


Learning protein dynamics from quantum fragments

Ab initio characterization of protein molecular dynamics with AI²BMD


Tong Wang et al. • Nature 635, 1019–1027 (2024)

Reusable quantum-trained fragments connect atomic forces to protein motion.

Sean Florez • CSCI 7000 • September 28, 2026




From a whole protein to fragment energies and forces




Atomic coordinates define each local environment. Combined forces update the protein's motion.

Reusable fragments make quantum training data tractable



A fixed vocabulary of local units supports evaluation across many protein sequences.

Overlapping fragments reconstruct protein energy




$$E_{\text{protein}} \approx \sum E_{\text{units}} - \sum E_{\text{overlaps}} + E_{\text{remaining interactions}} \quad | \quad 20 \text{ capped amino-acid units} + 1 \text{ shared Ace-Nme unit}$$

 Capped units
 Ace-Nme overlap
 --- Coulomb + LJ


Subtracting the shared contribution corrects double counting; additional interactions connect separate regions.

Different conformations teach the model how energy changes

A conformation is one 3D arrangement of the same bonded atoms.



Rotating the **N-Cα bond (ϕ)** changes shape, energy and forces.




*Two backbone angles at 5° intervals, per non-proline unit. Grid and conformations are schematic.

†Derived: $1,036,800 \text{ per unit} \div 5,184 = 200$.


Angles vary continuously. These examples sample the space of possible shapes.

$$\begin{aligned}
 & \times 20 \text{ amino-acid units} \\
 & \times 5,184 \text{ starting conformations per unit } (72 \times 72)^* \\
 & \times 200^\dagger \text{ snapshots per start} \\
 & = 20,736,000 \\
 & + 144,000 \text{ Ace-Nme } (72 \times 2,000) \\
 & = \mathbf{20.88M} \text{ labelled examples} \\
 & 80\% \text{ train } \approx 16.7M \\
 & 10\% \text{ validation} \cdot 10\% \text{ test}
 \end{aligned}$$

ViSNet learns energy and geometry-consistent forces



AI²BMD combines learned terms with classical interactions




- 1** Local overlapping units
ViSNet terms + overlap correction
- 2** Remaining protein interactions
Coulomb + Lennard-Jones terms
- 3** Solvent and coupling
AMOEBA polarizable force field

Reported simulation results describe this combined system.

Protein benchmarks show lower error against DFT


5 smaller proteins: whole-protein DFT reference



AI²BMD MEAN ENERGY MAE

0.038

$\text{kcal mol}^{-1} \text{atom}^{-1}$



MEAN FORCE MAE


1.974 8.094

AI²BMD MM

$\text{kcal mol}^{-1} \text{\AA}^{-1}$

4 larger proteins use a fragmented DFT reference. Nine proteins, 175–13,728 atoms; 200 configurations each.

Energy evaluation is much faster than DFT



TRP-CAGE, 281 ATOMS

0.072 s AI²BMD**21 min** DFT

Reported system: NVIDIA A6000 + 32 CPU cores

- Measured · DFT, first four proteins
- -○- - Extrapolated · DFT, last five proteins

Large-protein DFT timings are extrapolated.

These measurements concern energy evaluation;
full solvated trajectories add further work.

NMR observables test sampled conformations against experiment

Backbone-angle samples





Predicted J couplings



NMR comparison

PEARSON CORRELATION

■ **0.924** AI²BMD
 ■ **0.543** MM


A measured observable checks one physical consequence of the simulated ensemble.


Short chignolin trajectories capture folding and unfolding

Chignolin 10 residues


60 starting structures

Trajectories up to 10 ns

Starts unfolded




Starts folded




Plotted curves • Relative potential-energy error over time, against a QM protein + AMOEBA solvent reference

Selected trajectories show transitions between folded and unfolded conformations.

Energy re-evaluation improves selected thermodynamic estimates



Accuracy depends on chemistry, sampling and the reference model



Reference


Learned DFT targets + fragmentation + hybrid interactions

DEMONSTRATED

Force MAE 1.974 (AI²BMD) and 8.094 (MM) kcal mol⁻¹ Å⁻¹ against whole-protein DFT

CONDITION FOR REUSE

Agreement measures fidelity to this reference system



Chemistry


A defined vocabulary of protein units

DEMONSTRATED

20 amino-acid units plus Ace-Nme cover the proteins tested

CONDITION FOR REUSE

Ligands, nonstandard residues, charge states or reactions need new validation



Sampling

Ensembles must cover states relevant to the observable

DEMONSTRATED


Chignolin folding and unfolding in trajectories up to 10 ns

CONDITION FOR REUSE

Converged ensembles for each property of interest

Each new application needs validation at these three levels.

Reusable learned physics still needs experimental checks

**1 Reuse**

Train local units, compose larger systems

2 Physical structure

Geometry and energy-derived forces

3 Evaluation

Reference errors, trajectories and experimental observables

My takeaway: trustworthy reuse depends on the full model-and-sampling system.


Thank you

Questions?

Ab initio characterization of protein molecular dynamics with AI²BMD

Tong Wang et al. • Nature 635, 1019–1027 (2024)

Sean Florez • CSCI 7000



Fragment and solvent contributions form the hybrid potential

UNITS

$$1 \quad E_{\text{units}} = \sum E_{\text{dipeptide}} - \sum E_{\text{Ace-Nme}}$$

PROTEIN


$$2 \quad E_{\text{protein}} = E_{\text{units}} + E_{\text{remaining Coulomb}} + E_{\text{remaining LJ}}$$

TOTAL SYSTEM

$$3 \quad E_{\text{total}} = E_{\text{protein}} + E_{\text{AMOEBA,all}} - E_{\text{AMOEBA,protein}}$$

AMOEBA


Coulomb + LJ



ViSNet units

Subtract overlapping/internal contributions to avoid double counting.

Training combines energy and force supervision



Separate ViSNet potential for each of **21** unit types

6 hidden layers • **128**-dimensional embeddings


5 Å cutoff • maximum **32** neighbours

Energy / force loss weights: **0.05** / **0.95**

Random per-unit train / validation / test split: **80** / **10** / **10**

Backbone angles determine predicted NMR couplings



$$J(\phi) = 7.09 \cos^2(\phi - 60^\circ) - 1.42 \cos(\phi - 60^\circ) + 1.55$$



Curve plotted from the equation above


Average J over sampled backbone conformations.

J is measured in Hz; ϕ is a backbone dihedral angle.



Correlation and χ^2 errors provide complementary views of agreement.

Thermodynamic comparisons depend on ensembles and targets



ΔC_p (kcal mol ⁻¹ K ⁻¹)			
	AI ² BMD	MM	EXPERIMENT
Barnase	1.4	1.0	1.4
CI2	0.7	0.5	0.8

WW-domain T _m			
	AI ² BMD	MM	EXPERIMENT
WW domain	359.06 K	353.69 K	371 ± 2 K

WW-domain improvement leaves about 12 K discrepancy.

Runtime and chemical coverage define practical use

Protein	Atom number of protein	Atom number of system	AI ² BMD (s)	DPMD (s)	Allegro (s)	Tinker Amoeba (s)	Amber FF19S B (s)
Chignolin	175	4,715	0.047	0.040	0.238	0.117	0.004
Trp-cage	281	6,067	0.052	0.055	0.322	0.136	0.005
WW domain	571	10,678	0.070	0.095	0.626	0.196	0.008
ABD	746	11,793	0.085	0.106	0.712	0.208	0.008
PACSIN 3	1,040	17,923	0.106	0.162	-	0.292	0.011
SSO0941	2,450	44,401	0.213	0.414	-	0.699	0.027
APC	5,292	54,999	0.449	0.580	-	0.938	0.033
Polyphosphat e Kinase	11,404	97,657	0.966	-	-	1.487	0.058

ENERGY TIMING

Fig. 2f

SOLVATED PER-STEP TIMING

Extended Data Table 1

DEMONSTRATED CHEMISTRY

**Protein-unit vocabulary and
specified interactions**

New chemistry and rare events require additional validation.

Sources and reading

PAPER AND CODE

Wang et al. (2024). Ab initio characterization of protein molecular dynamics with AI²BMD. Nature 635, 1019–1027.

DOI: [10.1038/s41586-024-08127-z](https://doi.org/10.1038/s41586-024-08127-z)

Author Correction (2025). DOI: [10.1038/s41586-024-08556-w](https://doi.org/10.1038/s41586-024-08556-w)

Authors' code and data: github.com/microsoft/AI2BMD

FIGURES IN THIS DECK




Fig. 1
pages [1](#) • [2](#) • [3](#)




Fig. 2
pages [2](#) • [8](#) • [9](#)




Fig. 3
pages [10](#) • [18](#)




Fig. 4
pages [11](#)






Fig. 5
pages [12](#) • [19](#)



Extended Data Fig. 1
pages [6](#) • [17](#) • [23](#)



Extended Data Fig. 7
pages [4](#) • [16](#)

Protein	Atom number of protein	Atom number of system	AI ² BMD (s)	DPMD (s)	Allegro (s)	Tinker (s)	Amber (s)
Chignolin	175	4,715	0.047	0.040	0.238	0.117	0.004
Trp-cage	281	6,067	0.052	0.055	0.322	0.136	0.005
VW domain	571	10,678	0.070	0.095	0.626	0.196	0.008
ABD	746	11,793	0.085	0.106	0.712	0.208	0.008
PACSIN 3	1,040	17,923	0.106	0.162	-	0.292	0.011
SSO941	2,450	44,401	0.213	0.414	-	0.699	0.027
APC	5,292	54,999	0.449	0.580	-	0.938	0.033
Polyphosphatase Kinase	11,404	97,557	0.966	-	-	1.487	0.058


Extended Data Table 1
pages [20](#)

Supplementary slides


23 ViSNet architecture: graph, embedding, message-passing layers, output

Backups B1 to B6 are pages 16 to 21.

ViSNet is a message-passing graph network over nearby atoms



- 1 Graph.** Atoms are nodes. Edges join atoms within 5 Å, at most 32 neighbours per atom. Each unit has 12 to 36 atoms.
- 2 Embedding.** Atomic numbers and 3D coordinates give h_i^0 and f_{ij}^0 , 128 dimensions each, and $\vec{v}_i^0 = \vec{0}$. Higher-order spherical harmonics expand the coordinates.
- 3 ViSNet block $\times N$.** One message-passing layer, shown at right. Messages from neighbours j are summed at atom i ; the update block then refreshes h_i , f_{ij} and $\vec{v}_i$ using angle and dihedral terms (RGC). The paper reports 6 hidden layers.
- 4 Output.** The output block gives one energy $\hat{E}$. Forces $\vec{F} = -\partial\hat{E}/\partial\vec{r}$ are its gradient with respect to the input coordinates. One model for each of the 21 unit types.



SYMBOLS IN THE FIGURE AND STEPS

i, j target atom and neighbour (source) atom

Å ångström, 0.1 nm

l layer index; layer 0 is the embedding output

N number of stacked ViSNet blocks

h_i^l, h_j^l scalar features of atoms i and j

f_{ij}^l scalar features of edge $i-j$

$\vec{v}_i^l, \vec{v}_j^l$ vector features of atoms i and j

$\vec{0}$ zero vector, the starting value of $\vec{v}_i$

$\vec{r}_{ij}$ relative position of atoms i and j

ϕ_m^s, ϕ_m^v learned scalar and vector message functions

m_{ij}^l scalar message from atom j to atom i

Σ_j sum over the neighbours j of atom i

$m_i^l, \bar{m}_i^l$ summed scalar and vector messages

ϕ_{un}^s, ϕ_{un}^v learned update functions for atom features

ϕ_{ue}^s learned update function for edge features

RGC runtime geometry calculation

RGC-Angle dot product $\vec{v}_i \cdot \vec{v}_j$, angle information at atom i

RGC-Dihedral $\vec{v}_i \cdot \vec{v}_j$ after removing parts along $\vec{r}_{ij}$

$\Delta h_i^{l+1}, \Delta \vec{v}_i^{l+1}, \Delta f_{ij}^{l+1}$ updates for layer $l+1$

$\hat{E}, \vec{F}$ predicted energy and forces

$\partial\hat{E}/\partial\vec{r}$ gradient of $\hat{E}$ with respect to atom coordinates $\vec{r}$