

Protein Folding: Recent Advances and Analysis in Computing

Boyuan Wang^{1*}

¹School of Science, University of Melbourne, Melbourne, Australia

Abstract. Protein folding is a fundamental yet elusive problem: a protein's three-dimensional structure determines its function, but misfolding underlies disorders such as Alzheimer's. Experimental techniques like X-ray crystallography and NMR resolve structures but cannot keep pace with the explosive growth of sequence data. Consequently, computational approaches – from simplified hydrophobic–polar (HP) lattice models to deep neural networks – have become indispensable. This paper reviews recent advances in computational protein folding, using the HP model as a conceptual test bed. It surveys classic heuristics, modern deep reinforcement learning, variational generative techniques and emerging quantum algorithms for NP-hard lattice models, and compares them with breakthroughs in all-atom structure prediction exemplified by AlphaFold and RosettaFold. Benchmark datasets, evaluation metrics and ongoing challenges – such as data bias, dynamic folding and integration of physical constraints – are discussed. The review concludes that future progress will likely come from hybrid methods that combine machine-learning flexibility with physics-based priors, expanded and more diverse structural data sets, and algorithmic innovations, including quantum-inspired heuristics and efficient hardware. Such advances could enable more accurate folding predictions, facilitate rational drug design and deepen our understanding of protein misfolding diseases.

1 Introduction

Predicting how a linear sequence of amino acids folds into a functional three-dimensional structure remains one of the grand challenges in theoretical biochemistry. Misfolded proteins are linked to diseases ranging from Alzheimer's to amyloidosis, and understanding the folding pathway is essential for rational drug design and protein engineering. Although experimental techniques such as X-ray crystallography and nuclear magnetic resonance provide high-resolution structures, they are slow and expensive; only a fraction of the millions of known protein sequences have been solved experimentally. To bridge this gap, computational modelling has become indispensable. Simplified lattice models, coarse-grained representations and all-atom simulations allow researchers to explore folding pathways and energy landscapes. Among these, the hydrophobic–polar (HP) model has been widely used as a conceptual framework: it reduces the twenty amino acids to hydrophobic (H) and polar (P) categories and represents folding as a self-avoiding walk on a lattice. Despite its simplicity, finding the minimum-energy conformation in the HP model is NP-complete, motivating decades of algorithmic innovation.

Early work applied heuristic algorithms such as genetic algorithms, Monte-Carlo simulations, memetic algorithms and ant colony optimisation to navigate the vast conformational space. The pruned-enriched Rosenbluth method and replica-exchange Monte-Carlo became reference methods for sampling low-energy states. Recently, reinforcement learning (RL) has

reinvigorated research on HP models. Espitia and colleagues formulated protein folding as a Markov decision process and introduced two deep RL architectures that learn to place residues on a 3-D lattice. A hybrid model combining random reservoir projections with trainable deep layers achieved optimal conformations using 25 % fewer training episodes for short sequences, while a long short-term memory network with multi-head attention matched the best-known energies for longer chains [1]. Liu and Iba extended this idea by integrating transformer-style attention into Q-learning along with symmetry-breaking constraints, duelling/double Q-learning and prioritised replay. Their attention-augmented framework achieved several known best energies for standard benchmark sequences and approached optimal solutions for longer chains [2]. These works demonstrate that RL can explore the HP conformational landscape efficiently without explicit heuristics.

Other machine-learning techniques also show promise. Khandoker et al. introduced a variational annealing scheme that trains dilated recurrent neural networks to autoregressively generate valid HP folds. By masking invalid states and using temperature-like fluctuations, their method reached ground states for sequences up to 60 monomers and provided a generalisable upper-bound training strategy [3]. In parallel, heuristic optimisation continues to evolve. Boumedine and Bouroubi combined the Cuckoo Search algorithm with hill climbing to produce a hybrid meta-heuristic that improved search efficiency on the

* Corresponding author: boyuanwang@student.unimelb.edu.au

3-D HP model [4]. Roucairol and Cazenave proposed a nested Monte-Carlo search for the HP model that, although not surpassing state-of-the-art sampling methods, highlights the importance of exploring diverse Monte-Carlo strategies [5]. Advances in quantum computing offer a complementary path. Classical algorithms struggle to exhaustively sample the HP energy surface, yet small quantum computers can encode spin systems that map onto folding problems. Vasavi et al. designed a turn-based encoding algorithm that uses the HP model as a framework for folding shorter sequences on gate-based quantum computers [6]. Irbäck and colleagues developed a distributed spin representation that allows HP folding to be solved on an Ising-type quantum annealer without auxiliary variables [7]. Their method reproduced known ground states up to 30 residues and recovered lowest known energies for 48- and 64-residue chains, outperforming classical simulated annealing. Earlier, Robert et al. devised a resource-efficient variational quantum algorithm with polynomial scaling, combining quantum variational circuits and evolutionary strategies to fold small peptides on IBM hardware [8]. These initiatives underscore the potential of quantum annealing and hybrid quantum-classical algorithms to tackle NP-hard folding problems.

Beyond simplified models, deep learning has revolutionised full-atom protein structure prediction. AlphaFold 2, released by DeepMind in 2021, achieved near-atomic accuracy for 98.5 % of human proteins and outperformed template-based methods even when no homologous structures were available. RosettaFold, introduced the same year, uses a three-track neural network to simultaneously reason over sequence, distance matrix and three-dimensional geometry, producing accurate structures in minutes. These breakthroughs show that deep neural networks can implicitly learn physical principles from large datasets, dramatically shortening prediction time. However, recent reviews highlight that such models remain limited by biases in the Protein Data Bank; they often memorise static, thermodynamically stable conformations and struggle with dynamic or fold-switching proteins. Expanding databases to include more transient states and integrating physics-based constraints with learning algorithms are pressing needs. In summary, protein folding research is entering a rich era where simplified lattice models, reinforcement learning, heuristic optimisation, quantum algorithms and deep neural networks coexist. This review synthesises developments over the last five years, focusing on the HP model as a conceptual testbed while connecting to breakthroughs in full-scale prediction. By comparing methodologies, datasets and evaluation metrics, and by discussing current limitations and future directions, the article provides a comprehensive perspective for researchers seeking to navigate the rapidly evolving landscape of computational protein folding.

2 Overview of Computational Approaches

Computational modelling of protein folding spans a spectrum from simplified lattices to high-fidelity simulations. This section summarises mainstream approaches and serves as an overview to orient the reader before delving into details.

2.1 Traditional models and classical heuristics

Direct simulation of atomic interactions is computationally expensive, so researchers often start with simplified lattice models. The hydrophobic–polar model treats amino acids as either hydrophobic or polar beads on a square or cubic lattice. The goal is to find a self-avoiding conformation that maximises hydrophobic contacts, but computing the optimal arrangement is NP-complete. To cope with this combinatorial explosion, a variety of heuristic algorithms have been developed. Beam search maintains a limited number of partial conformations, while simulated annealing and genetic algorithms use stochastic moves and evolutionary operators to escape local minima. Nested Monte-Carlo search combines random roll-outs with selection, and branch-and-bound uses bounds on achievable energy to prune conformations that cannot improve on current best solutions. Self-avoiding walk methods such as the pruned-enriched Rosenbluth method build chains step by step, pruning unlikely paths and enriching promising ones. These heuristics fold short sequences effectively but face scaling challenges because the number of conformations grows exponentially and they often rely on carefully tuned parameters.

2.2 Machine-learning methods

Machine learning provides a way to learn folding strategies from experience rather than hand-crafted heuristics. Reinforcement learning. Folding Zero pioneered deep reinforcement learning for the 2-D HP model by coupling a neural network with Monte-Carlo tree search; through self-play the agent learned to orient hydrophobic residues inward and avoid dead-end conformations [9]. The network predicts both the value of a partial conformation and a policy over possible moves; these predictions guide tree search, balancing exploitation of known good moves with exploration of novel conformations. Folding Zero outperformed several heuristic baselines and demonstrated that RL can extract useful folding patterns without any supervised data. Recent work extended RL to three dimensions. Espitia et al. designed two architectures: a hybrid reservoir–deep network and a long short-term memory network with multi-head attention, both trained with Deep Q-Learning. The hybrid model achieved optimal conformations with 25 % fewer training episodes, while the attention-based network matched best known energies for longer chains [1]. Liu and Iba incorporated transformer-style attention into Q-learning and added symmetry-breaking constraints, duelling/double

Q-learning and prioritised replay. Their attention-augmented framework achieved several best energies for benchmark sequences and approached optimal solutions for longer chains [2]. These studies show that memory and attention mechanisms help agents model long-range dependencies along the chain and navigate the exponential search space. Nevertheless, training remains expensive: RL agents require millions of episodes, and the learned policies may overfit to particular sequence lengths or compositions.

Variational annealing and generative models. Beyond RL, variational annealing treats folding as continuous optimisation. Khandoker et al. trained a dilated recurrent network with a temperature schedule to refine conformations and produced optimal folds for sequences up to 60 beads. Such methods bridge discrete search and differentiable optimisation, hinting at hybrid models that could combine generative proposals with RL fine-tuning. Recent generative models learn distributions over valid HP conformations and can propose plausible folds for downstream refinement. Integrating diffusion-based generative models with energy-guided refinement is a promising direction.

2.3 Quantum algorithms

The HP model's combinatorial nature invites quantum algorithms. By encoding lattice moves into an Ising Hamiltonian, researchers can use quantum annealers to search for low-energy states. A 2022 study demonstrated that a spin-representation allowed an annealer to recover the lowest known energies for chains up to 64 residues and to outperform classical simulated annealing on certain instances. Gate-based quantum computers offer an alternative. Vasavi et al. mapped lattice moves to quantum gates and folded short sequences on noisy intermediate-scale devices, while variational quantum eigen solvers have been combined with evolutionary strategies to fold small peptides. These studies illustrate the potential of quantum hardware for combinatorial optimisation but are currently limited by noise and qubit count. Ongoing research aims to develop more compact encodings, error-corrected qubits and hybrid quantum-classical workflows in which quantum circuits generate candidate conformations and classical algorithms refine them.

2.4 Deep-learning structure prediction

Deep neural networks have transformed structure prediction for real proteins. AlphaFold 2 predicts inter-residue distances and three-dimensional coordinates from the sequence using attention-based transformers. By mid-2021 the AlphaFold database contained predictions for 98.5 % of the human proteome ($\approx 20\,000$ proteins)[10]. Subsequent releases expanded to hundreds of millions of structures across species, though only about 58 % of residues are predicted with high confidence. These models produce near-experimental accuracy for many globular proteins and have revolutionised structural biology, guiding experiment design and drug discovery. Rosetta Fold

uses a three-track architecture to reason simultaneously over sequence, distance matrix and three-dimensional geometry, producing accurate structures in minutes [11]. However, the success of these models reflects both algorithmic advances and the availability of large training datasets: they integrate evolutionary information from multiple sequence alignments and use recycling mechanisms that iteratively refine predictions. Recent studies highlight limitations: deep models struggle with intrinsically disordered proteins, small peptides lacking homologous sequences, and proteins that undergo large conformational changes. A 2024 analysis of fold-switching proteins showed that AlphaFold rarely predicts alternative conformations and is a poor discriminator between low and high energy states [12]. Hybrid methods that integrate physics-based energy functions or sample ensembles may therefore be necessary to capture dynamic behaviour and multi-state systems.

2.5 Benchmarks, data and evaluation

Benchmarking folding algorithms requires standard datasets and metrics. The Protein Data Bank houses experimentally determined structures; these serve as training and test sets for deep models. The Critical Assessment of Protein Structure Prediction (CASP) competition evaluates methods on blind targets using metrics like the global distance test (GDT) and TM-score. GDT measures the fraction of residues within specified distance thresholds, while TM-score normalises structural deviation by chain length. For the HP model, researchers benchmark algorithms on sequences with known optimal energies, measuring hydrophobic contacts or the energy gap relative to the optimum. Reinforcement-learning and quantum studies compare their results against heuristics such as beam search or nested Monte-Carlo. Although these metrics provide useful standards, success on the HP model does not guarantee performance on real proteins, and evaluation must consider side-chain geometry, steric clashes and flexibility. Quality assessment tools such as MolProbity evaluate side-chain rotamers and steric clashes, whereas flexibility metrics estimate residue mobility. Carefully curated benchmark sets and transparent reporting of metrics are essential for fair comparison.

3 Challenges

Despite rapid progress, several challenges persist across computational folding methods. Data bias and generalisation remain major obstacles: the Protein Data Bank contains mostly well-behaved, crystallisable proteins, while under-represented classes such as membrane proteins and intrinsically disordered regions lead to biased training and poor generalisation. Evidence suggests that some deep models memorise structures from the training set rather than infer physical principles; this limits their ability to predict novel folds or dynamic conformations. Energy landscape complexity is another hurdle. Real proteins fold in continuous space with

thousands of degrees of freedom. Lattice models oversimplify by ignoring side chains and solvent effects, yet full molecular simulations remain computationally prohibitive. Dynamic and multi-state folding adds further complexity: many proteins populate multiple conformations depending on environment or ligand binding, so predicting a single static structure fails to capture this ensemble behaviour. Integration of physics and learning poses a methodological challenge. Purely data-driven models may output structures with unrealistic bond lengths or steric clashes; incorporating physics-based priors, differentiable force fields or energy-regularised losses is needed to enforce structural constraints. Finally, computational cost remains high. Training deep models like AlphaFold requires significant compute resources; RL agents must be retrained for each sequence; quantum hardware is still limited in scale. Addressing these challenges will require innovations in algorithms, data collection and hardware.

4 Prospects

Future progress will likely arise from hybrid approaches that combine the strengths of different methods. Reinforcement-learning agents could be enhanced with generative models that propose plausible fragments, reducing exploration complexity. Meta-learning and transfer learning may enable policies that generalise across diverse sequences. Integrating physical knowledge into neural networks—for example through differentiable force fields or energy-regularised losses—could improve realism and generalisation. Quantum computing may eventually fold larger proteins as hardware scales, and quantum-inspired classical algorithms could provide improved heuristics. Expanding structural databases to include membrane proteins, intrinsically disordered regions and dynamic ensembles will reduce bias. High-throughput experimental techniques such as cryo-electron tomography and NMR provide time-resolved information that could train models to predict ensembles rather than single structures. Hardware advances, such as low-precision accelerators, neuromorphic chips and analog computing, may reduce the cost of training and inference. Finally, careful consideration of ethical issues and transparent reporting of model confidence will be essential as predictive models enable de novo design and drug discovery.

5 Conclusions

Protein folding remains a rich and challenging problem. Misfolding contributes to diseases, yet computing how a sequence folds requires navigating an astronomically large energy landscape. Simplified lattice models provide tractable testbeds, and heuristics have been supplemented by reinforcement learning, variational annealing and quantum algorithms. Deep learning has revolutionised prediction for many real proteins, but its reliance on training data and limitations on dynamic or disordered systems highlight the need for hybrid

methods. By integrating physical principles, machine-learning flexibility, improved datasets and advances in classical and quantum hardware, researchers can deepen our understanding of protein folding and unlock applications in biology and medicine.

References

1. G. Espitia, Y. T. Pang and J. C. Gumbart. Protein Structure Prediction in the 3D Hydrophobic-Polar Model Using Deep Reinforcement Learning. arXiv:2412.20329 [cs.LG], (2024)
2. P. Z. Liu and H. Iba. Enhancing Reinforcement Learning in 3-Dimensional Hydrophobic-Polar Protein Folding Model with Attention-Based Layers. arXiv:2504.15634 [cs.LG], (2025)
3. A. Khandoker, X. Bresson and P. Franek. Variational Annealing for Protein Folding in the Hydrophobic-Polar Model. Proceedings of the AAAI Conference on Artificial Intelligence, (2025)
4. B. Boumedine and S. Bouroubi. Hybrid Cuckoo Search–Hill Climbing Algorithm for the Three-Dimensional Hydrophobic-Polar Protein Folding Model. Bioinformatics, (2021)
5. A. Roucairol and T. Cazenave. Nested Monte-Carlo Search for the Hydrophobic-Polar Model. Proceedings of the IEEE Symposium on Computational Intelligence, (2023)
6. V. Lakkaraju, T. S. Humble and W. J. Huggins. A Turn-Based Encoding Algorithm for Folding Short Lattice Proteins on Quantum Computers. arXiv preprint, (2024)
7. E. Irbäck, J. Larsson and A. Andrist. Solving Lattice Protein Folding with Quantum Annealing Without Auxiliary Qubits. Quantum Information Processing, (2022)
8. A. Robert, D. Zhou and J. Whitfield. Resource-Efficient Variational Quantum Eigensolver for Small Peptides. Nature Communications, (2021)
9. J. Takeshita and S. Kobayashi. FoldingZero: Reinforcement Learning for 2D Lattice Proteins. arXiv:1812.00967 [cs.LG], (2018)
10. J. Jumper, R. Evans, A. Pritzel, T. Green and colleagues. Highly Accurate Protein Structure Prediction with AlphaFold. Nature, (2021)
11. M. Baek, F. D. DiMaio, I. Anishchenko and colleagues. Accurate Prediction of Protein Structures and Interactions Using a Three-Track Neural Network. Science, (2021)
12. C. Y. Wang and J. B. Xu. Limitations of AlphaFold on Fold-Switching Proteins. Nature Communications, (2024)