BraintankAI

EXPLAINER · INTERATOMIC POTENTIALS · 5 MIN

Two ways of learning atomic forces mapped onto one blueprint

A 2025 study maps atom-centred descriptor models and equivariant neural networks onto one design space, then tests which parts of a leading network actually matter.

18 November 2026 · based on a paper in Nature Machine Intelligence

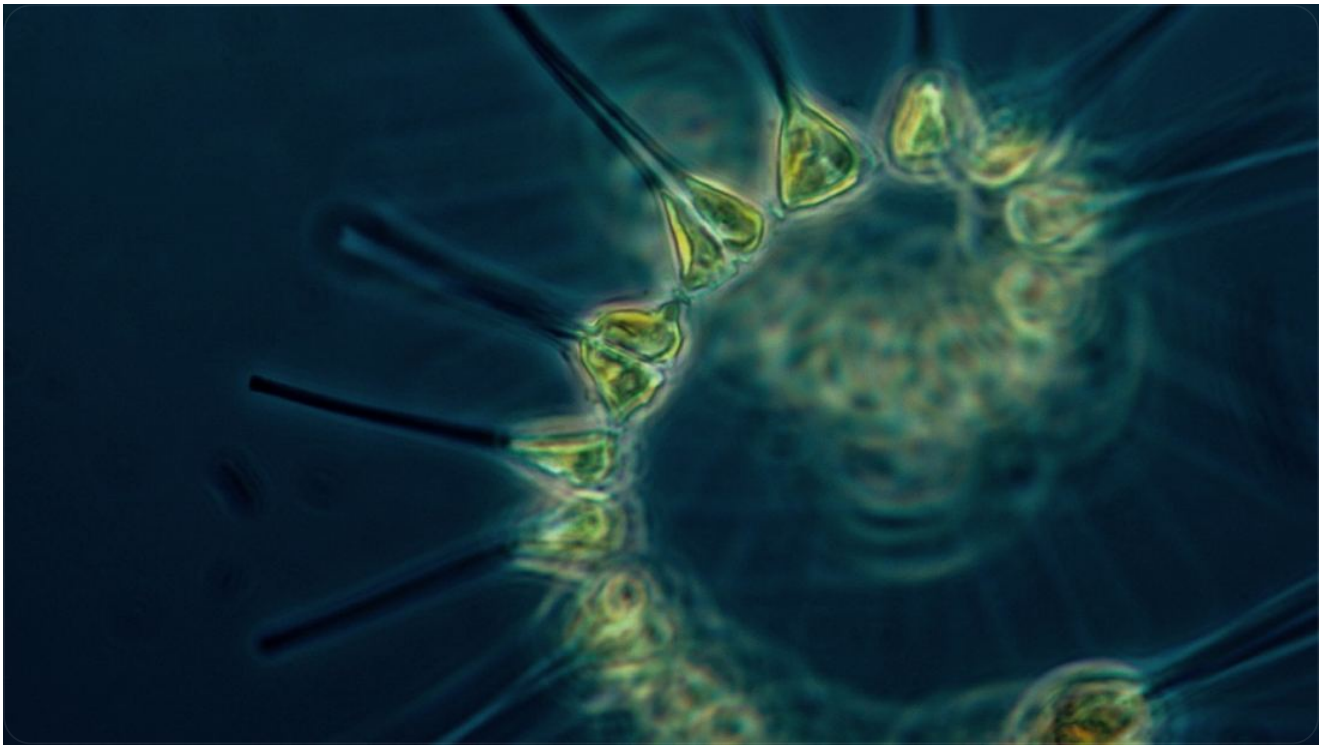


Photo: Noaa

To simulate how a drug molecule flexes or how a material holds together, you need the forces between atoms at every instant. Quantum-mechanical calculations supply those forces accurately but slowly. An interatomic potential is a shortcut. It is a function that takes atomic positions and returns energy and forces fast enough to run molecular dynamics, the step-by-step simulation of atoms in motion. Machine learning has become a common way to build such functions by fitting them to quantum calculations.

Yet machine-learned potentials developed along two separate lines. One builds fixed mathematical descriptions of each atom's surroundings. The other uses neural networks that pass messages between neighbouring atoms. New designs arrived quickly. In a 2025 paper in Nature Machine Intelligence, Ilyes Batatia and colleagues at institutions including the University of Cambridge and Harvard University argue that both lines sit in one design space. They then use that space to ask which ingredients of a leading network actually matter.

The idea

Any potential must respect basic physics. If you rotate, reflect or shift a molecule, its energy should not change, while directional quantities such as forces should turn with it. Models whose internal features transform consistently under these operations are called E(3)-equivariant, after the mathematical group of rotations, reflections and translations in three dimensions. The models here are also atom-centred: the total energy is a sum of contributions from each atom, each built from information about the surrounding atoms.

The first line of work describes an atom's neighbourhood with fixed functions. The Atomic Cluster Expansion (ACE) is a complete, systematic set of polynomial building blocks organised by body order, the number of atoms that interact in one term. A three-body term, for example, depends on triplets of atoms. Earlier work had shown that older descriptors such as SOAP are special cases of ACE. ACE uses a trick that was first developed for earlier descriptors. It computes a function for each pair of atoms, sums it over all neighbours, and then multiplies those sums together. This yields many-body features without looping over every triplet or quadruplet.

The second line is message-passing neural networks. Atoms become nodes in a graph, linked when they lie within the cutoff, and each layer updates every atom's features using those of its neighbours. Early versions such as SchNet used only quantities that do not change under rotation, such as distances. Equivariant successors also pass directional information between layers. The authors credit one of them, NequIP, with improving accuracy across datasets by about a factor of two.

How this paper puts it to work

The authors' framework, Multi-ACE, describes a message-passing network in which each layer's message is itself an ACE layer. For every pair of neighbours, each layer multiplies three things: a distance-dependent radial function, a spherical harmonic (a standard function describing direction), and an embedding of the atoms' elements and current features. These products are summed over neighbours. Multiplying ν such sums together gives correlation order ν, which equals body order $\nu+1$ because the central atom counts too. Fixed numbers called Clebsch–Gordan coefficients then combine the products so the result behaves correctly under rotation.

A model is then fixed by a few dials:

- the number of layers - the correlation order in each layer - the angular resolution - the shape of the radial functions - the nonlinearities - how chemical elements are handled

On the last dial, ACE couples elements, giving every combination its own basis function, so the count explodes as elements are added. Most message-passing networks instead use learned element embeddings of fixed length. The authors place 13 published models on this map. With one layer, the framework reduces to linear ACE. NequIP uses two or more layers with correlation order one. Most published message-passing models, the authors note, take this many-layer, low-correlation route, whereas descriptor models use few layers with high correlation order.

The authors then show that, if the network is assumed to be linear, a message-passing model with T layers can be viewed as a pared-down version of a single ACE layer with a cutoff T times larger. It keeps only chain-like paths that hop from neighbour to neighbour towards the central atom, rather than every cluster in the wider sphere. The authors argue this tracks chemical bonding more naturally.

They also show that common nonlinear functions destroy a clean body order. Their simplified NequIP, called BOTNet (Body-Ordered Tensor Network), therefore removes nonlinear activations from its first five message-passing layers. It adds one nonlinear readout at the end and uses a radial basis that depends on each neighbour's element.

What it shows, and what it does not

The main experiments use a single benchmark, 3BPA, a flexible drug-like molecule. Models were trained on configurations at 300 K and tested at 300 K, then at the hotter 600 K and 1,200 K to probe extrapolation. Force errors below are root-mean-square values in meV/Å, thousandths of an electronvolt per ångström.

Widening NequIP helped. At 300 K, force error fell from 12.9 for the smallest model, with 437,336 parameters, to 10.6 at 3,415,832 parameters. The largest model, with 11,580,440 parameters, scored 10.7. The authors report the same direction at the higher temperatures, where every model's errors were several times larger. Only one model size was trained from three random starting points, so the other results come from single runs.

Activation functions, the nonlinear operations inside a network, mattered. Swapping NequIP's SiLU function for tanh made results significantly worse, worse even than using no activation at all. The authors suggest this is probably because tanh flattens out for large inputs, starving training of gradient signal. In BOTNet, the final nonlinear readout significantly improved on a purely linear version.

The authors also report results from supporting experiments. They find that BOTNet matches NequIP's accuracy on benchmark datasets. Learned element embeddings give sensible predictions for pairs of elements never seen together in training. And the element-dependent radial basis fits the data better, while NequIP's version holds up better in extreme cases such as bond breaking.

The limits are clear. The main tests rest on one molecule trained at one temperature. The activation comparison mixes two code bases, different library versions and different normalisation choices. The equivalence between message passing and ACE holds only under the linearity assumption, and nonlinearities were left out of the design-space comparison for simplicity. The best degree of element coupling is posed as an open question. The paper's main contribution is a shared language. The authors note that a later model, MACE, builds on it and reaches high accuracy with two layers at reduced computational cost.

PAPER	*The design space of E(3)-equivariant atom-centred interatomic potentials*
AUTHORS	Ilyes Batatia, Simon Batzner, Dávid Péter Kovács, Albert Musaelian, Gregor N. C. Simm, Ralf Drautz, Christoph Ortner, Boris Kozinsky, Gábor Cśanyi et al.
VENUE	Nature Machine Intelligence · 15 January 2025
LINK	doi:10.1038/s42256-024-00956-x
INDEX	stacks.braintank.ai