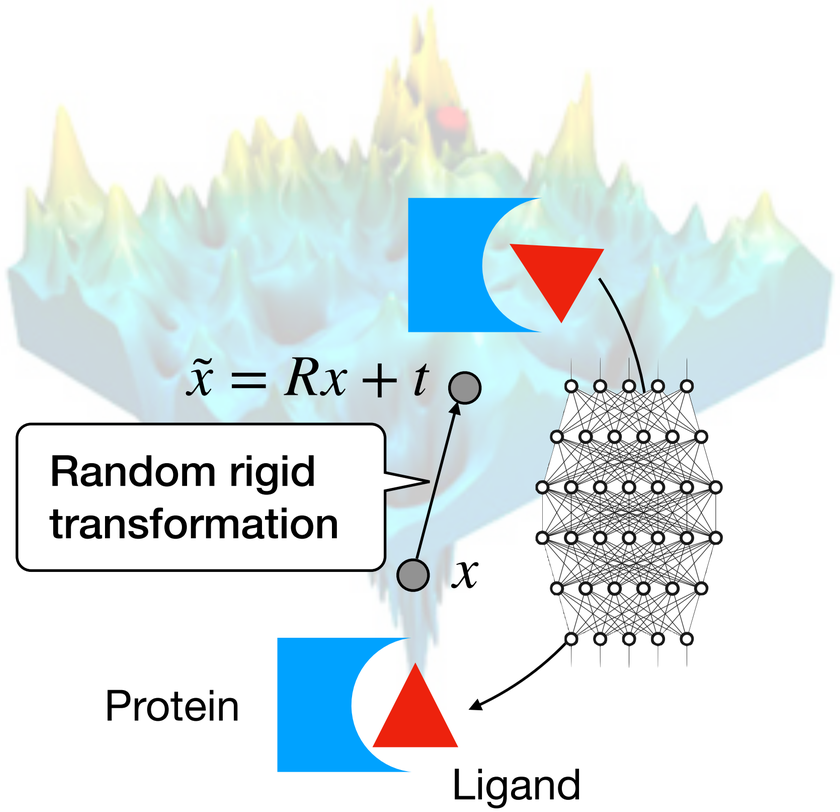

- Step 1: Sample rotation from SO(3) Gaussian distribution $\mathcal{N}_{\mathrm{SO}(3)}$
  - Sample a random direction e from unit sphere
  - Sample an angle θ with density $f(\theta)$ (Isotropic Gaussian)
  - The score of $\mathcal{N}_{\mathrm{SO}(3)}$ is $\nabla_\theta \log f(\theta) \cdot e$
- Step 2: Calculate the energy $E_\theta(x)$ and its score $\nabla_x E_\theta(x)$ (force)
  - Calculate the torque τ
  - Infer angular velocity $\omega = I^{-1}\tau\Delta t$
- Step 3: Compute SE(3) DSM Loss
  - $\|\omega - \nabla_\theta \log f(\theta) \cdot e\|$

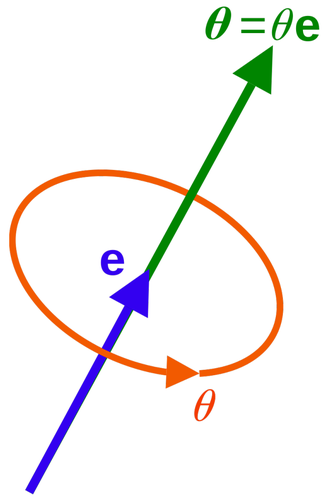

Results: protein-ligand binding

Log-likelihood is strongly correlated with binding affinity

- Training set: 5237 protein-ligand complexes in PDBBind refined set (without using binding affinity data)
- Test set: 285 complexes from CASF challenge evaluation set [1]. Measure the Pearson correlation between predicted and true affinity
- Supervised models are trained on ~18000 binding affinity data in PDBBind
- SE(3) DSM outperforms Gaussian noise DSM and other unsupervised models

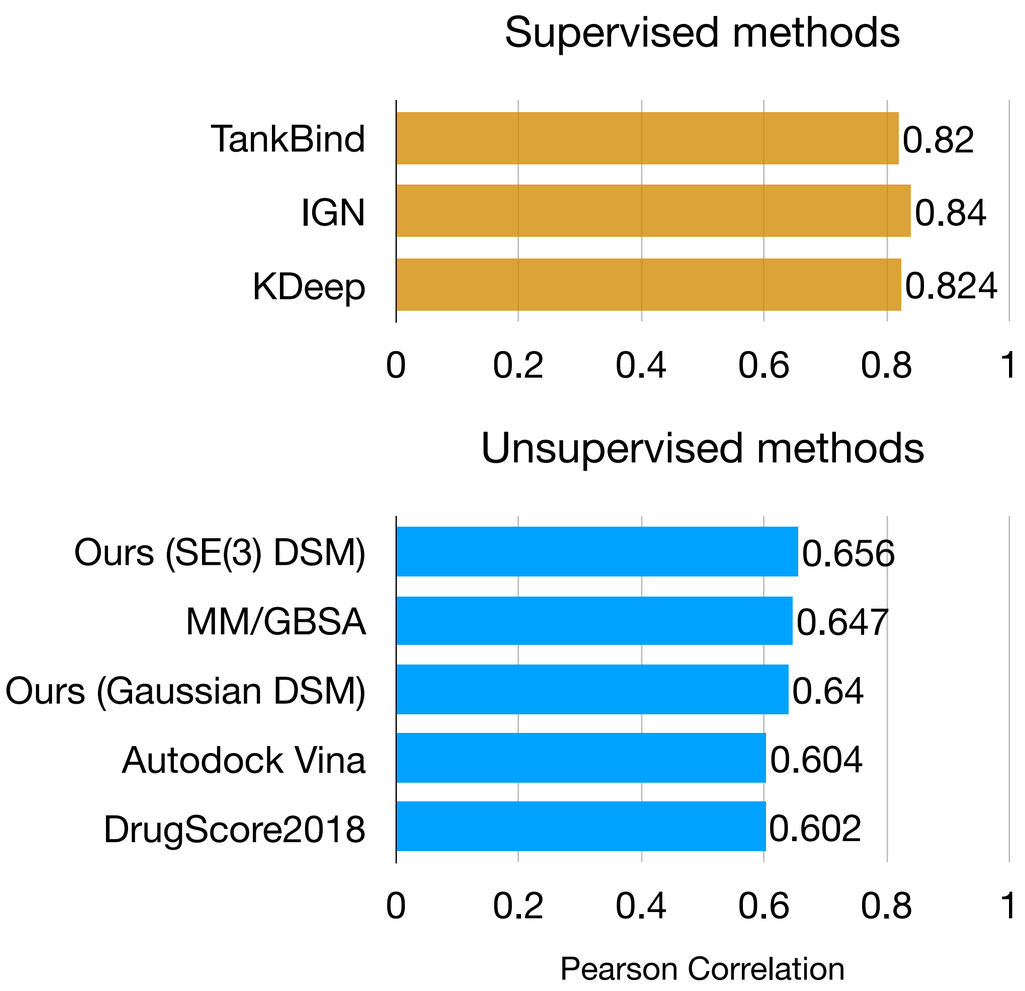

Results: antibody-antigen binding

Supervised models suffer from lack of binding affinity data

- Training set: 3416 complexes from Structure Antibody Database (SAbDab).
- Test set: 566 complexes from SAbDab that have binding affinity labels
- We compare with several biophysical potentials (unsupervised), and a supervised neural network trained on only 100 binding affinity data points
- We outperform supervised baseline because we can leverage more unlabeled antibody-antigen complexes

	Crystal
ZRANK	0.318
ZRANK2	0.176
RosettaDOCK	0.064
PYDOCK	0.248
SIPPER	-0.138
AP_PISA	0.323
FIREDOCK	0.101
FIREDOCK_AB	0.199
CP_PIE	0.234
NERE (ours)	$\mathbf{0.340}_{.029}$
- standard DSM	$0.335_{.038}$
Supervised NN	$0.295_{.098}$

Visualization: binding energy attention

What residues contribute the most?

- $E_\theta(x) = \sum_{i,j} \phi_o(h_i, h_j) \mathbb{I}[D_{ij} < d]$
- $\phi_o(h_i, h_j)$ is the binding energy between two residues
- We plot $\phi_o(h_i, h_j)$ for all pairs within the distance threshold d
- Interestingly, the model pays the most attention to CDR-H3 and CDR-L3 residues, which are most critical to binding.

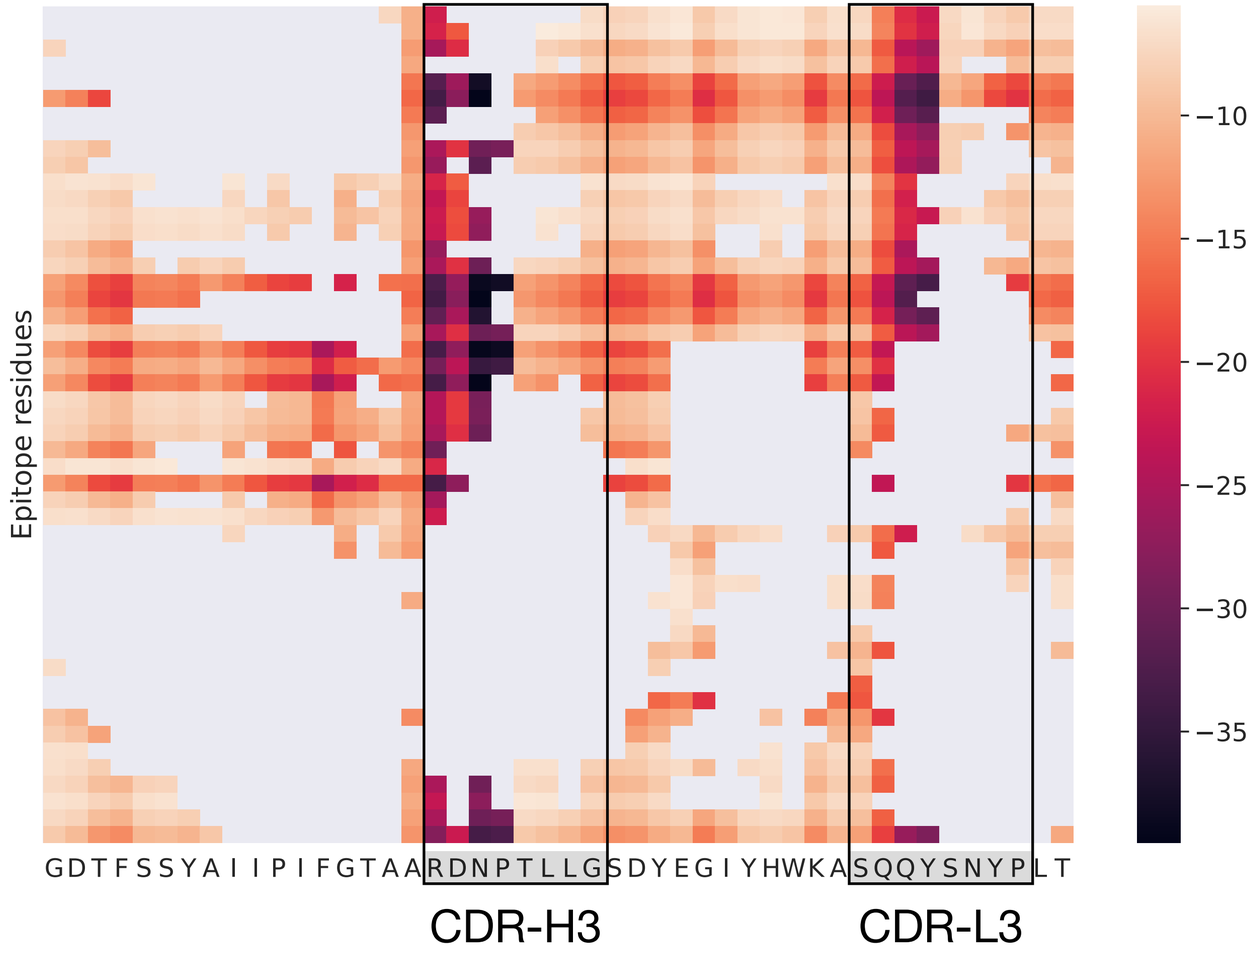

Visualization: binding energy landscape

How binding energy changes with respect to ligand orientation?

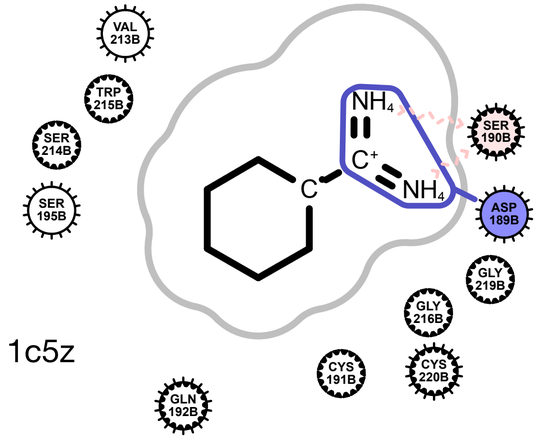

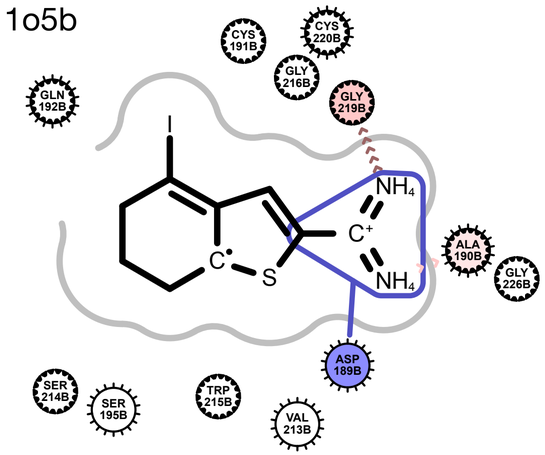

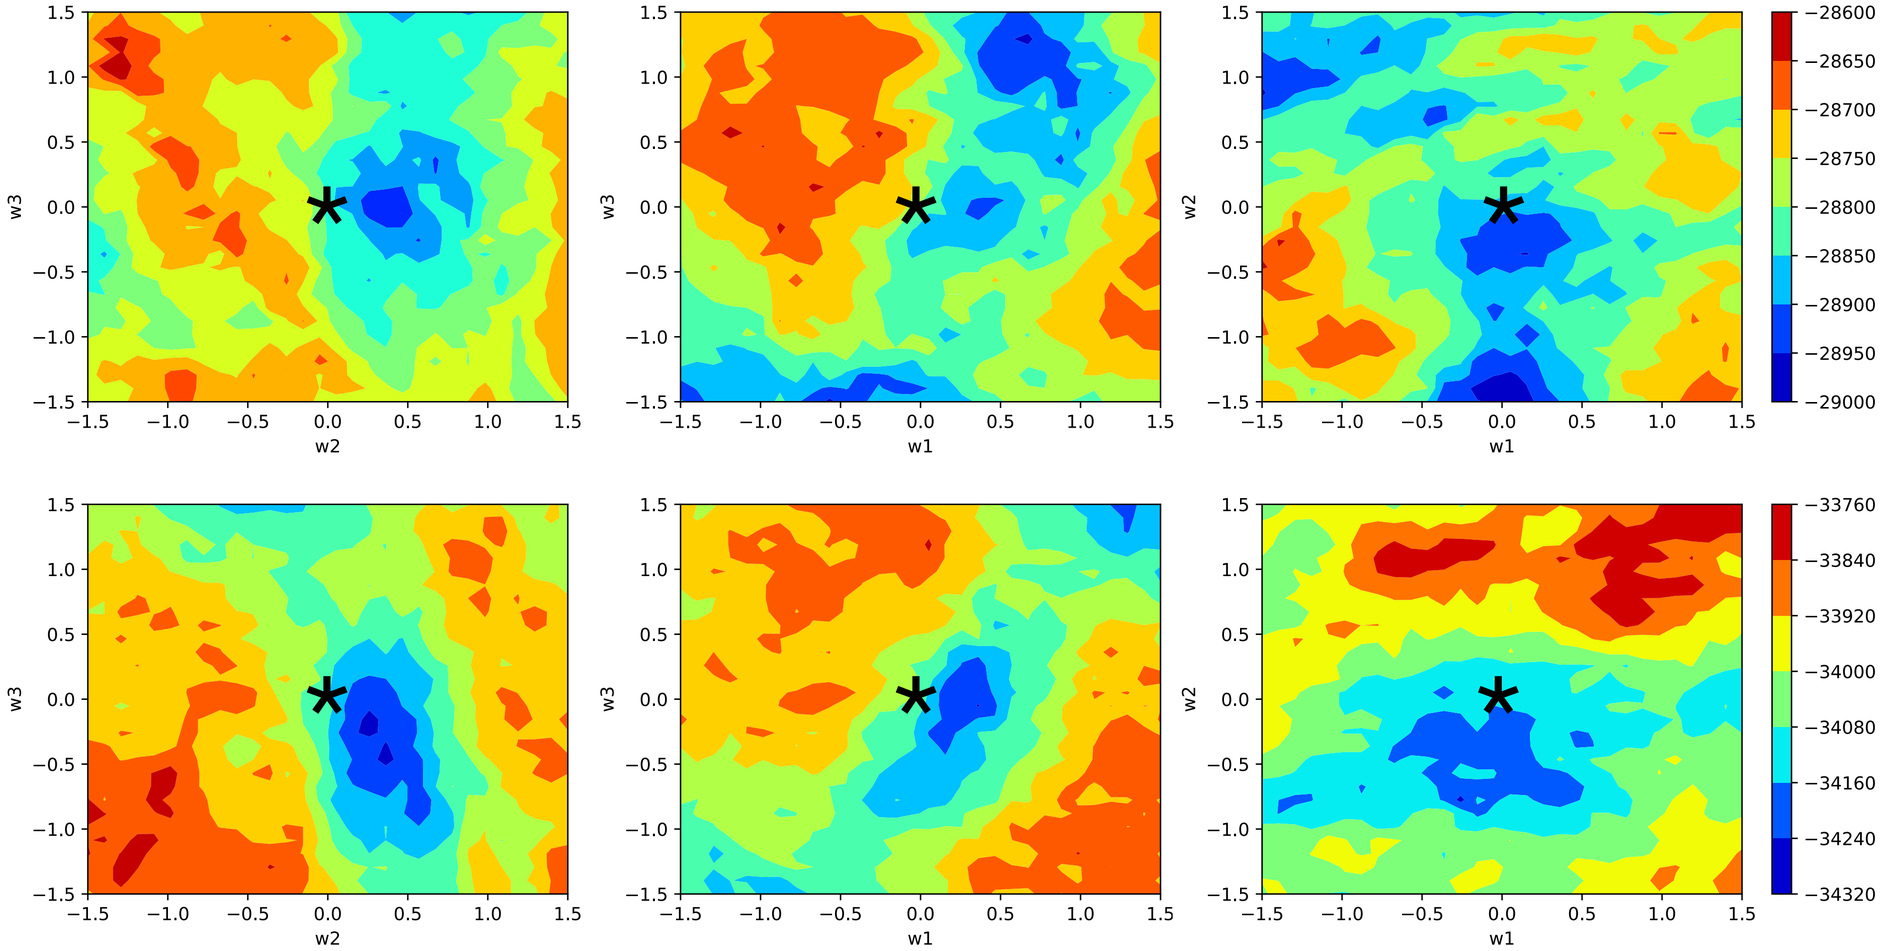

Conclusion & acknowledgements

Towards unsupervised models for protein-ligand binding

Main contribution

1. Formulate binding affinity prediction as a generative modeling problem
   - Train the generative model using SE(3) denoising score matching (DSM)
2. Propose a **simple** & **equivariant** rotation prediction module for SE(3) DSM
   - Embed Euler's rotation equation into neural networks (adding physical prior)

Acknowledgements:

Funding: Eric and Wendy Schmidt Center

Collaborators: Sisi Sarkizova, Xun Chen, Nir Hacohen, Caroline Uhler