Introducing Physics into Representation Learning

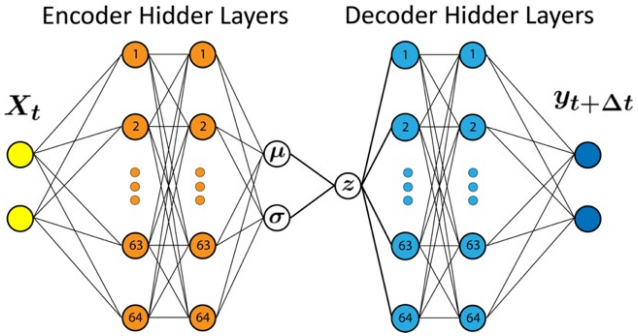

Dedi Wang

University of Maryland, College Park

Advisor: Pratyush Tiwary

@tiwarylab

D. Wang, P. Tiwary *J. Chem. Phys. 2021*

D Wang, Y Wang, Evans & Tiwary arXiv:2209.00905

Dimension reduction can be challenging

Looks like a transition state?

2D -> 1D

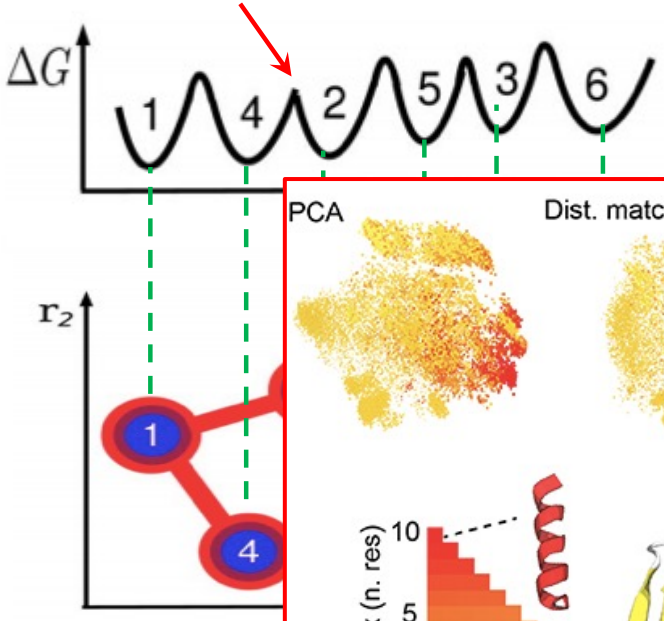

Why should you trust these projections?

N-D -> 2D

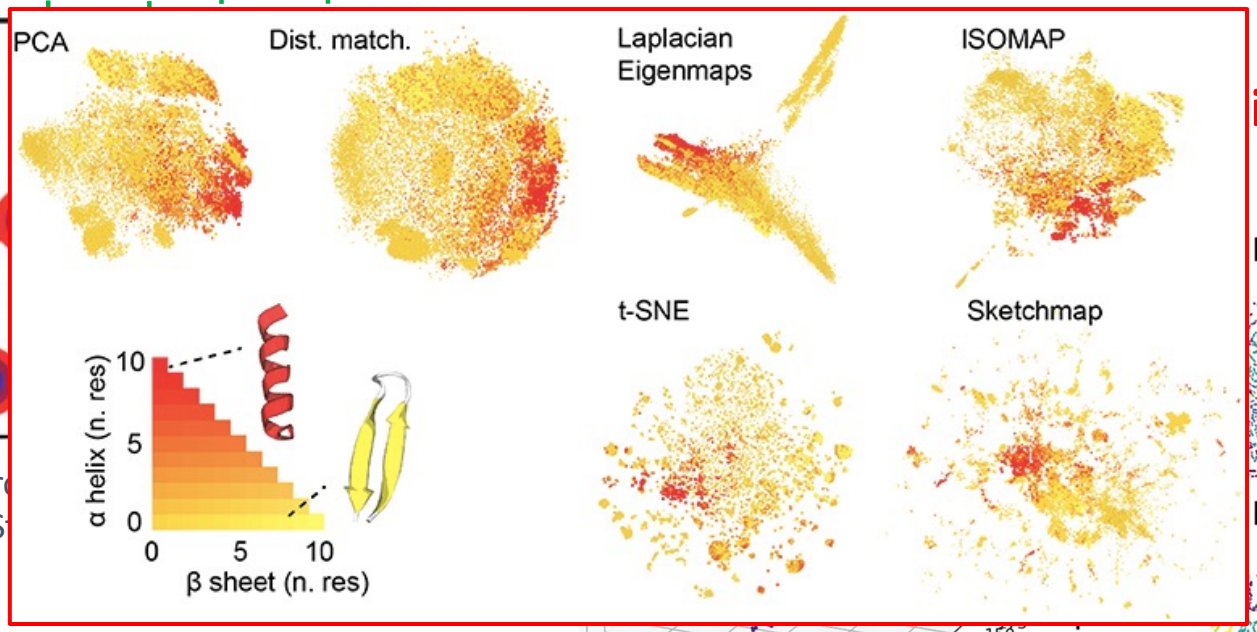

3D -> 2D

it represents?

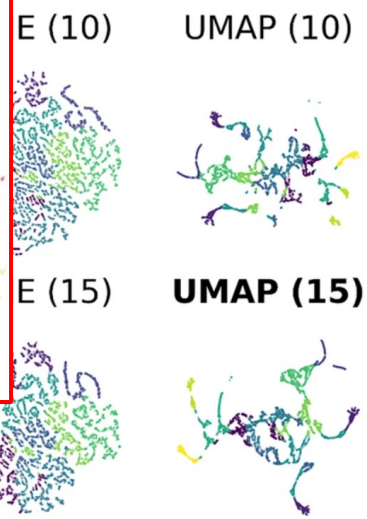

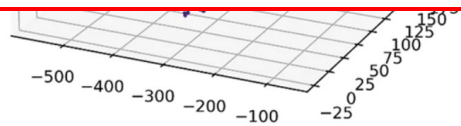

Tribello & Gasparotto Frontiers in Mol. Bio. 2019

Huang H, et al. Communications biology. 2022

Learning useful and meaningful low-dimensional representations via AI + physics

- **State Predictive Information Bottleneck (SPIB)**

  D Wang & Tiwary, J. Chem. Phys. 2021

- **Dynamics Constrained Autoencoder (DynAE)**

  D Wang, Y Wang, Evans & Tiwary, arXiv:2209.00905

Learning useful and meaningful low-dimensional representations via AI + physics

- **State Predictive Information Bottleneck (SPIB)**

  D Wang & Tiwary, J. Chem. Phys. 2021

Introducing "metastable states" into representation learning

1. Identifying metastable states

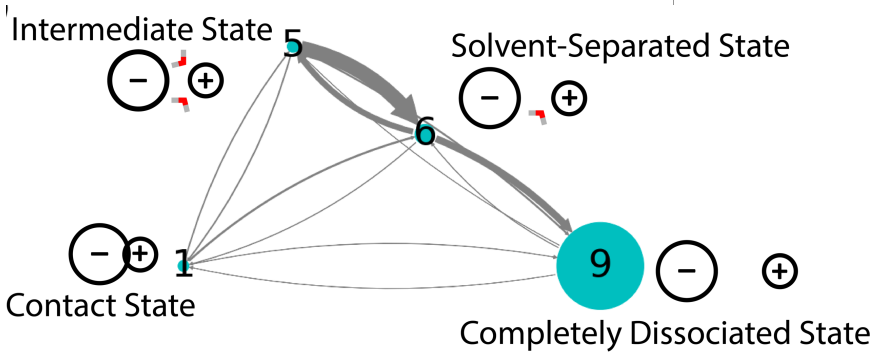

D Wang, Zhao, Weeks, Tiwary. JPCB. 2022

2. Predicting transition states

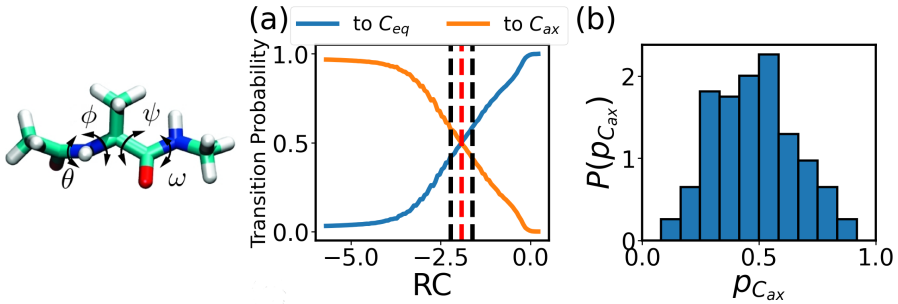

D Wang & Tiwary, J. Chem. Phys. 2021

3. Enhancing sampling

SPIB + metadynamics

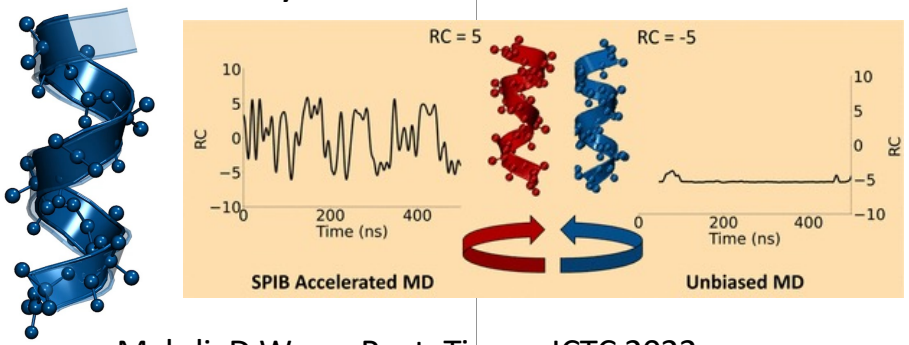

Mehdi, D Wang, Pant, Tiwary JCTC 2022

SPIB + weighted ensemble

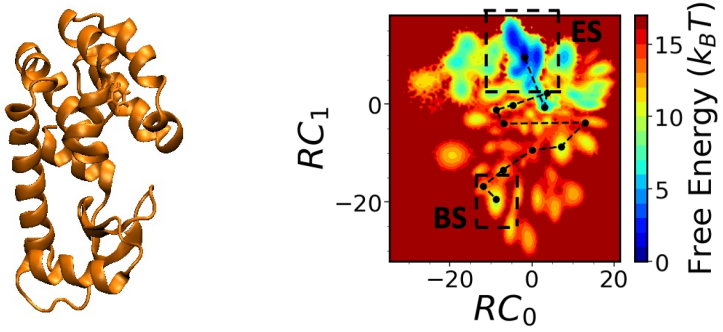

D Wang & Tiwary (under preparation)

Learning useful and meaningful low-dimensional representations via AI + physics

- **State Predictive Information Bottleneck (SPIB)**

  D Wang & Tiwary, J. Chem. Phys. 2021

- **Dynamics Constrained Autoencoder (DynAE)**

  D Wang, Y Wang, Evans & Tiwary, arXiv:2209.00905

Introducing the law of dynamics into representation learning

1. Recovering kinetics from distorted data

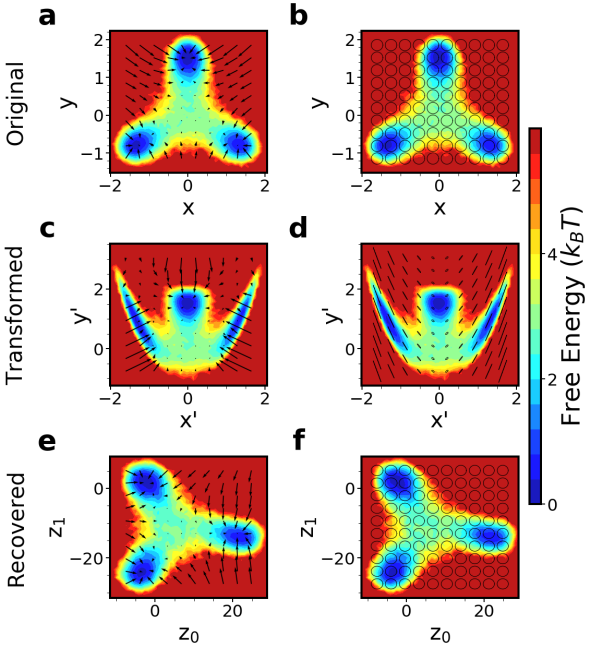

2. Learning dihedrals from atomic coordinates

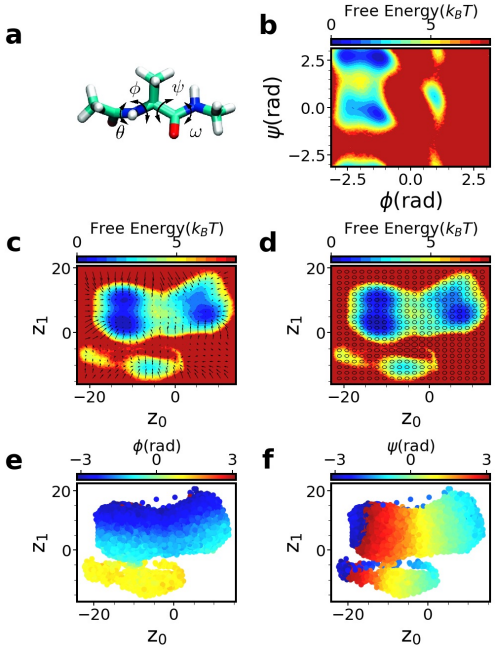

3. Interpreting experimental data

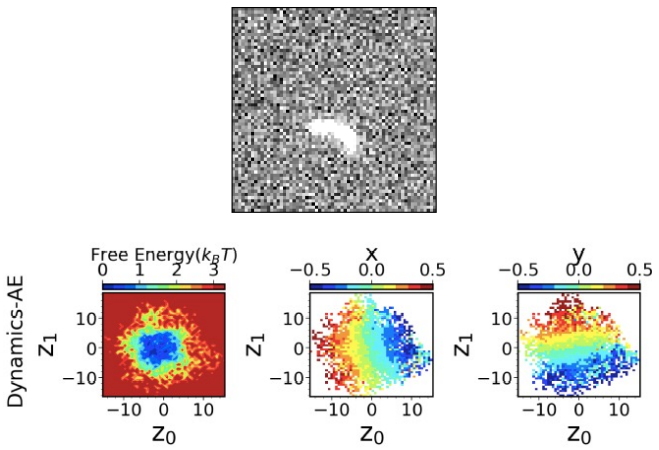

Learning useful and meaningful low-dimensional representations via AI + physics

- **State Predictive Information Bottleneck (SPIB)**

  D Wang & Tiwary, J. Chem. Phys. 2021

- **Dynamics Constrained Autoencoder (DynAE)**

  D Wang, Y Wang, Evans & Tiwary, arXiv:2209.00905

Information bottleneck - an approach to model human and artificial intelligence

Input

$\mathbf{X} \in \mathbb{R}^N$

Encoder

Bottleneck

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

Decoder

Target

$\mathbf{y} \in \mathbb{R}^D$

"Guitar"

Picture from Ribeiro, Singh, Guestrin. 2016

Rate Distortion Theory (Shannon); Information bottleneck (Tishby, Pereira, Bialek 1999)
Wang, Ribeiro, Tiwary (Nat. Comm. 2019); D. Wang, Tiwary (J. Chem. Phys. 2021)

Information bottleneck - an approach to model human and artificial intelligence

Input

$\mathbf{X} \in \mathbb{R}^N$

Encoder

Bottleneck

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

Decoder

Target

$\mathbf{y} \in \mathbb{R}^D$

"Dog"

Picture from Ribeiro, Singh, Guestrin. 2016

Rate Distortion Theory (Shannon); Information bottleneck (Tishby, Pereira, Bialek 1999)
Wang, Ribeiro, Tiwary (Nat. Comm. 2019); D. Wang, Tiwary (J. Chem. Phys. 2021)

Information bottleneck - an approach to model human and artificial intelligence

Input

$\mathbf{X} \in \mathbb{R}^N$

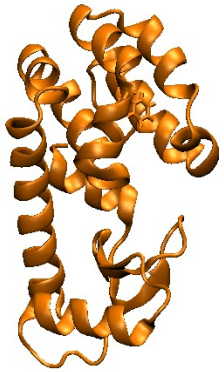

Bottleneck

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

Large variance

Encoder

Slow large variance

Small fluctuations

Filter out

Fast processes

Time delay Δt

Decoder

Target

$\mathbf{y} \in \mathbb{R}^D$

$\mathbf{y} = \mathbf{X}$

$\mathbf{y} = \mathbf{X}_{\Delta t}$

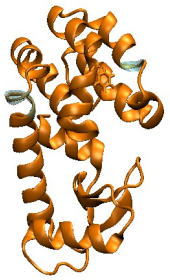

Too expensive to decode the full configuration

Is it necessary?

Rate Distortion Theory (Shannon); Information bottleneck (Tishby, Pereira, Bialek 1999)
Wang, Ribeiro, Tiwary (Nat. Comm. 2019); D. Wang, Tiwary (J. Chem. Phys. 2021)

Committor: A useful projection

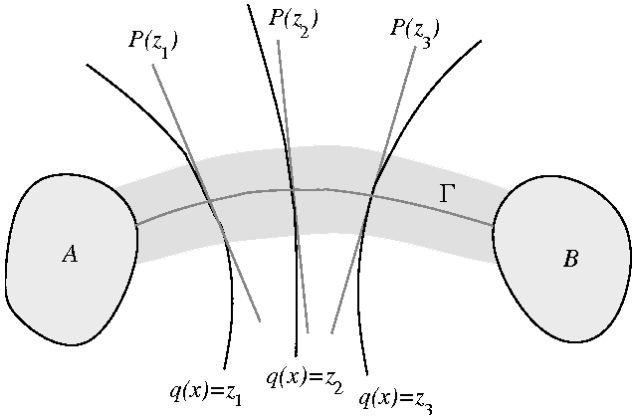

Onsager 1939; Vanden-Eijnden 2006

$q(\mathbf{X})$: the probability that trajectories initiated at configuration ***X*** reach state B prior to state A

- Committor encapsulates mechanisms and is **arguably** the perfect coordinate for biased molecular dynamics
  - Requires good state definition
  - Becomes difficult for multi-state systems

Lesson:

Instead of reconstructing the whole configuration, predicting states is enough

How to get the state labels?

Learn metastable states on-the-fly

Central idea:

If the system was initially located at state *i*, then after the time delay Δt, it should still have the largest probability to be found at state *i*.

(Metastability of state *i* with respect to Δt)

Δt filters out states with very short life-time

controls the level of coarse-graining

D. Wang, P. Tiwary (J. Chem. Phys. 2021)

Learn metastable states on-the-fly

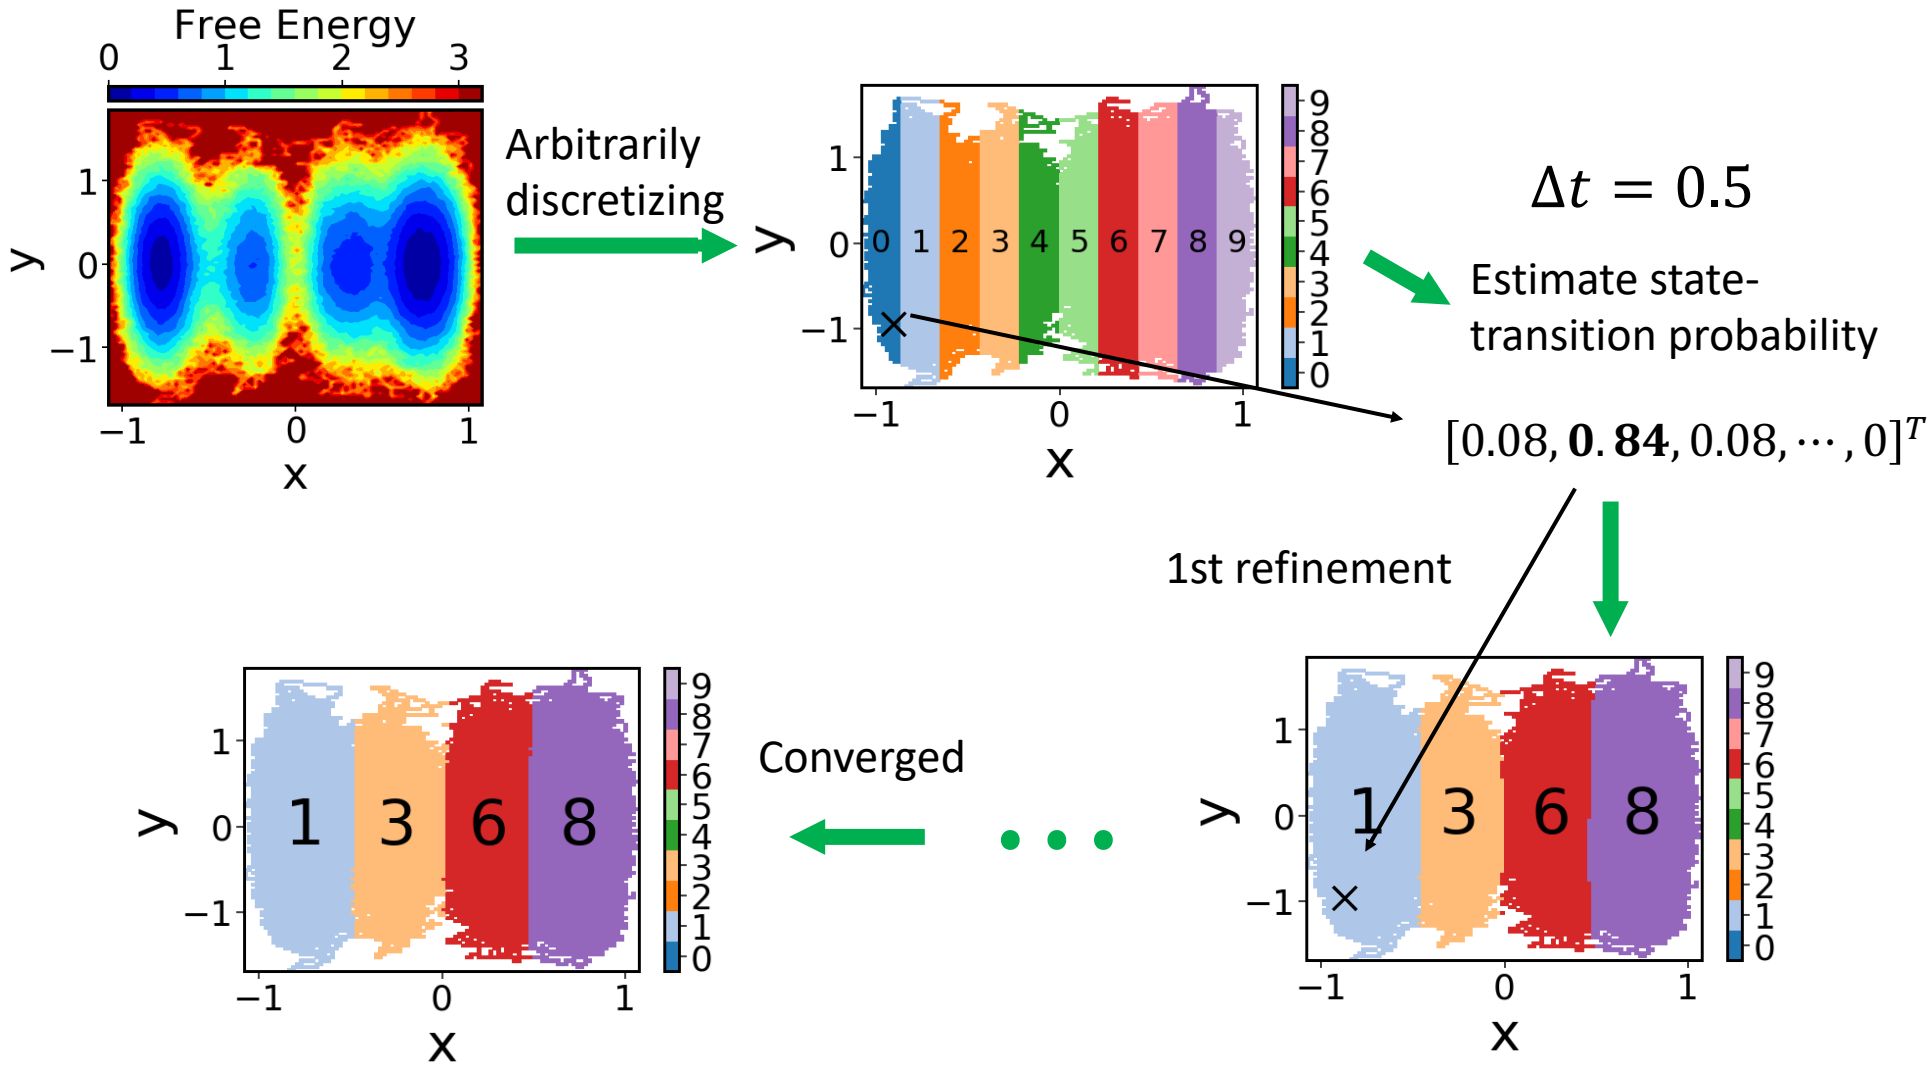

D. Wang, P. Tiwary (J. Chem. Phys. 2021)

Learn metastable states on-the-fly

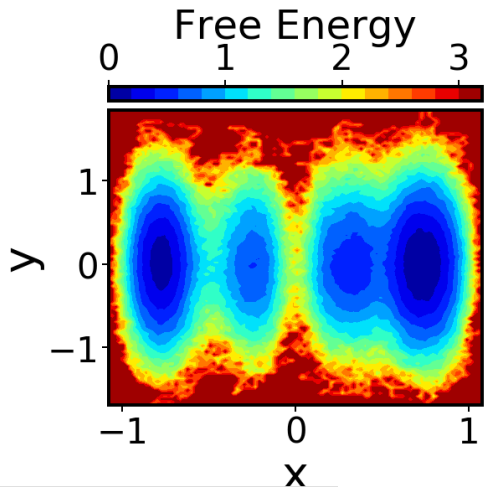

Arbitrarily discretizing

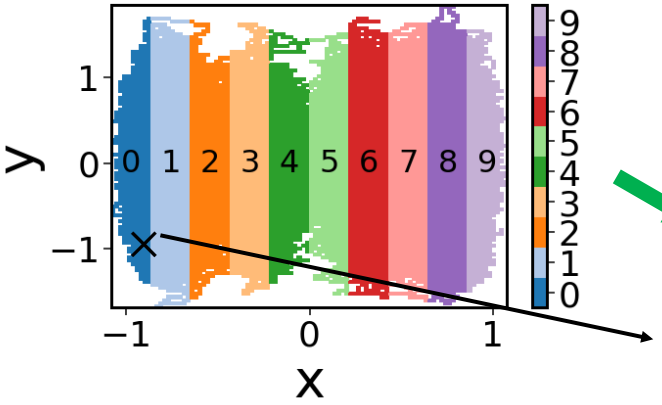

$\Delta t = 0.5$

Estimate state-transition probability

$[0.08, \mathbf{0.84}, 0.08, \cdots, 0]^T$

1st refinement

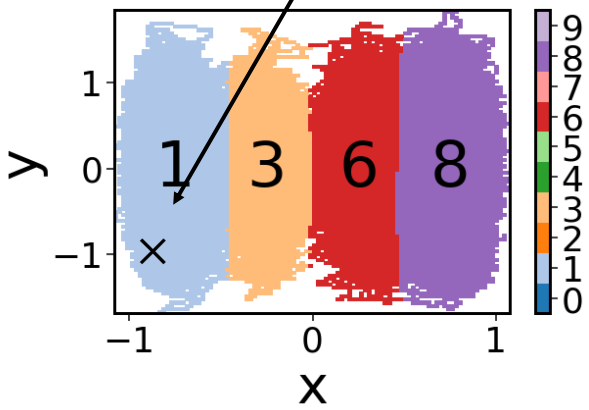

Converged

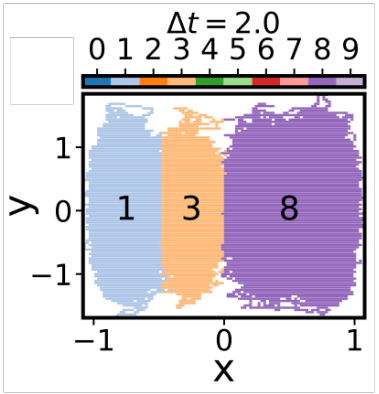

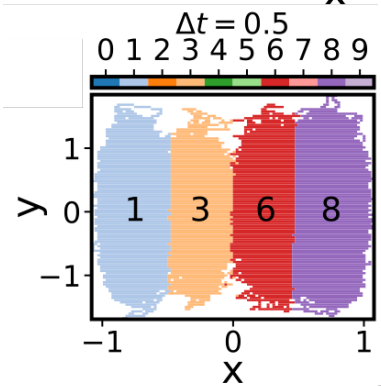

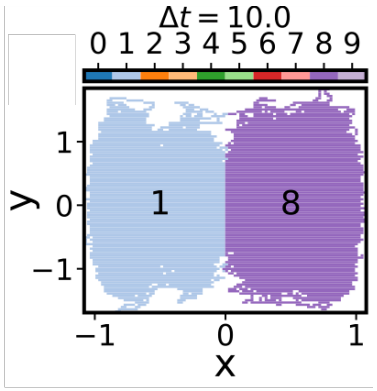

D. Wang, P. Tiwary (J. Chem. Phys. 2021)

State Predictive Information bottleneck (SPIB)

Introducing "metastable states" into representation learning

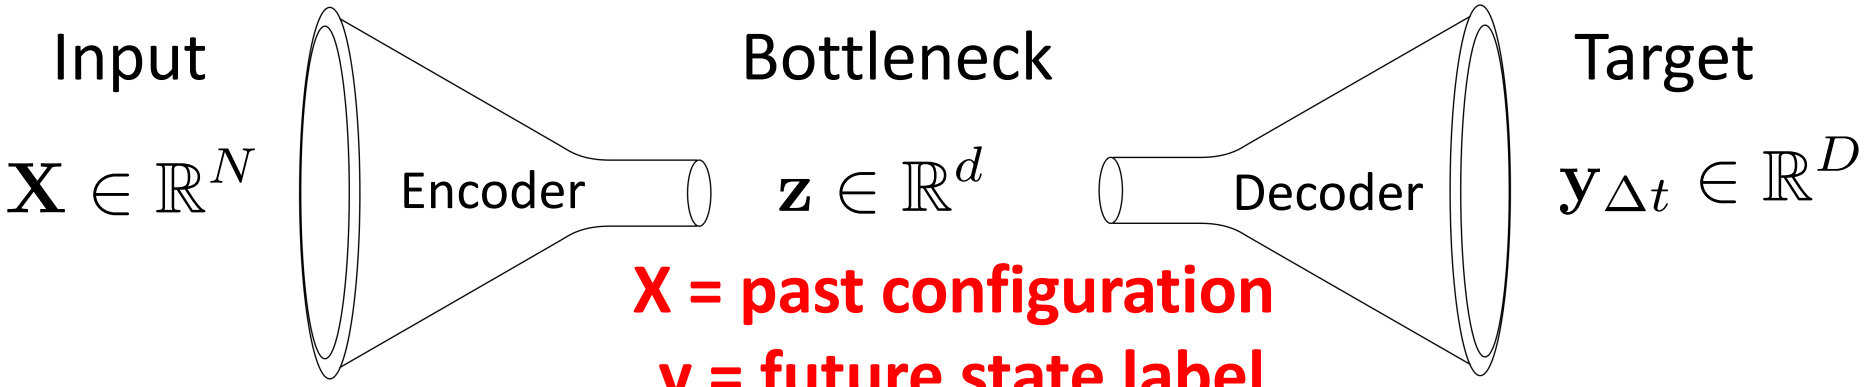

X = past configuration

y = future state label

z *is bottleneck or Reaction Coordinate (RC)*

Compresses X maximally*: low I*(X, z)

Predicts y maximally*: high I*(y, z)

where *I* is mutual information

To find z, maximize objective function* $\mathcal{L}_{IB} \equiv I(\mathbf{z}, \mathbf{y}_{\Delta t}) - \beta I(\mathbf{X}, \mathbf{z})$

β controls complexity of RC

Use a variational approach to construct a lower bound for optimization

-> Alemi, Alexander A., et al. "Deep variational information bottleneck." (2016) *arXiv:1612.00410*

D. Wang, P. Tiwary (J. Chem. Phys. 2021)

State Predictive Information bottleneck (SPIB)
Introducing "metastable states" into representation learning

Lower bound:

$$\mathcal{L}_{IB} \geq \mathcal{L} = \frac{1}{M}\sum_{n=1}^{M}[\log q_\theta(\mathbf{y}^{n+s}|\mathbf{z}^{(n)}) - \beta \log \frac{p_\theta(\mathbf{z}^{(n)}|\mathbf{X}^n)}{r_\theta(\mathbf{z}^{(n)})}]$$

$$\text{if } \mathbf{y} = \mathbf{X} \longrightarrow \beta\text{-VAE}$$

$$\text{if } \mathbf{y} = \mathbf{X},\ \beta = 1 \longrightarrow \text{VAE}$$

Variational autoencoder (VAE)

D. Wang, P. Tiwary (J. Chem. Phys. 2021)

State Predictive Information bottleneck (SPIB)

Introducing "metastable states" into representation learning

$$\mu = \mu_\theta(\mathbf{X})$$
$$\sigma = \sigma_\theta(\mathbf{X})$$

$$\log p_\theta(\mathbf{z}^n|\mathbf{X}^n) = \log \mathcal{N}(\mathbf{z}^n; \mu, \sigma I)$$

Lower bound:

Gaussian encoder

$$\mathcal{L}_{IB} \geq \mathcal{L} = \frac{1}{M}\sum_{n=1}^{M}[\log q_\theta(\mathbf{y}^{n+s}|\mathbf{z}^{(n)}) - \beta \log \frac{p_\theta(\mathbf{z}^{(n)}|\mathbf{X}^n)}{r_\theta(\mathbf{z}^{(n)})}]$$

categorical decoder

(future state prediction)

$$\log q_\theta(\mathbf{y}^{n+s}|\mathbf{z}^n) = \sum_{i=1}^{D} y_i^{n+s} \log \mathcal{D}_i(\mathbf{z}^n;\theta)$$

One-hot vector

e.g. [0, 0, 1, 0]

softmax-output

weighted VampPrior

(multi-modal distribution)

$$r_\theta(\mathbf{z}) = \sum_{k=1}^{K} \omega_k \; p_\theta(\mathbf{z}|\mathbf{u}^k)$$

mixture of Gaussians

Tomczak, Welling PMLR 2018

D. Wang, P. Tiwary (J. Chem. Phys. 2021)

1. SPIB: A future state predictor

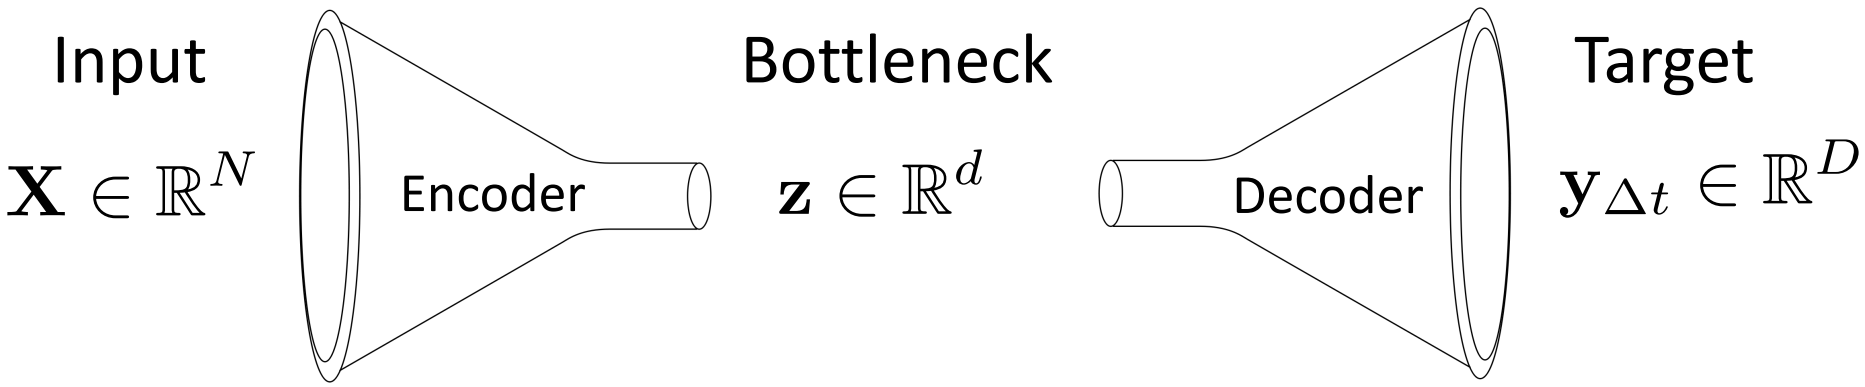

SPIB

State-transition density

$$P_{\Delta t}(y = j; \mathbf{X})$$

Markov state model (MSM)

State-transition matrix

$$P_{\Delta t}(y = j; y = i)$$

- Refines state labels on-the-fly
- Identifies the “boundary” -- transition states

$$\text{transition state} \equiv \{X \mid P_{\Delta t}(y = i; \mathbf{X}) = P_{\Delta t}(y = j; \mathbf{X}),\ i \neq j\}$$

State predictive information bottleneck converges to 0.5 committor as time-delay increased

Muller potential

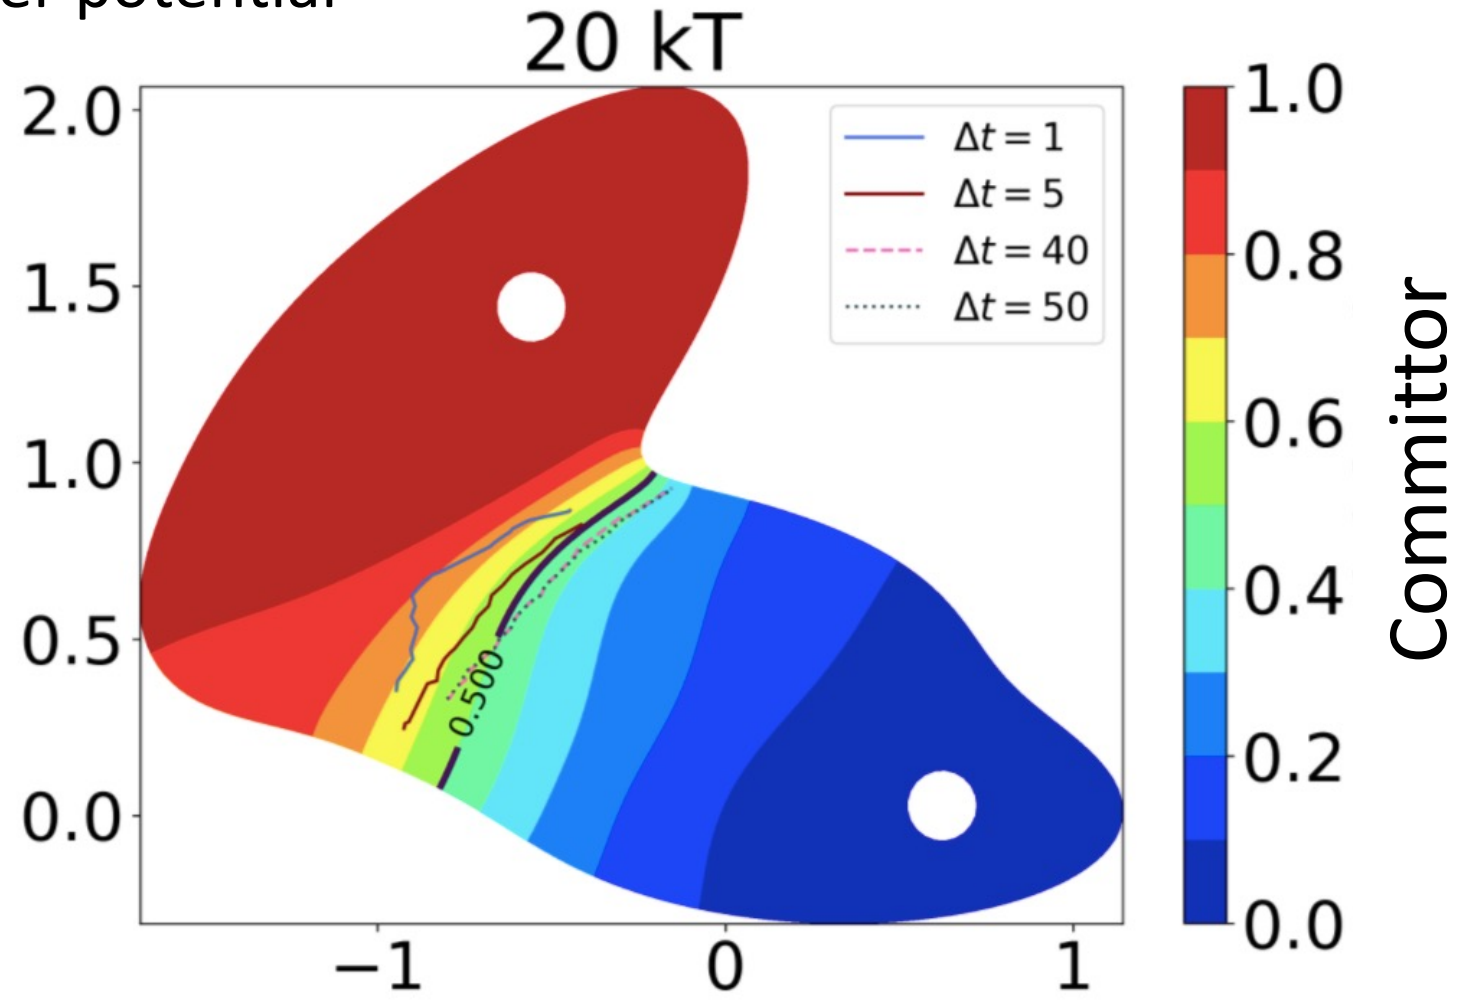

Calculations by Vanessa Meraz
Reference isocommittor surfaces from finite element code
thanks to Luke Evans and Maria Cameron

State predictive information bottleneck accurately identifies transition states for alanine dipeptide in vacuum

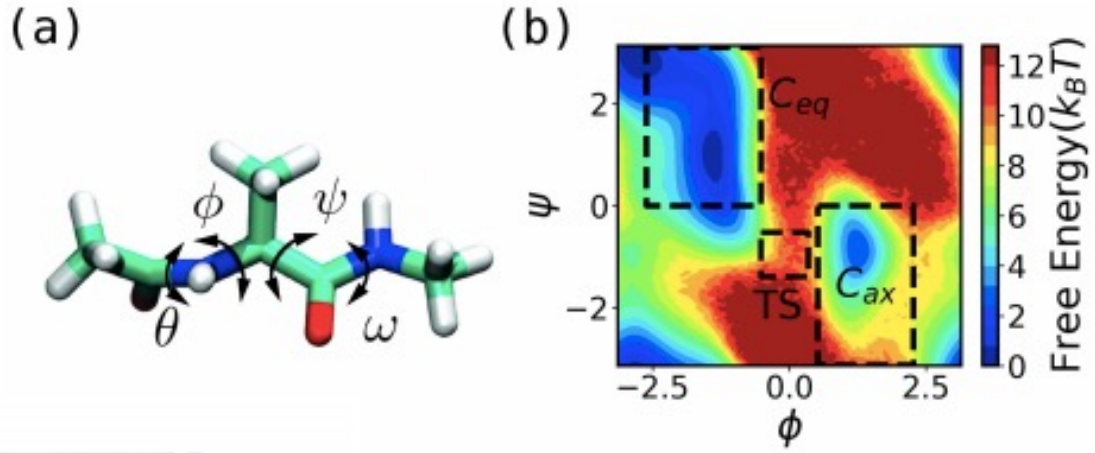

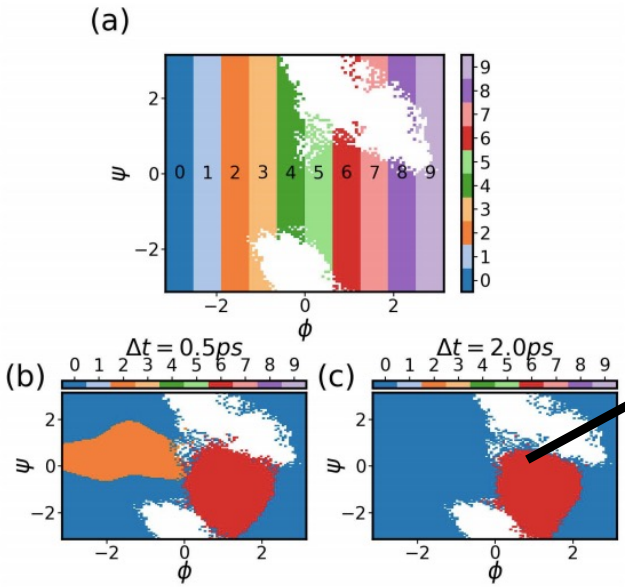

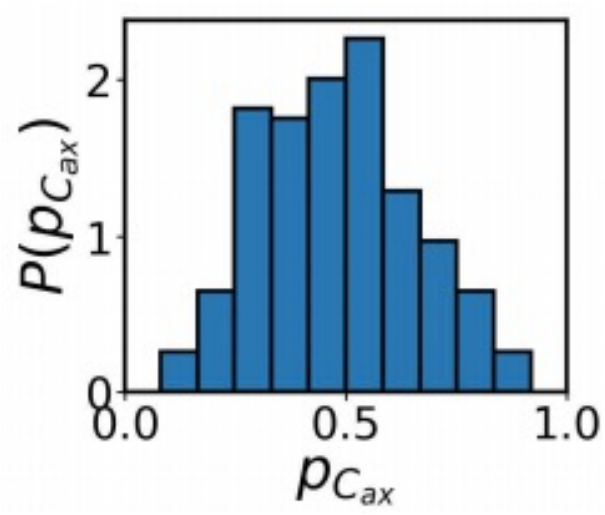

Arbitrary number of starting states give 2/3 states for long/short Δt

Lag time is fast mode filter & works with any number of states

Information bottleneck gives sharply peaked committor

"State predictive information bottleneck", Wang, Tiwary (JCP 2021)

Correct transition states for NaCl ion-pair dissociation

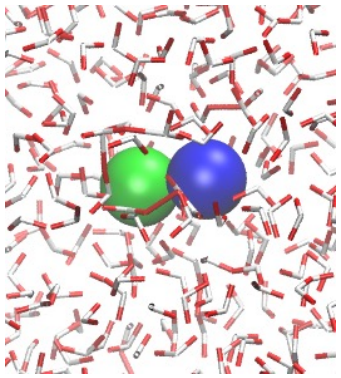

Δt controls the level of coarse-graining

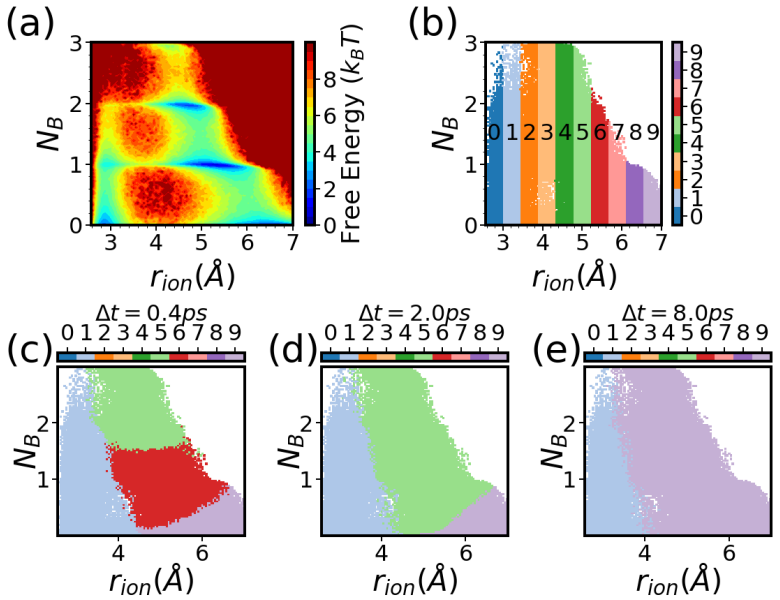

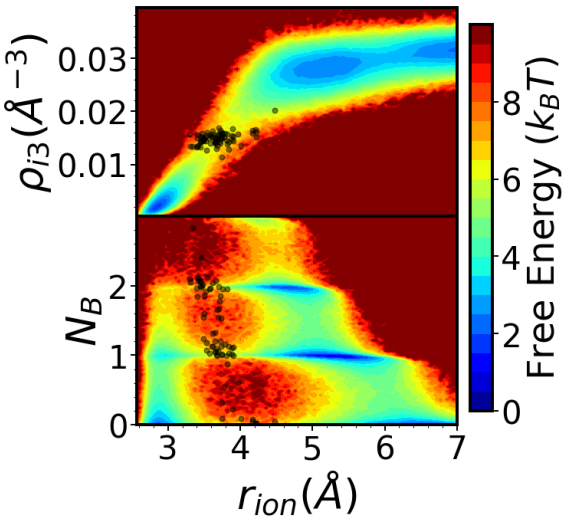

Transition states from SPIB
(25 ns equilibrium traj~40 transitions)

Wang, Zhao, Weeks, Tiwary JPCB 2022

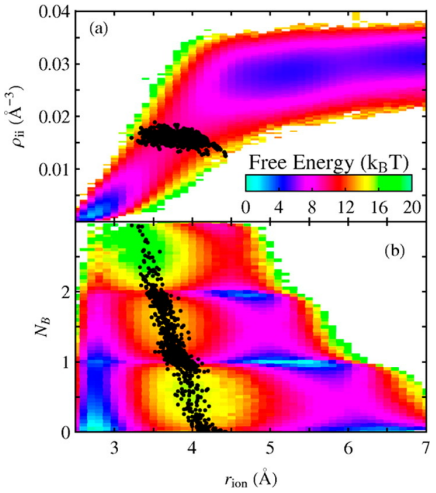

Transition states from Likelihood Maximization (24000 trajs ~ 72ns)

Mullen, Shea, Peters (2014)

2. SPIB: A continuous embedding of Markov state model that preserves maximum information about state transitions

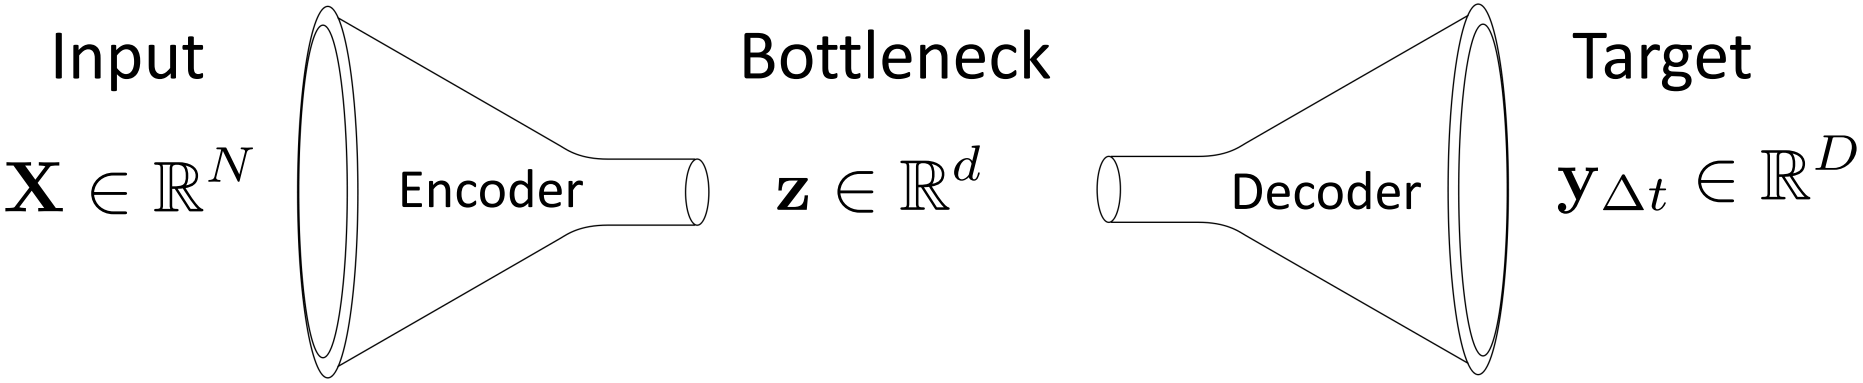

Na-Cl Ion Pair Dissociation

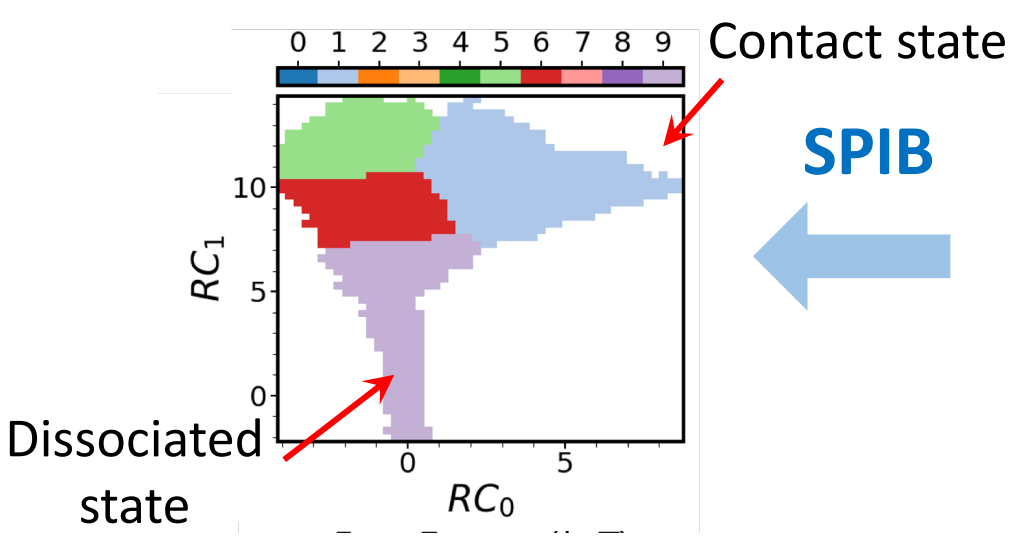

Markov state model (MSM)

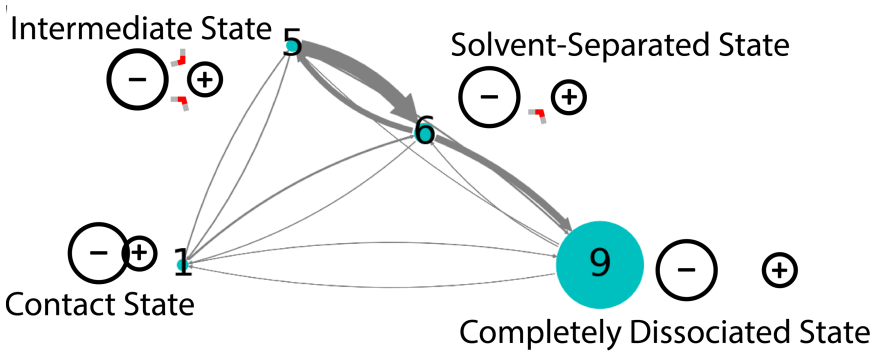

Wang, Zhao, Weeks, Tiwary JPCB 2022

2. SPIB: A continuous embedding of Markov state model that preserves maximum information about state transitions

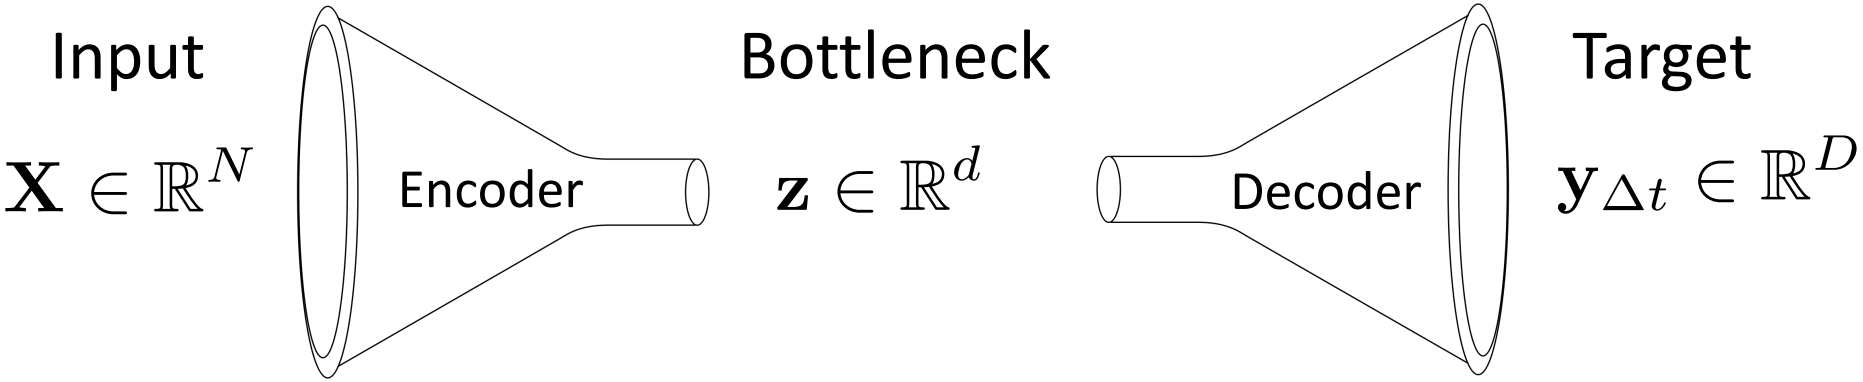

Na-Cl Ion Pair Dissociation

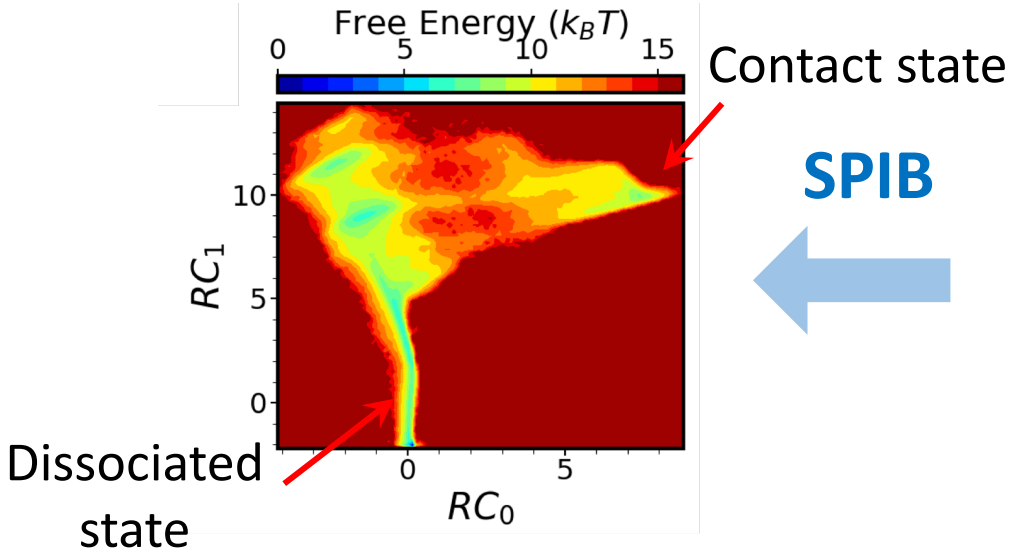

SPIB

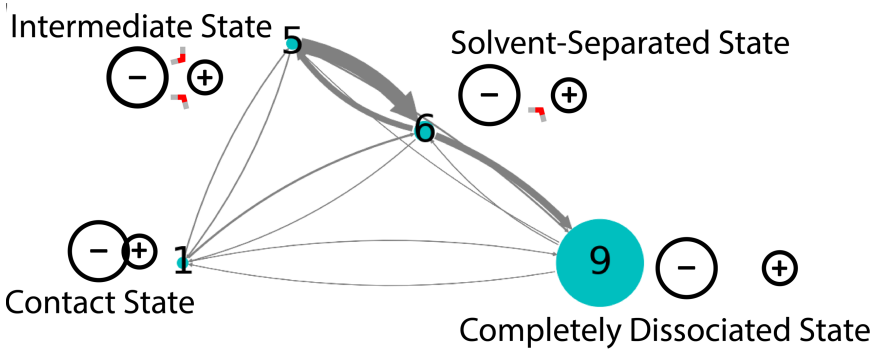

Wang, Zhao, Weeks, Tiwary JPCB 2022

3. SPIB: A blueprint for sampling the learned RC could be used to enhance molecular dynamics

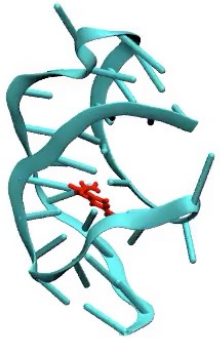

$$\mathcal{L}_{\mathcal{IB}} \equiv I(\chi, \mathbf{y}) - \beta I(\mathbf{X}, \chi)$$

Forget noise and learn features that predict future

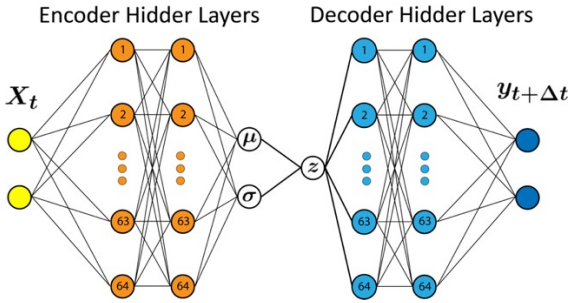

Machine Learning

+

Enhanced Sampling

RC = Information Bottleneck

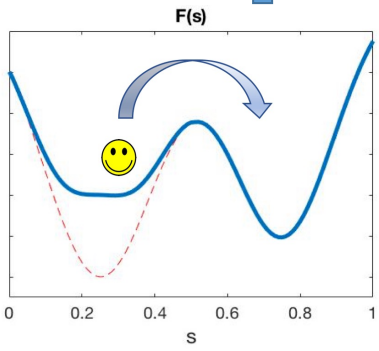

Run enhanced sampling using current RC estimate

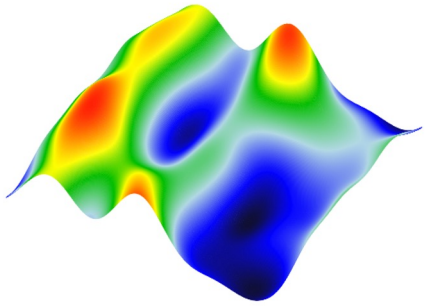

Agnostic to the enhanced sampling method used

Y Wang, Ribeiro, Tiwary, Curr. Op. Struc. Bio. 2019
Mehdi, D Wang, Pant, Tiwary *JCTC* 2022

Y Wang, Ribeiro, Tiwary, Nat. Commun. 2019

Applications: Overview

SPIB + metadynamics

(1) Small molecule permeation through lipid bilayer

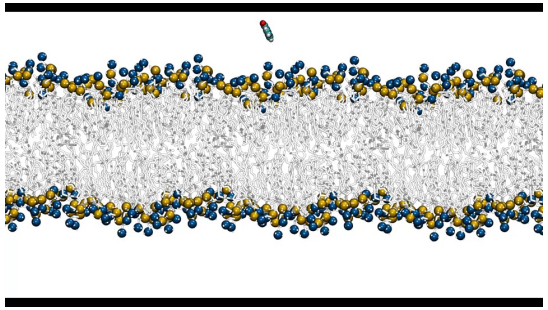

S Mehdi, D Wang, S Pant, P Tiwary ***JCTC*** 2022

(2) Nucleation of urea and glycine in water

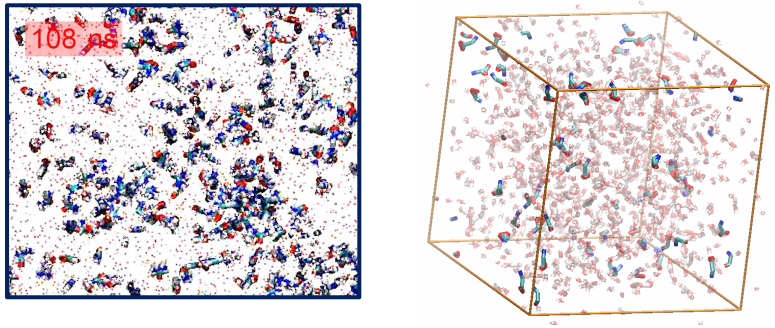

Zou, Beyerle, Tsai, Tiwary. ***PNAS*** 2023

(3) Imatinib dissociation from Abl kinase

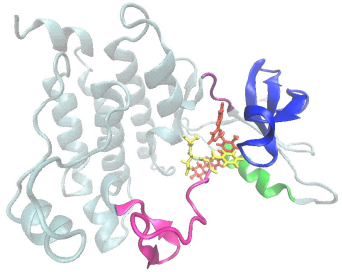

Lee, D Wang, Tiwary (under preparation)

Alphafold2 + SPIB + metadynamics

(4) Helix position changes of SSTR2 GPCR

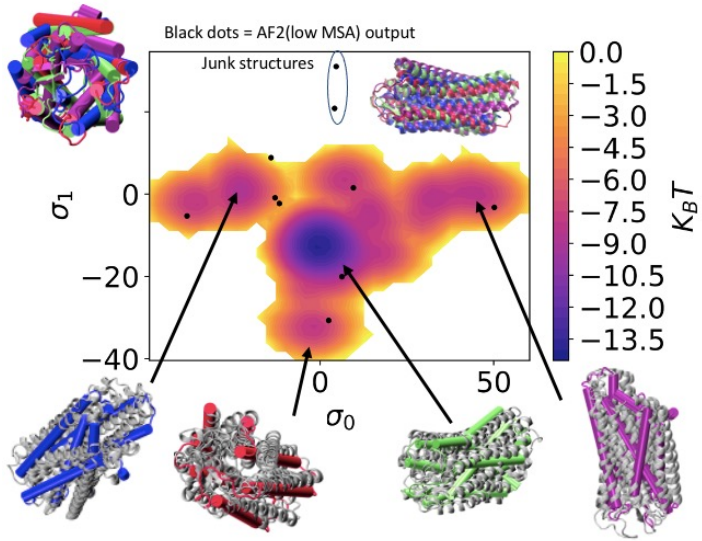

Vani, Aranganathan, D Wang, Tiwary **(bioRxiv** 2022)

SPIB + weighted ensemble

(5) Protein folding of chignolin

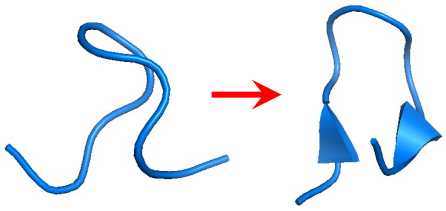

(6) Loop dynamics of T4 lysozyme

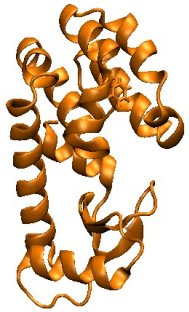

D Wang & Tiwary (under preparation)

Learning useful and meaningful low-dimensional representations from physics

- **State Predictive Information Bottleneck**

  D Wang & Tiwary, J. Chem. Phys. 2021

Introducing "metastable states" into representation learning

1. Identifying metastable states

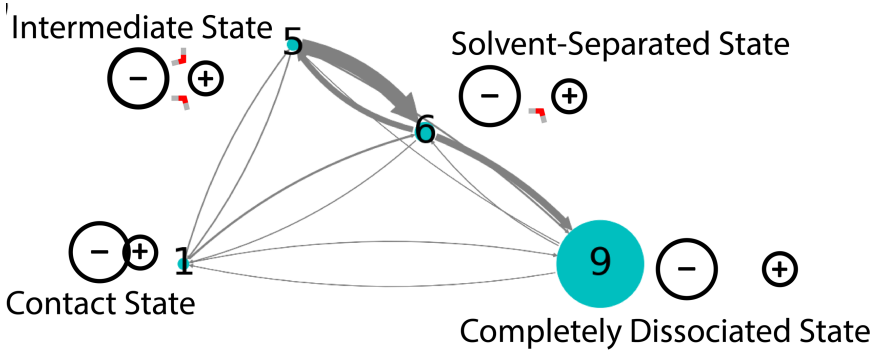

D Wang, Zhao, Weeks, Tiwary. JPCB. 2022

2. Predicting transition states

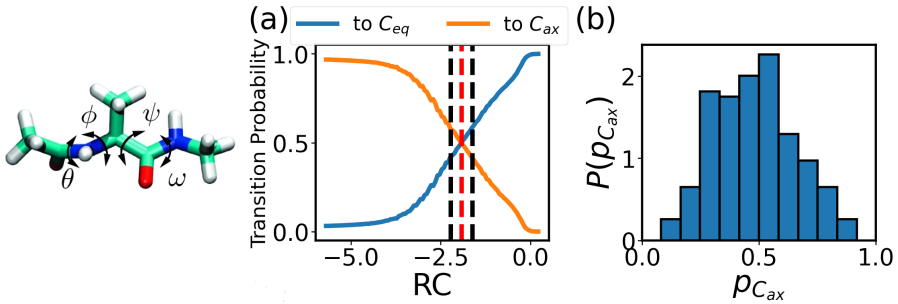

D Wang & Tiwary, J. Chem. Phys. 2021

3. Enhancing sampling

SPIB + metadynamics

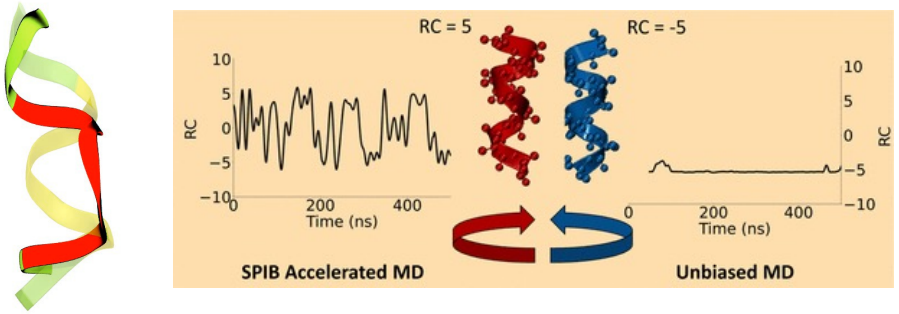

Mehdi, D Wang, Pant, Tiwary JCTC 2022

SPIB + weighted ensemble

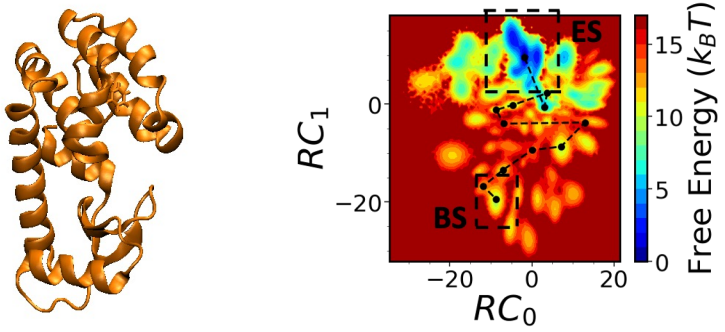

D Wang & Tiwary (under preparation)

Learning useful and meaningful low-dimensional representations from physics

- State Predictive Information Bottleneck

  D Wang & Tiwary, J. Chem. Phys. 2021

- **Dynamics Constrained Autoencoder**

  D Wang, Y Wang, Evans & Tiwary, arXiv:2209.00905

Introducing the law of dynamics into representation learning

1. Recovering kinetics from distorted data

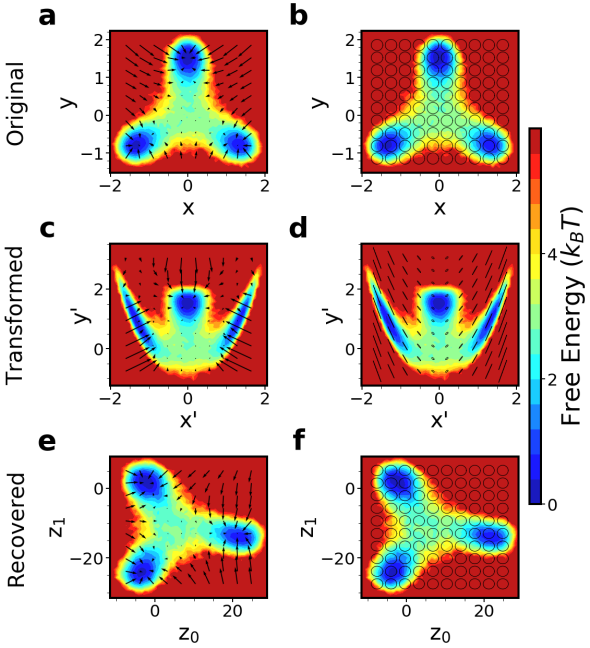

2. Learning dihedrals from atomic coordinates

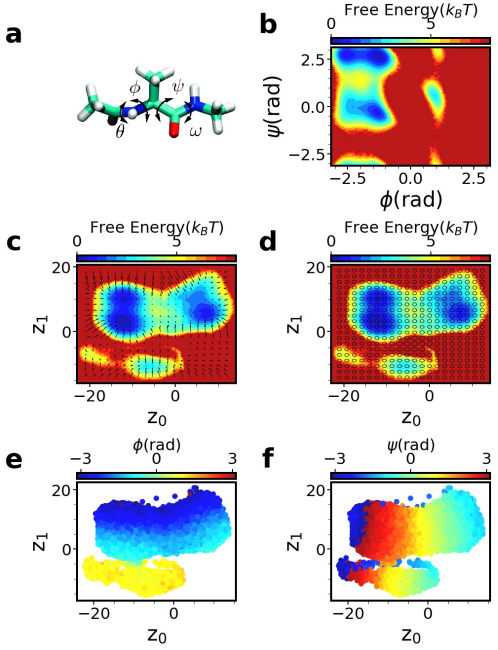

3. Interpreting experimental data

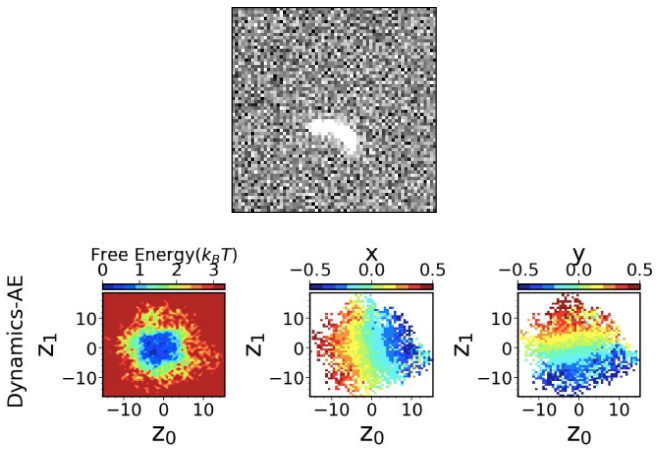

Dilemma in Representation Learning

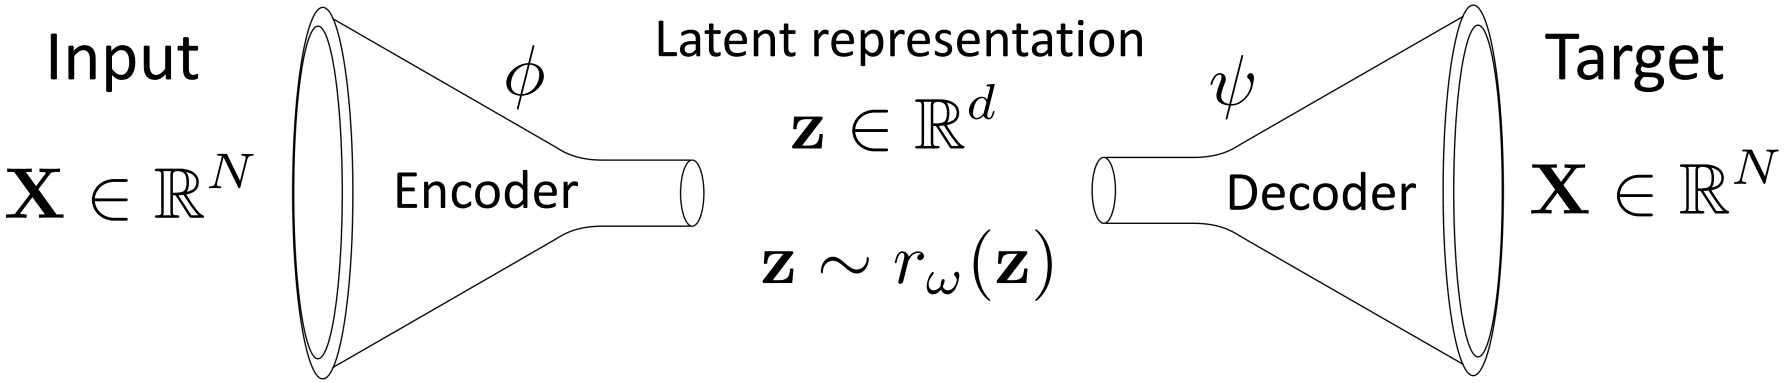

$$\mathcal{L}(\phi, \psi, \omega) \equiv$$

$$\mathbb{E}_{\hat{p}(\mathbf{X})}\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{X})}\left[-\log q_\psi(\mathbf{X}|\mathbf{z}) + \beta \log \frac{q_\phi(\mathbf{z}|\mathbf{X})}{r_\omega(\mathbf{z})} \right]$$

$$= \mathcal{L}_{REC} + \beta \mathcal{L}_{REG},$$

Reconstruction error Regulariser

Variational autoencoder (VAE)

Kingma & Welling. (2013). *arXiv:1312.6114*.
Dilokthanakul et al. (2016). *arXiv:1611.02648*.
Kobyzev, Prince, & Brubaker. (2020).

Traditional prior:
Standard Gaussian -> **over-regularization**
VampPrior
Mixture of Gaussians
Normalizing flow -> **under-regularization**

A lesson learned from physics

Falling apples

Pendulum system

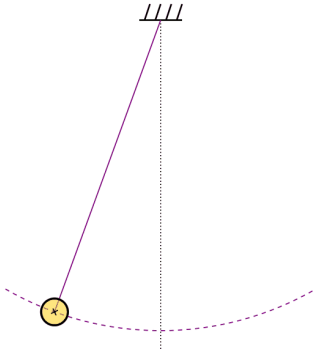

Solar system

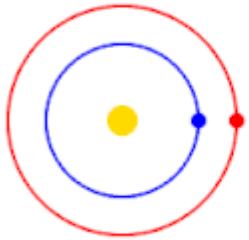

Changing: $H = \frac{p^2}{2m} + mgh$ $H = \frac{p_\theta^2}{2m} + mgl(1-\cos\theta)$ $H = \frac{p_1^2}{2m_1} + \frac{p_2^2}{2m_2} - G\frac{m_1 m_2}{r_{12}}$

Unchanging: $\mathbf{F} = m\mathbf{a}$ or $\frac{d\mathbf{q}}{dt} = \frac{\partial H}{\partial \mathbf{p}}, \frac{d\mathbf{p}}{dt} = -\frac{\partial H}{\partial \mathbf{q}}$

Lesson:

Though different systems can have different probability or free energy distributions, they typically obey the same dynamics, such as Newton's second law in classical mechanics and Schrödinger's equation in quantum mechanics

Wang, D., Wang, Y., Evans, L., & Tiwary, P. (2022). *arXiv:2209.00905*.

Solution: Learning latent representations from dynamics

Enforce latent representation to follow a specific probability distribution

Enforce latent representation to follow a specific yet generic class of dynamics

Impose a Gaussian

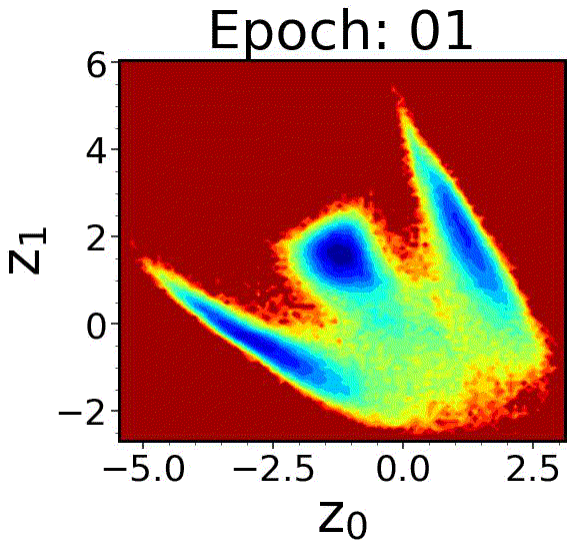

Impose overdamped Langevin dynamics

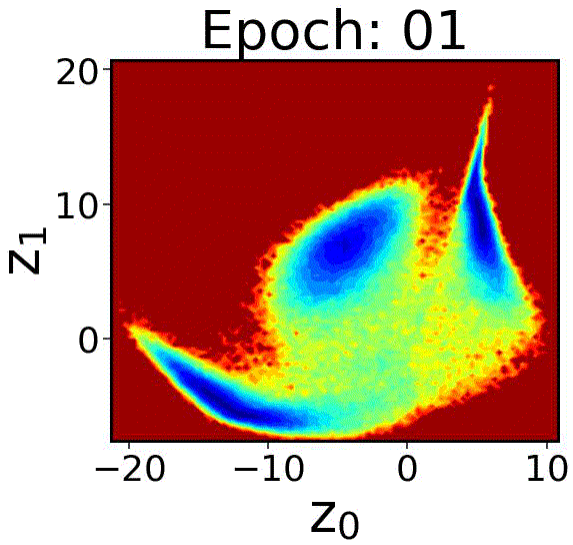

Representation learning for a three-state model potential

Wang, D., Wang, Y., Evans, L., & Tiwary, P. (2022). *arXiv:2209.00905*.

Enforcing the latent dynamics by minimizing the Wasserstein distance

For a Markovian process,

Transition density

$$p(\Delta \mathbf{z}|\mathbf{z}_t) \quad \text{where } \Delta \mathbf{z} \equiv \mathbf{z}_{t+\Delta t} - \mathbf{z}_t$$

Minimize

$$\mathcal{L}_{REG} = \mathcal{D}_{SW}(q_\phi(\Delta \mathbf{z}_t|\mathbf{z}_t), r_\omega(\Delta \mathbf{z}_t|\mathbf{z}_t))$$

encoded transition density

prior transition density

(Sliced) Wasserstein distance

Kolouri, S., Pope, P. E., Martin, C. E., & Rohde, G. K. (2018). *arXiv:1804.01947*.

"Euclidean distance" for probability distributions

Wang, D., Wang, Y., Evans, L., & Tiwary, P. (2022). *arXiv:2209.00905*.

Learning Langevin dynamics prior through likelihood maximization

To get the prior transition density $r_\omega(\Delta \mathbf{z}_t | \mathbf{z}_t)$

Overdamped Langevin dynamics

$$d\mathbf{z}_t = [\mathbf{M}(\mathbf{z}_t)\mathbf{f}(\mathbf{z}_t) + \nabla \cdot \mathbf{M}(\mathbf{z}_t)]\, dt + \sqrt{2\mathbf{M}(\mathbf{z}_t)} d\mathbf{w}_t,$$

Assume a diagonal diffusion matrix -> dynamically "uncorrelated" latent representation

$$dz_i = \left[M_{ii}(\mathbf{z}_t) f_i(\mathbf{z}_t) + \frac{\partial M_{ii}(\mathbf{z}_t)}{\partial z_i} \right] dt + \sqrt{2M_{ii}(\mathbf{z}_t) dt}\, \epsilon_i$$

where $\mathbf{f}(\mathbf{z}) = \mathbf{f}_\omega(\mathbf{z}),\ \mathbf{M} = \mathbf{M}_\omega(\mathbf{z})$

In practice, $M^*_{ij} \equiv \delta_{ij}$ a constant diffusion will be enforced

Likelihood maximization

$$\mathcal{L}_{prior}(\omega; \phi, \psi) = -\mathbb{E}_{\hat{p}(\mathbf{z}_t, \mathbf{z}_{t+\Delta t})} \log r_\omega(\Delta \mathbf{z}_t | \mathbf{z}_t)$$

$$\log r_\omega(\Delta \mathbf{z}_t \equiv \mathbf{z}_{t+1} - \mathbf{z}_t | \mathbf{z}_t) = -\frac{1}{2} \sum_i \left[\log M_{ii}(\mathbf{z}_t) + \frac{\left(\Delta z_i - M_{ii}(\mathbf{z}_t) f_i(\mathbf{z}_t) - \frac{\partial M_{ii}(\mathbf{z}_t)}{\partial z_i} \right)^2}{2M_{ii}(\mathbf{z}_t)} \right]$$

Wang, D., Wang, Y., Evans, L., & Tiwary, P. (2022). *arXiv:2209.00905*.

Dynamics-constrained autoencoder (DynAE)

Introducing the laws of dynamics into representation learning

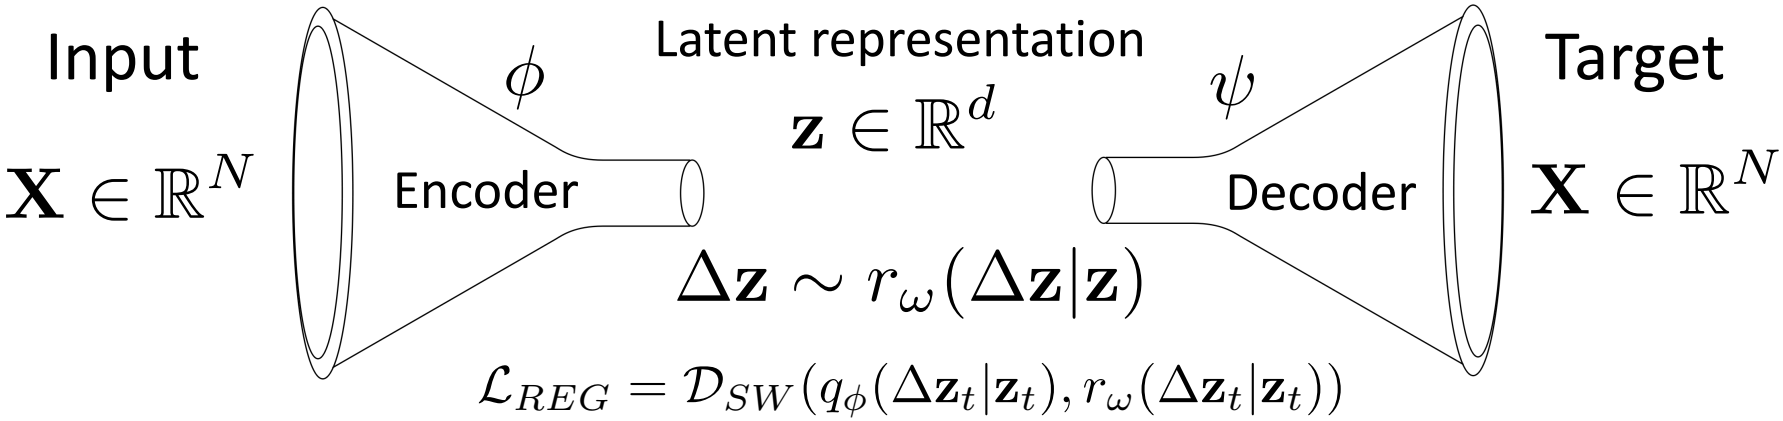

- Constraining the latent dynamics is enough to **uniquely** identify the latent representation up to an **isometry**

  (Theoretically proved by Hasan, et al. 2021)

- Decorrelates the latent representation by assuming a diagonal diffusion matrix -> dynamics-based disentanglement

- Only dynamics matters -> more flexible sampling strategies

  Draw samples from a flatter distribution to focus more on rarely sampled regions

  $\hat{p}_{data}(\mathbf{z}) \propto e^{-\beta F(\mathbf{z})}$

  $\hat{p}_{train}(\mathbf{z}) \propto e^{-\beta F(\mathbf{z})/\gamma}$ where $\gamma > 1$

Wang, D., Wang, Y., Evans, L., & Tiwary, P. (2022). *arXiv:2209.00905*.

Recover underlying kinetics from distorted data

Three-state model potential with constant diffusion

A nonlinear mapping function is used to warp the data into a "half-moon" shape and make diffusion anisotropic

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = f(\mathbf{X}) = \begin{bmatrix} x \\ y + x^2 \end{bmatrix}.$$

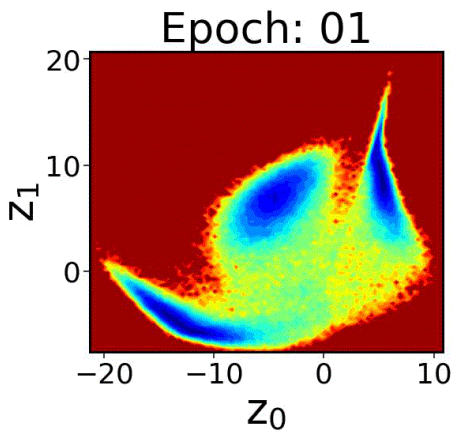

Dynamics-AE recovers ground truth

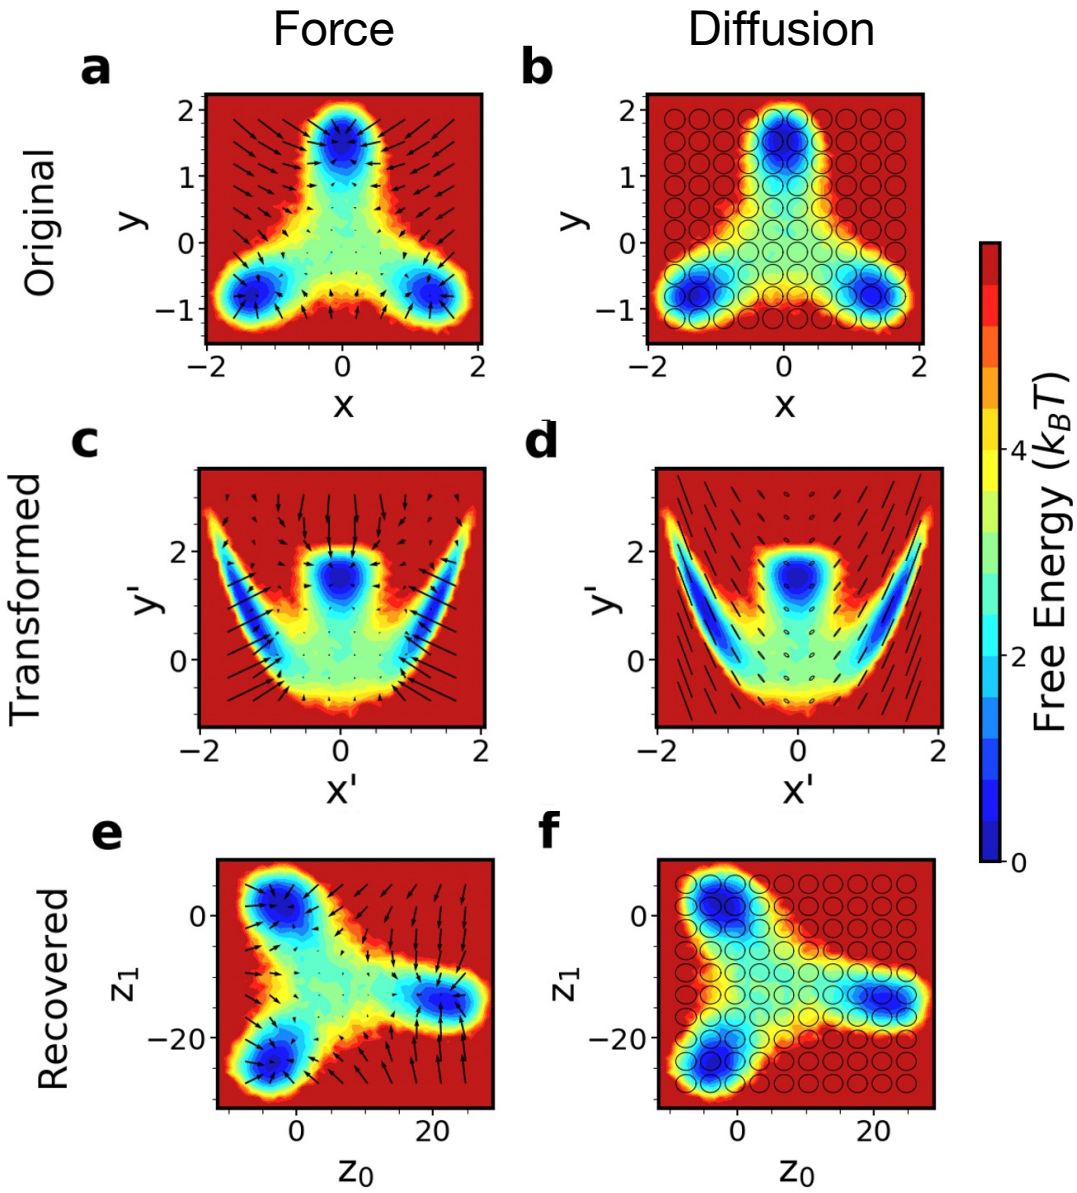

* Dynamics-AE: Wang, Wang, Evans, Tiwary (2022) *arXiv:2209.00905*

Learning dihedrals from atomic coordinates

Alanine dipeptide in water

Dynamics-AE recovers dihedral angles ϕ and ψ from a 30-dimensional space comprising three-dimensional coordinates of 10 heavy atoms

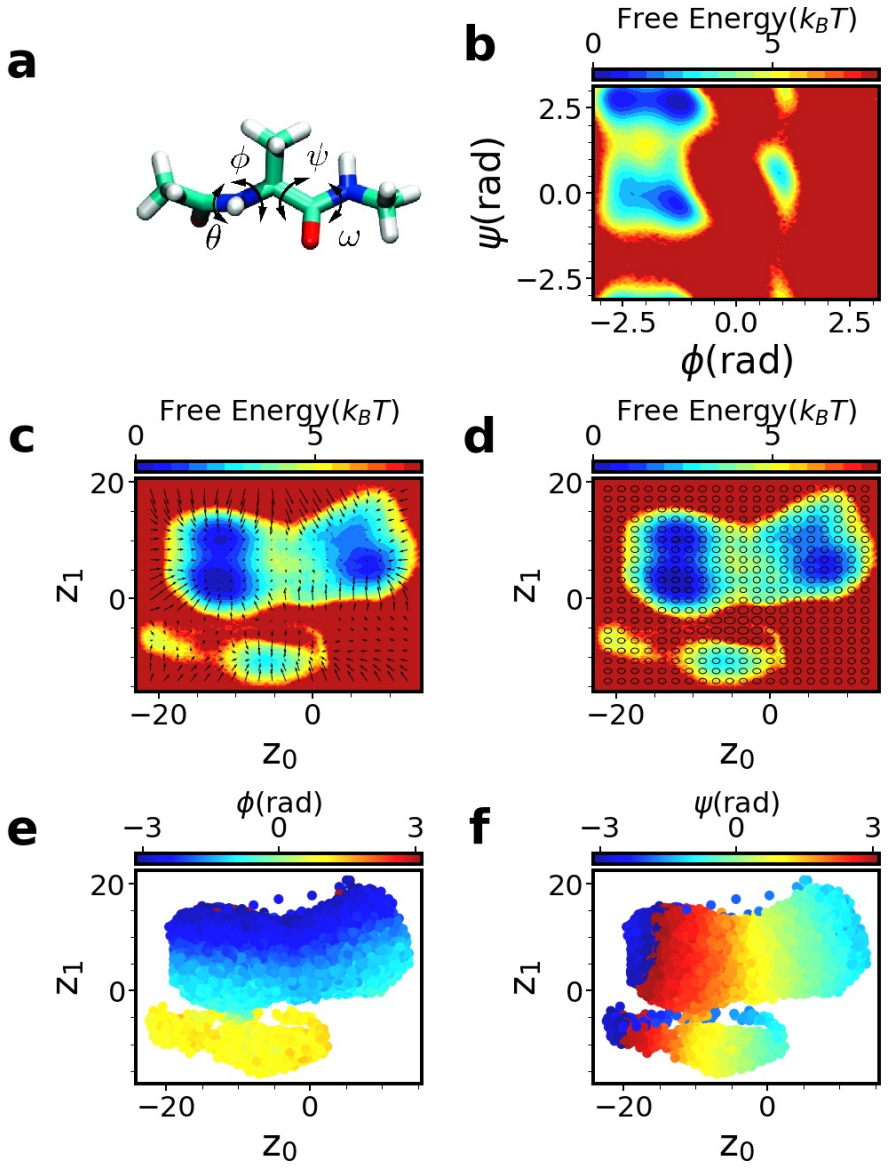

* Dynamics-AE: Wang, Wang, Evans, Tiwary (2022) *arXiv:2209.00905*

Learning *meaningful* representations from experiments

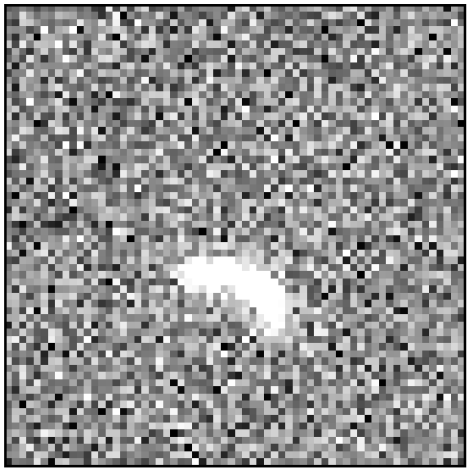

DNA Brownian motion in solution
Fluorescence Optical Microscope
Lameh, Ding, Stein Phys. Rev. Appl. 2020

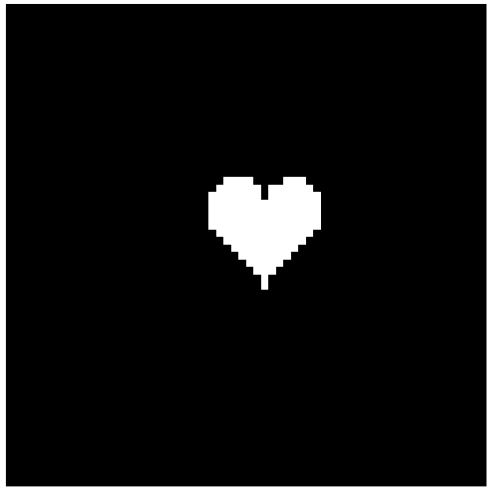

dSprites dataset
github.com/deepmind/dsprites-dataset

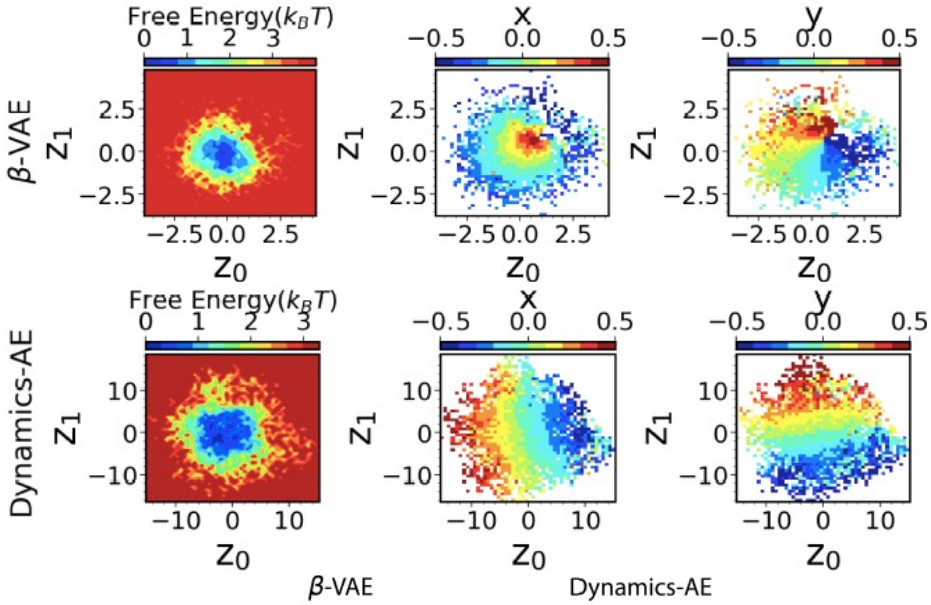

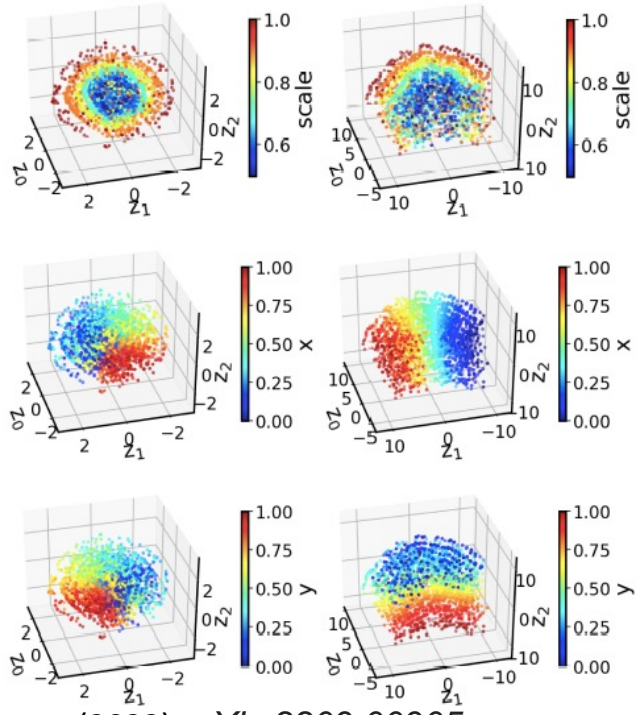

* Dynamics-AE: Wang, Wang, Evans, Tiwary (2022) *arXiv:2209.00905*

***Take-home:** We can use ideas from physics to guide the design of artificial intelligence (AI) methods to learn **useful** and also **meaningful** representations.*

- **State Predictive Information Bottleneck**

  D Wang & Tiwary, J. Chem. Phys. 2021

  Introducing "metastable states" into representation learning

- **Dynamics Constrained Autoencoder**

  D Wang, Y Wang, Evans & Tiwary, arXiv:2209.00905

  Introducing the law of dynamics into representation learning

All codes open-source @ github.com/tiwarylab

Advisor:
Pratyush Tiwary

@tiwarylab

Postdocs

Eric Beyerle

Bodhi Vani

Ruiyu Wang

Undergrads

Eric Fields
(Biochem+CS) → **Harvard**

Fiona Mon
(Biochem) → **BU**

Michael Strobel
(CS+Statistics) → **UCI**

Disha Sanwal
(Chemistry+Appl Math)

Highschooler

Thalia Haines

PhD students

Yihang Wang
(Biophysics) → Chicago

Sun-Ting Tsai
(Physics) → UMich

Zachary Smith
(Biophysics)

Luke Evans
(Appl Math)

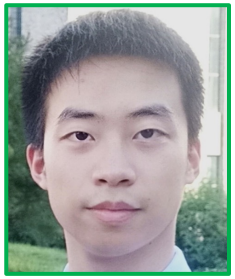

Dedi Wang
(Biophysics)

Postdog

Pakora the Dog

Connor Zou
(Chemistry)

Shams Mehdi
(Biophysics)

Lukas Herron
(Biophysics)

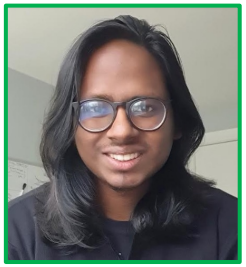

Akashnathan
Aranganathan
(Biophysics)

Suemin Lee
(Biophysics)

Vanessa Meraz
(Chemical Physics)

Thanks for $ and 💻 support