

# Physics-Guided Equivariant Neural Networks for Protein Microenvironment Modeling and pKa Prediction

George Eerguson

Department of Computer Science, Binghamton University, Binghamton, NY, USA.  
ferguson1993@binghamton.edu

Mailk Kalkarna

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.  
mihirwork@oregonstate.edu

Logan Vega

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.  
loganv@unr.edu

Kenneth Riley

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.  
rileykenneth@buffalo.edu

## Abstract

Predicting the pKa values of ionizable residues in proteins is a fundamental challenge in biophysics, with profound implications for drug design, enzyme engineering, and understanding biological mechanisms. The accuracy of pKa prediction depends critically on the faithful representation of the protein microenvironment, a three-dimensional chemical context characterized by electrostatic fields, hydrogen-bonding networks, and solvent exposure. Recent advances in geometric deep learning, particularly equivariant neural networks that respect the symmetries of physical space, offer a promising avenue for learning protein microenvironments directly from atomic coordinates. This paper presents a system-level analysis of physics-guided equivariant neural networks for microenvironment modeling and pKa prediction. We examine the architectural integration of SE(3)-equivariant message passing with physically inspired features such as continuum electrostatics and solvation free energies, discussing trade-offs between model expressiveness, interpretability, and computational efficiency. Beyond algorithmic design, we address the broader infrastructure required for training and deploying these models, including data curation from structural databases, high-performance computing considerations, and robustness to conformational variability. Furthermore, we explore governance and policy dimensions: the imperative for standardized benchmarks, fairness across diverse protein families, reproducibility through open-source release, and the ethical implications of precisely predicting protein properties that could be misused. We also consider sustainability metrics and the potential of amortized inference to mitigate the carbon footprint of large-scale training. By connecting architectural innovation with system-level thinking, this work charts a roadmap for responsible

development of physics-guided equivariant models that can transform molecular therapeutics while adhering to principles of equity and scientific rigor.

## Keywords

Equivariant neural networks, protein pKa prediction, microenvironment modeling, deep learning, biophysics, SE(3) equivariance, physics-guided machine learning.

## 1. Introduction

The ionization state of amino acid side chains within a protein is a dynamic property that governs folding stability, enzymatic catalysis, molecular recognition, and pH-dependent conformational switches. Accurate prediction of pKa values—the pH at which a titratable group is half-protonated—remains a cornerstone of computational biophysics. Traditional physics-based methods, such as empirical models exemplified by PROPKA [1] and continuum electrostatics approaches like those implemented in H++ [2], have served as workhorses for decades. These methods approximate the complex electrostatic interactions in a protein through simplified representations of solvent screening and dielectric heterogeneity. While computationally efficient and often interpretable, they rely on handcrafted features and linear response assumptions that can fail in buried or highly coupled active sites. The advent of deep learning has shifted the paradigm toward data-driven representation learning. Early adopters employed three-dimensional convolutional neural networks to model local structural motifs [3], capturing steric and chemical patterns from volumetric grids. More sophisticated geometric deep learning architectures evolved to operate on graph representations of protein surfaces and atomic neighborhoods, learning interaction fingerprints directly from point clouds [4]. These models demonstrated improved accuracy but lacked explicit physical constraints, leading to potential violations of fundamental symmetries such as rotation and translation invariance.

The emergence of equivariant neural networks has addressed this limitation. Models like Tensor Field Networks [5] and SE(3)-Transformers [6] embed rotational equivariance into their message-passing operations, ensuring that predictions transform consistently with the input structure under rigid motions. This inductive bias dramatically improves data efficiency and generalization, as demonstrated by the E(3)-equivariant interatomic potentials in NequIP [7]. In the protein domain, the invariant point attention module in AlphaFold2 [8] leverages SE(3) equivariance for end-to-end structure prediction, showcasing the power of symmetry-aware architectures. Translating these successes to pKa prediction requires careful consideration of the protein microenvironment, a concept that extends beyond local geometry to encompass electrostatic potentials, hydrogen-bond energies, and solvent accessibility. A physics-guided approach aims to marry the representation capacity of equivariant networks with physically meaningful features, producing hybrid models that respect both data and domain knowledge. The present work does not propose a novel algorithm but rather provides a comprehensive system-level examination of the design, deployment, and governance of such physics-guided equivariant neural networks. We critically assess the architectural decisions that balance predictive performance with physical consistency, the computational infrastructure needed to sustain large-scale training, and the socio-technical implications of deploying these models in drug discovery pipelines. By framing the problem as one of system design, we illuminate the trade-offs and responsibilities that accompany the development of advanced molecular modeling tools.

## 2. Background and Motivation

The accurate prediction of protein pKa values from structure alone is an enduring challenge because ionization equilibria are influenced by a multitude of coupled effects. The local dielectric environment, the proximity of permanent charges and dipoles, desolvation penalties upon folding, and the coupling of multiple titratable sites all contribute to shifts from model compound pKa values. Empirical predictors have long exploited statistical relationships between structure-derived descriptors and experimental pKa shifts [1]. Meanwhile, physics-based approaches solve the Poisson–Boltzmann equation or conduct free energy calculations, providing physically grounded estimates but at the cost of substantial computational effort and sensitivity to the choice of protein dielectric constant. The continuing influx of experimentally determined pKa data from databases such as PKAD [10] has created fertile ground for supervised learning. Early deep learning efforts such as DeepKa [9] used voxelized representations and 3D CNNs to learn from local structural neighborhoods. However, these grid-based methods discretize space, breaking rotational symmetry unless heavy data augmentation is applied. Graph-based architectures circumvent this issue by operating on invariant features such as pairwise distances and angles, forming a natural bridge to equivariant designs.

The key motivation for physics-guided equivariant networks is the desire to encode physical laws directly into the model architecture while still learning complex patterns from data. Unlike classical force fields, which are parameterized through expert knowledge, a neural network can discover non-additive and many-body effects that matter for protonation states. However, without physical guidance, a purely data-driven model may misinterpret the role of buried water molecules, ignore the long-range nature of electrostatic interactions, or fail to extrapolate to protein folds unseen during training. The incorporation of precomputed electrostatic potential maps, continuum solvation energies, or hydrogen-bond counts as auxiliary input features or as training targets has been shown to improve generalization [9], [16]. Indeed, a recent study employing graph-based deep learning with physically inspired feature engineering achieved state-of-the-art pKa prediction accuracy on multiple benchmark sets [16]. Such hybrid strategies align with the broader trend of augmenting equivariant networks with physical priors, as seen in interatomic potentials where forces and energies are jointly learned [7]. The challenge lies in designing an architecture that seamlessly integrates scalar and vector features under SE(3) constraints while allowing for the flexible injection of domain knowledge without compromising end-to-end differentiability.

### **3. Physics-Guided Equivariant Architectures: Design Principles and Trade-offs**

Designing a neural network that is both equivariant to three-dimensional rotations and translations and responsive to long-range electrostatic interactions is a non-trivial engineering problem. The core building blocks of modern equivariant architectures are derived from the theory of steerable features and tensor products. In practice, a protein microenvironment is represented as a local point cloud centered on a titratable residue, where each neighbor atom carries a feature vector that transforms under rotations according to its irreducible representation. Message passing proceeds via convolution operations that compute tensor products between node features and spherical harmonic filters, ensuring that output features respect the desired transformation rules. Physics guidance can be integrated at multiple levels. At the input level, scalar features such as partial charges, Born radii, or solvation free energies from a Poisson–Boltzmann calculation [11] can be concatenated to the invariant channels. Vector features can be initialized from the electrostatic field gradient, providing directional information about charge distributions. At the intermediate level, the network can be

regularized by auxiliary heads that predict local electrostatic potentials or hydrogen-bond energies, encouraging hidden representations to capture physically relevant interactions. Finally, at the output level, the predicted pKa value can be constrained to lie within physically plausible bounds and linked to a thermodynamic cycle that includes desolvation and Coulombic contributions.

A central design trade-off is between the expressiveness of high-order tensor features and the computational and memory footprint of the network. Including tensors of order higher than the scalar level permits the modeling of anisotropic charge distributions but dramatically increases the number of channels and the cost of tensor product operations, especially when applied to atomic systems with many neighbors. Strategies such as using the envelope function of NequIP [7] to localize interactions, or employing attention mechanisms like in SE(3)-Transformers [6] to selectively weight neighbors, provide adjustable complexity. Another critical consideration is the treatment of solvent. Explicit water molecules are often absent in static crystal structures, yet buried water molecules can mediate proton transfer and significantly shift pKa values. Physics-guided models may incorporate an implicit solvent model's reaction field as a feature or learn to predict water occupancy as an auxiliary task, effectively encoding solvation thermodynamics in an equivariant representation. The system architect must decide whether to precompute continuum electrostatics grids, which may be computationally expensive and break rotational equivariance at the grid level, or to compute equivariant approximations on the fly using fast multipole-like learned functions. The choice influences both training speed and inference latency in high-throughput screening pipelines. From a system perspective, the integration of physics modules often requires interfacing deep learning frameworks with molecular mechanics software, raising software engineering challenges in maintaining reproducible environments and handling versioning of force field parameters. Containerized deployment using technologies like Docker and workflow managers such as Snakemake have become essential for robust pipelines that can be replicated across institutional clusters.

The training data itself poses a system design puzzle. High-quality experimental pKa data are biased toward certain well-studied proteins and ionizable residue types, with cysteine, tyrosine, and histidine often underrepresented relative to aspartate and glutamate. This imbalance can propagate into model predictions, leading to fairness issues where pKa shifts for less common residue types are systematically less accurate. Data augmentation through physics-based simulation, such as constant-pH molecular dynamics, can generate synthetic training examples that help balance the dataset. Methods that incorporate extensive conformational sampling, as advanced in tools like MCCE2 [14], provide routes to generate physically plausible protonation states across alternative side-chain rotamers and backbone fluctuations. However, synthetic data may introduce systematic biases if the underlying force field is inaccurate, leading to a data governance problem regarding the weighting of experimental versus simulated data. Multi-task learning frameworks that jointly predict pKa, stability changes upon mutation, and catalytic activity have been proposed to leverage correlated biophysical data, but the optimal strategy for balancing losses and avoiding negative transfer remains an open system design question.

#### **4. System-Level Infrastructure and Deployment**

Deploying physics-guided equivariant neural networks at scale requires a well-orchestrated infrastructure that spans data ingestion, model training, and inference serving. The primary data source remains the Protein Data Bank, from which experimentally resolved structures

must be cleaned, missing atoms reconstructed, and protonation states assigned. Tools such as PDB2PQR [11] and PROPKA [1] can generate initial protonation assignments and electrostatic features, feeding into a preprocessing pipeline that extracts local microenvironments as graph tensors. The volume of data is considerable; with over 200,000 structures in the PDB, each containing hundreds of residues, the total number of microenvironment samples can reach millions. Managing this data lifecycle demands robust version control, metadata tracking, and compliance with FAIR principles [20]. Training such networks typically leverages distributed GPU clusters, with mixed-precision training and model parallelism to handle large protein graphs. The carbon footprint of training a state-of-the-art equivariant model with physics features can be substantial, comparable to that of large language models. Amortizing this cost over many downstream predictions through transfer learning is a critical sustainability strategy. A pre-trained microenvironment encoder can be fine-tuned for pKa prediction, binding affinity estimation, or mutation effect scoring, reducing the need for repeated costly training cycles [18]. However, the transferability of learned representations across tasks depends on the alignment of physical features; a model trained with electrostatic features specific to pKa may not transfer well to tasks dominated by dispersion forces.

Inference deployment in industrial drug discovery must meet stringent latency requirements. High-throughput virtual screening scenarios demand pKa evaluation for millions of compounds or mutants in a matter of hours. This necessitates optimized inference engines that can batch graph computations and exploit hardware acceleration such as TensorRT or custom CUDA kernels. The integration of equivariant models into existing CADD platforms involves building microservices that expose gRPC or REST endpoints, with attention to fault tolerance and load balancing. Given that protein structure predictions from AlphaFold2 [8] are now widely available, the pKa prediction pipeline can ingest both experimental and predicted structures, but the model must be robust to structural inaccuracies. Robustness can be enhanced by input dropout during training, ensemble methods, and by providing confidence scores through Bayesian approximations or deep ensembles [19]. A systematic evaluation of robustness to alternative conformations, pH-dependent structural changes, and crystal packing artifacts is essential for clinical and regulatory acceptance.

Governance of a deployed pKa prediction system involves continuous monitoring and model updating as new experimental data become public. A feedback loop that retrains the model on emerging data, such as newly characterized mutants or pH-dependent functional assays, can maintain accuracy and mitigate concept drift. However, this raises the question of version control for predictive models used in regulatory submissions. If a model's outputs influence drug candidate selection or safety assessment, its provenance, training data, and validation protocols must be transparently documented, aligning with guidelines such as those proposed in the DOME framework [13]. Furthermore, the community must converge on standardized benchmark tasks and evaluation metrics. While root-mean-square error and correlation coefficients are common, they do not capture the clinical relevance of pKa errors in specific pH ranges. Developing a suite of benchmarks that span membrane proteins, intrinsically disordered regions, and large macromolecular assemblies will ensure that model improvements translate to real-world impact.

## **5. Fairness, Governance, and Policy Implications**

Bias in pKa prediction models is an underexplored facet of algorithmic fairness in computational biology. If training data skew toward soluble globular proteins, the model may

exhibit degraded performance on membrane proteins, which have a distinct hydrophobic environment and dielectric profile. This disparity could inadvertently disadvantage research on ion channels, transporters, and G-protein-coupled receptors—protein classes that are among the most important drug targets. Ensuring representational parity in training data requires deliberate collection and curation of pKa measurements from underrepresented families, a goal that intersects with broader open science initiatives and structural genomics efforts. Policy mechanisms, such as funder mandates for data deposition and model sharing, can accelerate progress. The open-source release of model weights, training code, and preprocessing pipelines under permissive licenses fosters reproducibility and allows independent auditing for biases. Nevertheless, dual-use concerns must be addressed: the ability to precisely tune protonation states could be misapplied to design toxins or optimize the stability of harmful enzymatic functions. While the immediate risk may be low, an anticipatory governance framework suggests that pKa prediction tools should be coupled while the immediate risk may be low, an anticipatory governance framework suggests that pKa prediction tools should be coupled with responsible use guidelines and educational outreach that emphasize the humanitarian applications of protein engineering. A dual-use analysis of artificial intelligence for drug discovery has highlighted that predictive models enabling precise manipulation of biomolecular properties could be misappropriated for the design of harmful agents [22]. Although the protein design landscape currently requires extensive multidisciplinary expertise and wet-lab validation, the accelerating pace of generative modeling for biomolecules could shorten the path from computational prediction to synthesis. Addressing such risks proactively involves the establishment of review boards within academic and industrial institutions, as well as the incorporation of ethical risk assessment into the publication and model release process. The broader governance ecosystem must also confront the subtle structural inequities that arise when computational tools are optimized exclusively for high-resource research settings. Ensuring that model validation includes proteins from organisms and environments underrepresented in the PDB, such as extremophiles, viral glycoproteins, and membrane complexes, is both a scientific and a fairness imperative. This extends to the geographical distribution of computational resources; cloud-based deployment options and precomputed pKa databases can democratize access for researchers at lower-resourced institutions, provided that funding agencies and industrial stakeholders prioritize open-access infrastructure.

Policy measures at the intersection of open science and model governance can accelerate the maturation of the field. The FAIR principles for data management [20] should be extended to include model cards and datasheets that detail the provenance of training data, known limitations, and performance characteristics across different protein architectures. The DOME framework [13] already provides a methodological scaffold for transparent machine learning reporting in biology, and its adoption within pKa prediction efforts would improve reproducibility. However, mere transparency is insufficient without intentional efforts to monitor and correct for distributional shifts induced by the ever-growing corpus of protein structures. The growing reliance on AlphaFold2-predicted structures raises a recursive validation problem: models trained and evaluated on the same predicted conformational ensembles may exhibit inflated accuracy metrics that do not translate to experimental reality. Governance mechanisms should incentivize continuous revalidation against newly measured pKa values, using platforms such as the PKAD database [10] as a living benchmark. The development of community-wide blind prediction challenges, analogous to the SAMPL series

for small molecules [15], would provide an external check on overfitting and spur methodologically diverse solutions.

The sustainability of pKa prediction platforms further touches on computational equity. Large-scale equivariant model training incurs a carbon footprint that must be weighed against its scientific return. While pre-training on general protein microenvironments amortizes this cost [18], the upfront energy investment is substantial. Reporting energy consumption and model efficiency metrics alongside accuracy benchmarks, as advocated by work on Green AI [24], enables the community to make informed trade-offs. Additionally, edge inference on modest hardware, enabled by distilled or quantized equivariant models, can reduce the dependence on power-hungry GPU fleets during deployment. The interplay of algorithmic efficiency, hardware design, and policy instruments such as carbon-aware scheduling in academic computing centers will shape the long-term viability of physics-guided deep learning for biophysics.

In the pursuit of unbiased and robust models, the technique of deep ensembles provides a practical means to estimate predictive uncertainty for individual residues, guiding expert users toward high-confidence predictions and flagging cases that require higher-level theory or experimentation [19]. Uncertainty quantification is not merely a methodological nicety; in a clinical or industrial context where decisions regarding lead optimization have significant financial and safety implications, knowing when a model is likely to be wrong is as important as its average accuracy. The development of principled uncertainty baselines on a diverse set of protein families, including those rich in metal-binding sites and allosteric networks, remains a frontier. Further, fairness audits that disaggregate performance by residue type, secondary structure context, and solvent exposure can illuminate hidden biases. Such audits could reveal, for example, that pKa shifts for histidine in catalytic triads are systematically underestimated because the training data lack sufficient examples of coupled protonation rearrangements. Addressing these blind spots may require targeted data collection campaigns or the integration of quantum chemical calculations for underrepresented microenvironments, a hybrid experimental-computational loop that exemplifies the need for interdisciplinary governance structures bridging life sciences, computer science, and ethics [23].

## 6. Conclusion

This paper has traversed the system-level landscape of physics-guided equivariant neural networks for pKa prediction, articulating both the architectural ingenuity and the broader infrastructural and governance challenges that accompany their deployment. The marriage of SE(3)-equivariant message passing with physically inspired features drawn from continuum electrostatics, solvation theory, and molecular mechanics holds the promise of pKa prediction models that are simultaneously data-efficient, generalizable, and physically interpretable. Yet the realization of this promise is contingent upon navigating a series of structural trade-offs. The choice between high-order tensor features and computational tractability, the integration of implicit solvent representations that respect rotational symmetry, the curation of balanced and representative training data, and the design of inference pipelines that meet throughput demands all demand cross-disciplinary systems thinking. The infusion of domain knowledge at multiple levels of the learning architecture—from input featurization through auxiliary training objectives to output constraints—emerges as a design principle that balances the flexibility of deep learning with the rigor of physical chemistry.

Beyond the technical architecture, the sustainability and fairness of these systems demand attention. The computational expense of training equivariant models on large macromolecular

datasets must be justified by the downstream value, and amortization through pre-training combined with energy-efficient deployment strategies can moderate the environmental impact. Fairness considerations require deliberate efforts to compile and maintain pKa datasets that encompass the full diversity of protein space, ensuring that model performance does not become yet another source of bias in drug discovery pipelines that are already skewed toward well-funded therapeutic areas. Open-source release, transparent model documentation, and community benchmarking form the bedrock of reproducible and equitable science. Governance frameworks must evolve to address the dual-use potential of accurate biomolecular property prediction, embedding responsible innovation practices into the research lifecycle without stifling scientific progress.

Looking ahead, the frontier of pKa prediction will likely be shaped by generative models that sample the thermodynamic ensemble of protonation states and coupled conformational changes, moving beyond static structure prediction toward a holistic description of pH-dependent dynamics. Boltzmann generators [25] and diffusion models for molecular systems point toward a future where microenvironment modeling and pKa prediction are integrated into end-to-end frameworks capable of sampling rare protonation events and allosteric coupling. Physics guidance will remain essential as the complexity of these models grows, providing inductive biases that steer learning toward physically plausible configurations. The convergence of equivariant deep learning, physical simulation, and experimental data under a shared governance charter offers a model for the responsible development of molecular AI systems that can genuinely accelerate therapeutics while upholding scientific integrity and societal trust.

## References

1. Olsson, M. H. M., Søndergaard, C. R., Rostkowski, M., & Jensen, J. H. (2011). PROPKA3: Consistent treatment of internal and surface residues in empirical pKa predictions. *Journal of Chemical Theory and Computation*, 7(2), 525–537.
2. Anandkrishnan, R., Aguilar, B., & Onufriev, A. V. (2012). H++ 3.0: automating pK prediction and the preparation of biomolecular structures for atomistic molecular modeling and simulations. *Nucleic Acids Research*, 40(W1), W537–W541.
3. Torng, W., & Altman, R. B. (2017). 3D deep convolutional neural networks for amino acid environment similarity analysis. *BMC Bioinformatics*, 18(1), 302.
4. Fout, A., Byrd, J., Shariat, B., & Ben-Hur, A. (2017). Protein interface prediction using graph convolutional networks. *Advances in Neural Information Processing Systems*, 30, 6530–6539.
5. Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., & Riley, P. (2018). Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. *arXiv preprint arXiv:1802.08219*.
6. Fuchs, F., Worrall, D., Fischer, V., & Welling, M. (2020). SE(3)-transformers: 3D roto-translation equivariant attention networks. *Advances in Neural Information Processing Systems*, 33, 1970–1981.
7. Batzner, S., Musaelian, A., Sun, L., Geiger, M., Mailoa, J. P., Kornbluth, M., Molinari, N., Smidt, T. E., & Kozinsky, B. (2022). E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13, 2453.

8. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Židek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., ... Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596, 583–589.
9. Shang, Y., Li, H., Zhang, H., & Zhou, H. (2022). DeepKa: A deep-learning-based protein pKa prediction model. *Journal of Chemical Information and Modeling*, 62(11), 2719–2730.
10. Pahari, S., Sun, L., & Alexov, E. (2019). PKAD: a database of experimentally measured pKa values of protein ionizable residues. *Bioinformatics*, 35(17), 3189–3191.
11. Dolinsky, T. J., Nielsen, J. E., McCammon, J. A., & Baker, N. A. (2004). PDB2PQR: an automated pipeline for the setup of Poisson–Boltzmann electrostatics calculations. *Nucleic Acids Research*, 32(Web Server), W665–W667.
12. Radak, B. K., Chipot, C., Suh, D., Jo, S., Jiang, W., Phillips, J. C., Schulten, K., & Roux, B. (2017). Constant-pH molecular dynamics simulations for large biomolecular systems. *Journal of Chemical Theory and Computation*, 13(12), 5933–5944.
13. Walsh, I., Pollastri, G., & Tosatto, S. C. E. (2021). DOME: recommendations for supervised machine learning validation in biology. *Nature Methods*, 18, 1122–1127.
14. Song, Y., Mao, J., & Gunner, M. R. (2009). MCCE2: improving protein pKa calculations with extensive side chain rotamer sampling. *Journal of Computational Chemistry*, 30(14), 2231–2247.
15. Schiffer, S., Ho, J., Grossfield, A., & Mobley, D. L. (2022). SAMPL8 pKa challenge: assessing computational methods for protein pKa prediction. *Journal of Computer-Aided Molecular Design*, 36, 459–478.
16. Song, Z., Wang, R., Jiao, X., & Huang, Z. (2026). Graph-Based Deep Learning Models for Predicting p K a Values of Protein-Ionizable Residues via Physically Inspired Feature Engineering. *Journal of Chemical Information and Modeling*.
17. Rives, A., Meier, J., Sercu, T., Goyal, S., Lin, Z., Liu, J., Guo, D., Ott, M., Zitnick, C. L., Ma, J., & Fergus, R. (2021). Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15), e2016239118.
18. Zhang, Z., Xu, M., Gabizo, A., & Berger, B. (2023). Protein representation learning by geometric structure pretraining. *International Conference on Learning Representations*.
19. Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 30, 6402–6413.
20. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., ... Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, 3, 160018.
21. Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature*, 533, 452–454.

22. Urbina, F., Lentzos, F., Invernizzi, C., & Ekins, S. (2022). Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence*, 4, 189–191.
23. Chen, I. Y., Pierson, E., Rose, S., Joshi, S., Ferryman, K., & Ghassemi, M. (2021). Ethical machine learning in healthcare. *Annual Review of Biomedical Data Science*, 4, 123–144.
24. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.
25. Noé, F., Olsson, S., Köhler, J., & Wu, H. (2019). Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457), eaaw1147.