Protein sequence-to-structure learning: Is this the end(-to-end revolution)?

Elodie Laine, Stephan Eismann, Arne Elofsson, Sergei Grudinin

► To cite this version:

Elodie Laine, Stephan Eismann, Arne Elofsson, Sergei Grudinin. Protein sequence-to-structure learning: Is this the end(-to-end revolution)?. *Proteins - Structure, Function and Bioinformatics*, 2021, 89 (12), pp.1770-1786. 10.1002/prot.26235 . hal-03342140

HAL Id: hal-03342140

<https://hal.science/hal-03342140>

Submitted on 13 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Protein sequence-to-structure learning: Is this the end(-to-end revolution)?

Protein sequence-to-structure learning

Elodie Laine^{1*} | Stephan Eismann^{2*} | Arne Elofsson^{3*}
| Sergei Grudinin^{4*}

¹Sorbonne Université, CNRS, IBPS, Laboratoire de Biologie Computationnelle et Quantitative (LCQB), 75005 Paris, France

²Dep. of Computer Science and Applied Physics, Stanford University, Stanford, CA 94305, USA

³Dep. of Biochemistry and Biophysics and Science for Life Laboratory, Stockholm University, Box 1031, 171 72 Solna, Sweden

⁴Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

Correspondence

Sergei Grudinin, Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

Email:

sergei.grudinin@univ-grenoble-alpes.fr

Funding information

EL was funded by the French national research agency grant ANR-17-CE12-0009. SE was supported by a Stanford Bio-X Bowes Fellowship. AE was funded by grants from the Swedish, E-science Research Center, Swedish National Infrastructure for Computing, and Swedish Natural Science Research Council No VR-NT 2016-03798.

The potential of deep learning has been recognized in the protein structure prediction community for some time, and became indisputable after CASP13. In CASP14, deep learning has boosted the field to unanticipated levels reaching near-experimental accuracy. This success comes from advances transferred from other machine learning areas, as well as methods specifically designed to deal with protein sequences and structures, and their abstractions. Novel emerging approaches include (i) geometric learning, *i.e.* learning on representations such as graphs, 3D Voronoi tessellations, and point clouds; (ii) pre-trained protein language models leveraging attention; (iii) equivariant architectures preserving the symmetry of 3D space; (iv) use of large meta-genome databases; (v) combinations of protein representations; (vi) and finally truly end-to-end architectures, *i.e.* differentiable models starting from a sequence and returning a 3D structure. Here, we provide an overview and our opinion of the novel deep learning approaches developed in the last two years and widely used in CASP14.

KEYWORDS

deep learning, protein structure prediction, CASP14, geometric learning, equivariance, end-to-end architectures, protein language models

1 | INTRODUCTION

In December 2020, the fourteenth edition of CASP marked a big leap in protein three-dimensional (3D) structure prediction. Indeed, deep learning-powered ap-

proaches have reached unprecedented levels of near-experimental accuracy. This achievement has been made possible thanks to the latest improvements in geometric learning and natural language processing (NLP) techniques, and to the amounts of sequence and struc-

* Equally contributing authors.

ture data accessible today. The fundamental basis for the revolution in structure prediction comes from the use of co-evolution. While traditional measures of co-variations in natural sequences led to a few successes [1, 2, 3], major improvements came from recasting the problem as an inverse Potts model [4, 5]. These ideas started to show their full potential about 10 years ago with the development of efficient methods dealing with large scale multiple sequence alignments [6, 7, 8]. They enabled the modelling of 3D structures for large protein families [9, 10, 11, 12, 13, 14].

Shifting from unsupervised statistical inference to supervised deep learning further boosted the accuracy of the predicted contacts, and extended the applicability of this conceptual framework to families with fewer sequences [15, 16] and to the prediction of residue-residue distances [17, 18]. These advances have significantly increased the protein structure modelling coverage of genomes [19, 20, 21], and also of bacterial interactomes [22, 23, 24]. Over the past years, the CASP community has contributed to these efforts, with an increasing number of teams developing and applying deep learning approaches.

The emergence of novel deep learning techniques has inspired a re-visit of the representations best suited for biological objects (protein sequences and structures). In particular, advances in the treatment of language [25] and of 3D geometry [26, 27, 28, 29, 30] by deep learning architectures have further benefited the field of protein structure and function prediction. Expanding on this progress, the DeepMind team demonstrated in CASP14 that it is possible to produce extremely accurate 3D models of proteins by learning end-to-end from sequence alignments of related proteins [31]. This implies being able to capture long-range dependencies between amino acid residues, to transform these dependencies into structural constraints, and to preserve the symmetry and properties of the 3D space when operating on protein structures.

This article is a follow-up to Kandathil et al. [32]. It aims at providing CASP participants and observers with some overview of the recent developments in deep learning applied to protein structure prediction, and some comprehensive description of key concepts we think have contributed to the formidable improvements we have witnessed in the latest CASP edition. We then discuss the implications of these improvements, the next-to-solve problems, and speculate about the future of structural (and computational) biology.

2 | END-TO-END LEARNING FOR PROTEIN STRUCTURE PREDICTION

One of the advantages of deep learning methods compared with traditional machine learning approaches is

FIGURE 1 Schematic representation of the inputs and outputs of deep learning-based methods in CASP14, excluding pipelines compiling several methods coming from different sources, and methods lacking a clear description. The blue and red lines indicate the input and output levels, respectively. Pretrained: sequence embeddings determined from NLP models pre-trained on huge amounts of sequence data. MSA: raw multiple sequence alignment. MSA-feat: MSA features (such as PSSMs, covariance and precision matrices). Contacts: contact or distance matrix. Geometry: geometrical features, typically including contacts/distances and torsion angles. Structure: 3D coordinates. QA: model quality. In case of several inputs and/or outputs, we report those closest to the "end". BrainFold is highlighted with a star as it takes only the query sequence as input, without using pre-trained embeddings. This classification is based on available information from CASP abstracts and publications/preprints. See Supplementary Table S1 for more details.

the ability to automatically extract features from the input data without the need to carefully handcraft them (and potentially miss salient information). Assuming sufficient training data is available, learned features are expected to better generalize to heterogeneous or novel datasets. In addition, it is generally accepted that *end-to-end learning*, where the network is trained to produce the exact desired output and not some sort of heuristic representation of it, is advantageous. Indeed, achieving a high accuracy on some intermediate result does not guarantee high accuracy on the final output. For instance, a learning algorithm may achieve a small loss on dihedral angles, and yet computing atomic coordinates from the predicted dihedral angles could lead to a high reconstruction error [33]. Nevertheless, introducing well-chosen intermediate losses in a so-called "end-to-end" architecture can help to produce better final outputs [31]. These auxiliary intermediate losses provide some guarantee that the method is not only able to produce an accurate final output (e.g. a protein 3D structure) but also to accurately model some other prop-

erties of the object under study (e.g. secondary structure, stereo-chemical quality...), and a mean to incorporate additional domain knowledge. While most protein structure prediction methods take pre-computed features as input and output a contact or distance map, possibly augmented with other geometrical features (Fig. 1, see iPhord, ProSPr [34], Kiharalab_Contact [35], Phamulator, DeepPotential, RaptorX [36], Galaxy, Triple-tRes [37], A2I2Prot, DESTINI2 [38], DeepHelicon [39], DeepHomo [40], ICOS, PrayogRealDistance [41, 42], RBO-PSP-CP [43], DeepECA, ropius0 [44], tFOLD, plus QUARK, Risoluto, Multicom [45] and those from the Zhang lab), several efforts have been recently engaged towards developing end-to-end architectures. Here, we will shortly review these efforts and try to identify the key components of what represents end-to-end learning in protein structure prediction (Table 1).

Ideally, the ultimate input would be the sequence of the query protein. So far, only a couple of learning methods have exploited solely and directly this information to efficiently fold proteins *de novo* [46, 47]. They rely on differentiable [46] and neural [47] potentials whose parameters are learnt from conformational ensembles generated by Langevin dynamics simulations. More commonly, the strategy of state-of-the-art methods is to leverage the very high degenerative nature of the sequence-structure relationship through the use of a multiple sequence alignment (MSA) of evolutionary-related sequences, or a pre-trained protein language model (see below). In this context, methods qualifying for "end-to-X" learning should take as input raw (possibly aligned) sequence(s), as opposed to features derived from them such as conservation levels (e.g. stored in a Position-Specific Scoring Matrix or PSSM) or co-evolution estimates (e.g. mutual information, direct pairwise couplings). One of the first examples of end-to-X method was rawMSA [48], which leveraged embedding techniques from the field of NLP, to map the amino acid residues into a continuous space adaptively learned based on the sequence context (Table 1). In DMPfold2 [49, 50], this idea was extended to MSAs of arbitrary lengths by scanning individual columns in the MSA with stacked Gated Recurrent Unit (GRU) layers. CopulaNet [51] adopts a query-centered view by expanding the input MSA to a set of query-homolog pairwise alignments prior to embedding it. In AlphaFold2 [31], the MSA embedding is obtained through several rounds of self-attention (see below) applied to the MSA rows and columns. Beyond computing MSA embeddings, rawMSA, CopulaNet and AlphaFold2 add an explicit step aimed at converting the information they contain into residue-residue pairwise couplings through an *outer product* operation on the embedding vectors. Recently, a compromise end-to-X solution where the computation of traditional hand-crafted features takes place on the GPU and is tightly coupled to the network was implemented into trRosetta [52], allowing for backpropa-

gating gradients all the way to the input sequences [53].

At the other end of the spectrum, the ultimate output is the 3D structure of the query protein. Thus, an "X-to-end" deep learning architecture should directly produce 3D coordinates and not some intermediate representation such as a contact map. M. AlQuraishi [54] was among the first to develop such a method in 2019 (Table 1). The model takes as input a PSSM, without accounting for any co-evolutionary information, and outputs the Cartesian coordinates of the protein. The torsion angles are predicted and used to reconstruct the 3D structure. Although novel, such an approach has so far not proven to perform better than earlier methods in CASP. One well-known problem is that internal coordinates are extremely sensitive to small deviations as the latter easily propagate through the protein, generating large errors in the reconstructed structure [33]. To overcome this problem, it is possible to efficiently reconstruct Cartesian coordinates from a distance matrix by using multi-dimensional scaling (MDS) or other optimization techniques as in CUTSP [55], DMPfold2 [49], or E2E and FALCON-geom methods of CASP14. In its classical formulation, used by both DMPfold2 and E2E, MDS extracts exact 3D coordinates (provided that the distance matrix is exact) through eigendecomposition of the centered distance matrix. Nevertheless, one issue with using MDS as the final layer in the network is that the output may be a mirror image (chiral version) of the protein. The most recent version of DMPfold2 (DMPfold2-new in Table 1 [50]) attempted to resolve this issue by adding an extra-GRU layer. AlphaFold2 takes a different route and elegantly solves the 3D reconstruction and the mirror-image problems jointly by learning spatial transformations of the local reference frames of each of the protein residues. Computing the geometric loss function in the local frames automatically distinguishes the mirror images, as one of the local axes is a vector product of the two others. Noticeably, even though X-to-end approaches generate a 3D structure, the latter is usually *refined* afterwards (for example through molecular dynamics simulations). For instance, relaxation of AlphaFold2's output with a physical force field is necessary to enforce peptide bond geometry [31].

Although the protein 3D structure appears as an obvious and legitimate target, one may wonder whether generating 3D coordinates confers any advantage, in terms of problem solving and performance, compared to a *perfect* 2D contact map. First, as mentioned above, efficient methods to use 2D information for generating 3D models exist [56, 52]. Further, the most popular residue- or even atom-level loss functions used in deep neural networks (DNNs) do not depend on the superposition of the predicted model to the ground-truth structure and are evaluated using the comparison of distance maps. The most illustrative example is the local distance difference test (lDDT) [57], which has been employed as a target function in CASP14 by some of the best perform-

ers including AlphaFold2 [31] and Rosetta. The value of this loss would not change if we swap the 3D and 2D representations. Nevertheless, it is not clear whether a perfect 2D map can be reached without using some 3D knowledge about the structure. Operating on 3D representations allows calculating global or local quality scores reflective of the structural accuracy in a way that 2D distance maps do not, as illustrated by the mirror-image issue mentioned above. The DNN can then learn to regress against these quality scores, and iteratively refine a first rough 3D guess by predicting (local) deformations to arrive at a better structure. However, operating in 3D poses specific challenges related to the preservation of symmetries, which we discuss in Section 5. So far, the only successful example of indisputable improvement of 3D structure representation over 2D maps is given by AlphaFold2 [31]. Whether similar performance can be achieved with 2D maps and whether 2D maps are needed at all in the predictive process remain open questions.

Being able to produce 3D models resembling experimental structures implies being able to tell apart "good" from "bad" models. Hence, protein model quality assessment (MQA or QA), now referred in CASP to as estimation of model accuracy (EMA), has always been an important step in protein structure prediction pipelines. It allows, in principle, to choose the best models (in case of *global* QA) and/or spot inaccuracies in the proposed models for a subsequent refinement (in case of *local* QA). In recent years, a large number of deep learning-based approaches have been specifically designed for this task. Classically, they take a 3D model as input and then assess its quality in a stand-alone fashion (Fig. 1). Alternatively, some teams proposed integrative approaches. For example, QDeep QA predictions [58] are based on distance estimations from DMPfold [21]. In GalaxyRefine2 [59], Refined [60], and Baker suite [61], the QA is incorporated into a model refinement pipeline. Finally, QA blocks may be used as an integral part of a sequence-to-structure prediction process, as is the case in DMPfold2 [49] and AlphaFold2 [31].

3 | THE IMPORTANCE OF DATA AND DATA REPRESENTATIONS

The success of deep-learning methods is heavily grounded in the availability of large amounts of data, and the development of suitable representations structuring and expressing the information they contain. The advent of high throughput sequencing technologies has widened the gap between the number of known protein sequences and known protein structures. Genomics has become pre-eminent in terms of data scale, with an exponential growth [64, 65]. These huge amounts of data offer unprecedented opportunities to develop high-capacity models detecting co-variation patterns

and learning the "protein language".

3.1 | Leveraging (meta-)genomics

In the last few years, the accessible resources for unannotated sequences coming from *metagenomics* experiments have multiplied. They include databases like NCBI GenBank [66], Metaclust [67], BFD [68], MetaEuk [69], EBI MGnify [70], and IMG/M [71]. Since CASP12, several teams attempted to exploit this type of data, mostly to increase the depth of the MSAs and obtain a more accurate estimation of (co-)evolutionary features. For example, RaptorX [36], methods from the Yang and Baker teams [72, 73], Multicom [45], and GALAXY exploited metagenome data for contact prediction and distance estimation between residue pairs in combination with residual convolutional neural networks (resCNNs). The HMS-Casper [54, 74], DMPfold2 [49] and AlphaFold2 methods [31] exploited them directly to predict 3D structures. Regarding QA, DeepPotential from the Zhang lab and QDeep [58] leverage generated MSA profiles from metagenome databases. To gather large amounts of sequences, coming from different sources, many teams relied on the DeepMSA algorithm [75]. Most of the time, the sequences were integrated altogether in a single MSA. However, some methods proposed to combine several MSAs with different weights (e.g. Kihara's lab) or to select a few of them with high depth and/or variability (e.g. DeepPotential). Noteworthy, deep learning is not only used to exploit sequence alignments, but also to generate them. For instance, the SAdLSA algorithm improves the quality of low-sequence identity alignments by learning the "protein folding code" from structural alignments [76]. NDThreader [77] and ProALIGN [78] are specifically designed to optimally align the query with the template in template-based modeling. Both methods exploit predicted or observed inter-residue distances to improve the sequence alignments, a strategy that proved powerful already in CASP13 [72, 79, 80].

3.2 | From MSA to query-specific embeddings

The most traditional way to extract information from an MSA is to compute a probabilistic profile or a PSSM reflecting the abundance of each amino acid at each position. This type of representation has been very popular from the very first CASP editions. Over the past 10 years, direct coupling analysis (DCA)-based models [12], including Potts model and pseudolikelihood maximization [8, 81, 82, 83], and Graphical lasso-based (low-rank) models [84, 85, 86] became widespread in the community. These statistical methods explicitly estimate residue pairwise couplings as proxies for 3D contacts. More recently, some meta-models [87, 88], correlation and precision matrix-based approaches [89, 90, 52], and

TABLE 1 Overview of X-to-end and end-to-X deep learning approaches for protein structure prediction.

End-to-end learning	
AlphaFold2[31]	The MSA, along with templates, is fed into a translation and rotation equivariant transformer architecture, which outputs a 3D structural model
DMPfold2 (new)[49, 50]	The MSA, along with the precision matrix, is fed into a GRU, which outputs a 3D structure
End-to-X learning	
MSA Transformer[62]	Transformer architecture
rawMSA[48]	The MSA is fed into a 2DCNN (the first convolutional layer creates an embedding) which outputs a contact map
CopulaNet[51]	Extracts all sequence pairs from the MSA and feeds them to a dilated resCNN
TOWER	The network is trained with a deep dilated resCNN to predict inter-residue distances directly from the raw MSA
trRosetta[52]	Computes traditional MSA features on the fly and passes them to dilated convolutional layers
X-to-end learning	
NOVA[63]	Adopts DeepFragLib from the same team which uses Long Short Term Memory units (LSTMs), to output a 3D structure
DMPfold2[49]	The MSA, along with the precision matrix, is fed into a GRU, which outputs distances and angles (version used in CASP14)
HMS-Casper[54]	Raw sequences plus PSSMs are given to a "Recurrent Geometrical Network" comprising LSTM and geometric units and outputting a 3D structure

a variety of of deep-learning models [91, 16, 92, 93, 21, 38, 73, 42, 41, 37, 45], including generative adversarial networks for contact map generation and refinement [94, 35], got widely used to capture the same type of co-evolutionary information. One limitation of these methods is that they estimate average properties over an ensemble of sequences representative of a protein family. Hence, they may miss information specifically relevant to the protein query. The DeepMind team circumvented this limitation with AlphaFold2 by computing embeddings for residue-residue relationships within the query and sequence-residue relationships between the sequences in the MSA and the query, and making the information flow between these two representations. Alternatively, one may *transfer* the knowledge acquired on hundreds of millions of natural sequences to generate query-specific embeddings (Table 2). Several models developed for NLP, including BERT [95], ELMo [96], and GPT-2 [97], have been adapted to the "protein language". During the semi-supervised training phase, the model attempts to predict a masked or the next token [98]. In CASP14, EMBER directly made use of ELMo and BERT while HMS-Casper [54] used a reformulated version of the latter, called AminoBert. A2I2Prot and CUTSP leveraged the TAPE initiative [98], which provides data, tasks and benchmarks to facilitate the evaluation of protein transfer learning.

3.3 | Representations of protein structure

Sequence-based protein representation may be enriched with different levels of *structural information*, for

example, some prior knowledge about secondary structure (SS) elements. In principle, some of these elements, such as alpha helices or beta strands, can be represented with 3D primitives. An interesting idea that we saw in CASP14 was the use of a discrete version of Frenet-Serret frames for the protein backbone parametrization by HMS-Casper. However, such a representation is very complex, and a much simpler way would be to abstract SS primitives with a hydrogen-bond (HB) 2D map. For example, the ISSEC network was specifically trained to segment SS elements in 2D contact maps [99]. Similarly, the protein 3D topology may be abstracted as a 2D contact map, or its probabilistic generalization, e.g. a matrix filled with continuous probabilities or contact propensities between protein atoms or residues. Beyond 2D contact maps, richer descriptions of the 3D structures can be achieved with 2D contact manifolds and protein surfaces, 3D molecular graphs, point clouds, sets of oriented local frames, volumetric 3D maps, or 3D tessellations, e.g. through Voronoi diagrams (Table 2). These different levels of protein representations and their applications in CASP are discussed in more details below and schematically shown in Fig. 2.

3.3.1 | Volumetric protein representations

The first attempt to train 3D CNNs on a volumetric protein representation dates back to CASP12, with the goal of assessing protein model quality [100]. The architecture was robust but had two major limitations. Specifically, it relied on a predefined protein's atom types, and the orientation of the protein model given as input had

FIGURE 2 Comparison between protein representations for human PCNA (PDB code:1AXC, chain A) .

TABLE 2 Overview of approaches transferring knowledge from large amounts of protein sequence data.

HMS-Casper (NLP)[54]	Sequence embeddings generated by a reformulated version of the BERT language model are given as input to a LSTM-based architecture
EMBER (NLP)	Sequence embeddings are generated by BERT and ELMo trained on protein sequence sets and given to a resCNN with dilations
A2I2Prot	A sequence embedding correlation map is fed into a resCNN
CUTSP[55]	Sequence embeddings, along with a MSA, are fed to a bi-directional GRU and LSTM with skip connections, followed by an Encoder-Decoder architecture

an influence on the output of the network. In other words, the network was not *rotation-invariant*. To cope with this issue, it had to be trained on the input data augmented by a set of rotations applied to each input protein model. In a follow-up work, Derevyanko and Lamoureaux [101] introduced an SE(3)-invariant architecture building on Weiler et al. [102].

The Ornate architecture overcomes both limitations [103] (Table 3). Ornate learns atom type embeddings and constructs *local* volumetric representations of each amino acid in a protein in a local coordinate system, thus achieving *local translation-rotation invariance* of the network. Sato-3DCNN by Sato and Ishida [104] used an idea similar to that of Ornate with oriented local frames but did not automatically learn the atom type embeddings. 3DCNN_prof (or P3CMQA) extended this network with additional input features including MSA pro-

file, predicted secondary structure, and solvent accessibility [105]. Finally, DeepAccNet showed remarkable performance in the CASP14 refinement category. This architecture extends Ornate by adding 1D and 2D input features coming from sequence and Rosetta energy terms [61] (Table 3). It predicts per-residue model accuracy and also inter-residue distance signed error, such that the network can be efficiently used for protein model refinement.

3.3.2 | Graph protein representations

A remarkable fact of the CASP14 edition is the emergence of graph representations as means to encode sequence and/or structure information. Indeed, graphs allow formally and compactly encoding diverse relationships between heterogeneous objects. The DeepMind

team was probably the one who best exploited this property, by using the graph representation to encode both sequence information taken as input and structural information learned by the architecture in an end-to-end fashion. Several other teams have made contributions toward deriving graph representations for protein data and developing algorithms operating on these graphs. For example, DeepML, GQArank, LAW, and GraphQA [106] applied classical graph convolutional networks (GCNs) at the residue level, where the convolution operator averages the features of each node's neighbours (Table 3). Spherical Graph Convolutional Network (S-GCN) made a step further and extended the graph convolution operator for spherical geometry in molecular graphs. This allowed to effectively encode mutual angular dependence of neighbouring graph nodes using spherical harmonics expansions [107]. In its turn, GN-Refine predicted distances between protein atoms using graph neural networks and then converted these distances into interatomic potentials and employed them for protein structure refinement [108]. A more recent method, GVP-GNN [109], augments graph networks with the ability to reason about protein features expressed as geometric vectors in an equivariant manner.

3.3.3 | Point clouds and oriented frames

Alternatively to a graph, one may want to make the protein topology evolve through the architecture, without explicitly fixing it. The protein is then seen as a set of isolated nodes with specific positions in 3D space, in other words, a *3D point cloud*. The EDN method in CASP14 was the first approach to describe a 3D protein structure as a set of points in 3D [110, 111]. In this setting, each point is associated with a set of features. At input, individual atoms are represented as points and the chemical element type is the sole associated feature. New point-based features are then calculated over a series of rotation-equivariant convolutions based on the 3D environment around each point. In addition, the network aggregates information at different levels of point hierarchy, from individual atoms over α -carbons to the whole protein.

In AlphaFold2, the protein structure is modelled through a related representation, a cloud of local coordinate system frames, each unambiguously associated with a residue. The method updates these frames indirectly by applying an attention mechanism to "3D points" generated from the query sequence embedding. Hence, at each step, the changes computed on the embedding, which is an abstract representation of the protein, are converted into concrete 3D coordinate changes. DeepMind team uses the term "residue gas" to refer to this frame cloud representation, to emphasize the fact that the peptide bond geometry is totally unconstrained (i.e. the peptide bonds are "broken").

3.3.4 | 3D tessellations

A somewhat similar concept to the graph description is a tessellation of the 2D or 3D space. Tessellations are partitions of the space (3D for most CASP applications) into regions (cells) with specific properties. A tessellation can be represented with a graph, where each node stands for a cell and each edge for the contact between two cells. A particularly useful type of tessellation is the *Voronoi tessellation*, or *Voronoi diagram*. Considering protein structure, the interior of the Voronoi cell around each protein atom must be closer to that atom than to any other. As the atoms are physical objects with different radii, the Voronoi cells are defined by intersecting pairwise bisector surfaces. In case of the additively-weighted Voronoi tessellation, a bisector surface is a part of a hyperboloid of two sheets, approaching a plane when the difference between atomic radii tend to zero (Fig. 2). The Voronoi tessellation turned out to be a very powerful description of protein structure and interactions and has been used in structural bioinformatics for several decades [112, 113, 114, 115, 116]. In CASP14, we saw this description incorporated into DNNs. VoroCNN was the first attempt to construct a deep network passing messages between the neighbouring Voronoi cells [117]. The network performs a hierarchical tessellation by starting at the atom level, and then aggregating features to the residue level. Another interesting idea was implemented in VoroMQA-dark [118], an extension of the VoroMQA method [119] where contact Voronoi areas and pseudo-energies are fed to a feed-forward network. An important particularity of VoroMQA, VoroCNN, and related methods is that their Voronoi tessellations are constrained by the solvent-accessible surface. Therefore, the Voronoi cells of the surface atoms are finite, and the corresponding contact surfaces abstract solvent-protein interactions.

3.3.5 | 2D manifolds

Another flavour of the VoroMQA method has also proved successful in scoring protein complexes since the CASP12-CAPRI experiment [120, 121]. Here, only the protein-protein contact areas contribute to the final score. In such an approach, the 3D protein structures are viewed as *2D surface manifolds*. In CASP14, this type of protein description has been combined with deep learning [118]. Other developments include the application of ideas from the recent MoNet manifold network architecture [122] toward learning protein surfaces with very exciting outcomes for protein binding sites and protein-protein complexes prediction [123, 124]. Overall, 2D surface manifolds seem very powerful and compact representations of 3D shapes, at least in the context of protein-protein interactions.

TABLE 3 Overview of deep learning QA approaches in CASP14.

Volumetric representations	
3DCNN[100]	A non-invariant 3D CNN
Ornate[103]	A local frame-based 3D CNN model with learned atom embeddings
Sato-3DCNN[104]	A local frame-based 3D CNN
3DCNN_prof (P3CMQA)[105]	Extends Sato-3DCNN with predicted features and PSSMs
SE(3)-3DCNN[101]	An invariant 3D CNN based on [102] trained for protein complexes
iPhord & DeepMUSICS	3D CNNs
TopQA[125]	3D CNN with explicit rotations and automatic scaling to fit into a unit cube
DeepAccNet[61]	Extends Ornate with 1D and 2D geometrical features to predict per-residue model accuracy and also inter-residue distance signed error
Graph representations	
Graph-QA[106]	GCN with representation learning, explicit modeling of both sequential and 3D structure, geometric invariance, and computational efficiency
S-GCN[107]	Molecular-graph-based method where angular information is accounted for using spherical harmonics
GQArank	GCN with many features, including PSSM and predicted geometrical properties
DeepML	A classical GCN
LAW	5-layer GCN followed by a 3-layer 1D CNN
Tessellations, 2D manifolds, and point clouds	
VoroCNN[117]	A CNN built on a hierarchical 3D Voronoi tessellation of a protein molecule
VoroMQA-dark[118]	A CNN-based extension of VoroMQA [119]
EDN[111]	A point-cloud representation of the atomic structure combined with rotation-equivariant, hierarchical convolutions

4 | FROM CONVOLUTIONS TO ATTENTION

The choice of the protein data representation is intimately linked to that of the deep learning architecture and operators. Historically, the first deep learning breakthrough in protein structure prediction came from CNNs, widely used for computer vision, applying multiple filters to protein "images". Each filter of a standard CNN aggregates information coming from a region of the input data, namely the receptive field (Fig. 3, area above the main diagonal). The filters in the first layer directly operate on the input data, while the filters in each of the subsequent layers apply some operation on the output of the previous layer. As the information is processed by the successive layers, the size of the receptive field increases, and, as a consequence, longer-range dependencies are captured. However, this accounting of long-range dependencies comes at the expense of precision, since it occurs only after a certain depth in the network. Indeed, the late layers corresponding to a large receptive field do not directly operate on the input data but on some abstract representation of it containing less information. This makes CNNs strongly dependent on the way the input observations are ordered or located with respect to each other. For example, when dealing with a 2D covariance matrix computed from a multiple sequence alignment, local patterns formed by residues adjacent in sequence will be captured with higher precision. This may constitute a limitation since protein 3D structures are also stabilized by interactions formed be-

tween distant amino acids in the sequence. When treating a raw MSA as a 2D image, the order of the sequences will also have an influence while this order may somewhat be arbitrary. In the case of a geometric representation of the 3D protein structure, the information encoded in local neighbourhoods of the Euclidian space will be aggregated first and thus more precisely captured than relationships between distant atoms.

One way to overcome such limitation is to introduce gaps (dilations) when defining the filters. With a dilation d , the window starting at location i of size k is $[x_i \ x_{i+d} \ x_{i+2d} \ \dots \ x_{i+(k-1)d}]$. Stacking dilated convolutions with increasingly large d allows operating on exponentially large receptive fields, while retaining short backpropagations [126, 127, 18]. In CASP14, *dilated convolutions* were used by several groups, including ProSPR [34], DESTINI2 [38], CopulaNet [51], PrayogRealDistance [41, 42], and also EMBER, TOWER, ICOS, and LAW/MASS. Another solution lies in the *self-attention mechanism*, where parametric filters capture high-order dependencies between the input observations at arbitrary range and with high precision (Fig. 3, area under the main diagonal). The intuition is to focus on the most relevant parts of the input with respect to a task or output (general attention) or to another part of the input (self-attention). Specifically, for each input point, a set of trainable attention weights determines the relative importance of each of the other input points. Attention mechanisms have made a major breakthrough in NLP, as they allow keeping in memory the sequence context (although limited in practical ap-

FIGURE 3 Comparison between convolutional filters and attention heads. The input data is either represented as a 2D image or as a graph. The information encoded may be for instance MSA-inferred covariances or template-derived Euclidean distances between protein residues. In the top left triangle of the image, the overlapping squares correspond to the increasing receptive fields obtained by stacking multiple layers of convolutional filters in a 2D-CNN. In the area under the main diagonal, the coloured squares represent the input points being the most important with respect to the one in the center of the circle, as the result of applying *attention filters*. In case the nodes represent residues, and the attention weights can be interpreted in terms of 3D distances or contact, only 2 links are necessary to infer a triangle (in blue). This observation was exploited by AlphaFold2 through the implementation of *triangular self-attention* [31]. In the graph on the left, the red arrows indicate a standard convolution aggregating information from neighbouring nodes. On the right, the attention mechanism puts more importance on a subset of neighbours (illustrated by arrow thickness).

plications) in translation tasks [128, 129, 25, 130]. They are particularly well suited for data whose underlying representation does not have a grid-like structure and rather lie in an irregular domain. Such data can often be represented in the form of *graphs*. While standard graph convolutions indifferently aggregate information from neighbouring nodes [131] (Fig. 3, left graph), the attention mechanism puts more importance on a subset of neighbours, without increasing the time complexity [132] (Fig. 3, right graph). In the extreme case, each node may attend to all other nodes, allowing for a full inference of the graph structure. This strategy has been employed by the DeepMind team in CASP14 in several places of their AlphaFold2 architecture. Particularly, when operating on the 3D structure, they impose a strong spatial/locality bias on the attention that depends on the relative positions of 3D points produced in the local frames of the residues (invariant point attention, IPA). Recent works have also shown that the residue-residue dependencies extracted by certain attention heads in transformers trained on large amounts of sequences can be directly interpreted as 3D contacts or distances [133, 134, 62].

5 | FROM INVARIANCE TO EQUIVARIANCE

Breakthrough applications of deep learning often have in common that the underlying methods cater to specific characteristics of the data domain, such as long-range dependencies in text and hierarchical features in images. Figure 4 considers relevant characteristics for macromolecular structure, namely invariance and equivariance with respect to translations and rotations in 3D. For illustration purposes, the figure includes a series of cat cartoons in 2D.

As training progresses, a neural network should learn to identify structural features helpful for the task it is designed to solve. We refer to such an informative feature as a *structural motif* — a specific arrangement of a set of atoms in 3D (Figure 4A). In the case of the cat cartoons, the mouth and nose of the cat shall correspond to “structural motifs”. Given a protein structure, a network should further be able to identify structural motifs independent of the orientation and position in which they occur. If this ability is not built into the network architecture, the network needs to learn it by seeing the same motif in different orientations and positions. This additional learning task does not only require more network parameters, but the network can also only learn an approximation of the desired detection ability.

The general idea is that incorporating specific domain knowledge into the architecture — here the assumption that a structural motif is the same independent of where and in which orientation it occurs — provides an advantage over a more flexible network architecture through a reduction in model parameters that ultimately translates to better network prediction accuracy given the same, finite amount of training data. Assuming the domain-specific assumptions are true, this reduction in model parameters does not result in a loss of expressive power, as it only prevents the network from learning functions inconsistent with the assumptions [30].

Feature detection independent of orientation and position is not enough in the case of a larger receptive field encompassing two structural motifs (Figure 4B). To aggregate information from local neighborhoods, the network also needs information on the relative orientation and position of the learned structural motifs. These geometrical aspects are crucial since they govern the intramolecular interactions. The importance of relative orientation is also apparent in the cat cartoons — rotating the mouth motif by 180° with respect to the nose turns the happy cat into a sad one. Here, the desired property is that of *equivariance*. Informally, a function or neural network layer is equivariant to some transformation (such as a rotation or translation) if a transformation of the input results in the same transformation of the output. For *invariance*, the function output does not vary with respect to transformations of the input. Standard CNNs and other neural network architectures al-

FIGURE 4 Symmetry considerations in learning from macromolecular structure (A) Given the 3D structure of a protein (grey surface), it is desirable that a neural network can identify a structural motif (blue) — a specific arrangement of atoms in 3D — independent of the position and orientation at which the motif occurs in the structure. Insert: Analogy with a cat face, where the ‘structural motif’ is represented by the mouth of the cat. (B) A larger receptive field encompassing two motifs (blue and red) is shown in green. In order to aggregate information from local neighbourhoods, independent motif detection is not enough. The network requires information on the relative orientation and position of the structural motifs, which is realised through translation and rotation equivariant features. Insert: The cat’s happiness changes with the relative orientation of the nose and mouth ‘motifs’. (C) At the global level, the accuracy of a protein model is rotation invariant. Insert: A global rotation leaves the cat’s happiness unaffected.

ready account for translational equivariance (e.g., by encoding only relative positions of atoms) but not for rotational equivariance. Rotational symmetry is specifically important in 3D: non-equivariant architectures require a factor of $O(\delta^{-1})$ more filters to achieve an angular resolution of δ in 2D but already $O(\delta^{-3})$ more filters in 3D [29].

Finally, invariance vs. equivariance can also depend on the perspective, as illustrated in Figure 4C. If the goal is to predict global protein model accuracy, such as measured by global IDDT, a network should provide equivariant outputs at the local level but a global prediction that is invariant under rotations and translations. Turning again to the cartoon cat, we similarly note that a global rotation of the cat leaves its happiness unaffected.

Historically, machine-learning-based scoring functions [135, 136] were inspired by statistical potentials [137, 138] relying on pairwise distance/ angular distributions or contact maps, which are perfectly rotation and translation invariant representations. The challenge of equivariance arose with the development of deep learning architectures operating directly on raw 3D geometry, rather than precomputed (primarily 1D or 2D) features [100]. Pagès et al. [103], Sato and Ishida [104] elegantly circumvented the need for rotational data augmentation in standard CNNs by leveraging a residue-level coordinate system to learn an invariant local quality metric. AlphaFold2 also represents the protein residues with local oriented frames. Both the frames and their relationships are learnt through a geometry-aware attention mechanism. Graph representations of protein structure generally similarly encode the local and global

3D geometry through rotation-invariant scalar features such as angles and distances [106]. Recent efforts include the use of spherical convolutions in combination with a residue-level coordinate system to learn a local quality metric [107], and the development of invariant volumetric [101] and equivariant point clouds representations in 3D [110, 111]. Specifically, in Eismann et al. [111], starting from the 3D coordinates and element type of each atom, the network first learns rotation-equivariant representations of local neighborhoods and then aggregates this information hierarchically to predict a rotation-invariant fingerprint at the level of the entire protein structure, reflecting the previously discussed symmetry considerations. The architecture builds on tensor field networks [29] in which points in 3D space are associated with tensor features (such as scalars and vectors) and these features are updated over consecutive network layers.

From a broader perspective, early work from Cohen and Welling [27] pioneered the use of tools from group representation theory to build a rotation equivariant neural network architecture. This idea has been followed by a rich body of publications, including translation and rotation equivariant architectures for 3D point clouds [29, 139, 140]. These architectures can be seen as equivariant extensions of neural-network-based radial and angular symmetry functions for molecular structure [141, 142, 143]. Weiler et al. [102] further proposed a rotation-equivariant architecture for continuous data in 3D. All these equivariant architectures share the use of spherical harmonics, a set of functions defined on the unit sphere that is intrinsically linked to 3D

rotation equivariance. Spherical harmonics have played a prominent role in molecular surface representations for several decades [144, 145] and are also at the heart of the classical fast multipole method [146].

We believe that equivariant architectures in learning from macromolecular structure will grow further in popularity due to their parameter-efficient expressive power and their ability to directly reason about, and also predict geometric quantities such as vectors. The recent work by Fuchs et al. [147] on small molecules and the success of AlphaFold2 at CASP14 further suggest promise for network architectures combining equivariance and attention mechanisms. This combination was also leveraged in the most recent RoseTTAFold network [148].

6 | CONCLUSION AND OUTLOOK

The success of CASP14 methods in general, and AlphaFold2 in particular, leads to the awareness of the community that highly-accurate protein structure predictions can be obtained for virtually all well-folded protein domains, but also that the performance gap between AlphaFold2 and other methods is significant. After CASP13, it took the community about 12 months to catch up with AlphaFold1. It is uncertain whether we can expect a similar duration, in part due to the combination of innovative approaches in AlphaFold2. We should also not forget that the computational costs engaged by DeepMind to develop the method, tune its hyperparameters and train it are substantial. A rough estimate for the cost of training the network architecture using cloud resources exceeds \$20K. This emphasises the need for a community-wide effort to catch up with AlphaFold2 and/or the design of shortcuts alleviating training costs.

Recently, a first attempt to reproduce DeepMind's work was presented. The RoseTTAFold method [148] is not as accurate as AlphaFold2, but still clearly better than all earlier methods. Further, it is rather cheap to train and use, providing hope that computational cost will not be a burden for large scale application of these methods. Nonetheless, it is clear to us that the entire community ought to come together with an open mind to develop next generation deep learning-based tools for protein structure prediction. Such an effort would not only have an impact on the field of structure prediction but also on related fields through the innovation of novel deep learning methods. Further, there are, as discussed below, still many challenges in computational structural biology that are not (yet) solved.

6.1 | The impact of accurate models in structural biology and bioinformatics

Accurate 3D structures provide valuable information about protein biological functions. They can be used

by themselves, and also as a starting point for further computational studies. For some studies, the accuracy of theoretical models has been sufficient, for others not. The improvements brought by CASP14 will, therefore, increase the number of suitable targets for tasks requiring a very high-quality models, such as mutational effects prediction, ligand binding site identification, molecular dynamics simulations, drug discovery and enzymatic reactions modeling, to list a few. They also open avenues for a tight cross-talk between structure *in silico* prediction and experimental determination. Already in CASP14, models from AlphaFold2 were used to phase crystals and thereby to solve protein structures. If this can be extended and done systematically, there are probably hundreds of unsolved protein crystals that could benefit from high-quality models. The latter can also help in the initial steps of single-particle Cryo-EM reconstructions. However, the full extent of the impact of computational models on structural biology and other fields will likely depend on their ability to provide profound novel biological insights, that are generally accepted by the community. When this will happen depends on how good the models are. One possible start here could be to examine what additional information was obtained by the experimental structures, phased by the AlphaFold2 models, over the models themselves. If predicted models are accurate enough, one major role for future structural biology might become to identify all the chemical compounds (proteins, ligands, lipids, co-factors) that interact and then use artificial intelligence methods to predict the structure of this ensemble.

6.2 | Learning the laws of physics?

Most methodologies in computational structural biology build on physics' first principles to describe individual atoms and how they interact. These laws are then used to model larger molecules such as proteins. One may wonder to what extent the emergent data-driven approaches that do not explicitly implement detailed physical descriptions of biomolecules are able to implicitly learn physics laws. For instance, end-to-end sequence-to-structure deep learning methods do not explicitly model water molecules, co-factors or partners. Yet, AlphaFold2 was able to determine the residue side-chain orientations competent for binding a zinc ion in the M23 peptidase (target T1056) and also the bound conformation of a cell wall surface anchor protein forming homo-trimers (target T1080). These conformations make sense, from a physical point of view, only when the co-factor or the partner is present. From a data science point of view, however, if some proteins are always found in complex with co-factors or partners, then the machine will learn to associate the matching sequence contexts with bound conformations. In other words, it implicitly learns the physical contexts compatible (at least in the experimentally data at hand) with a partic-

ular sequence context. This ability may be further exploited to discover ligand- or partner-binding sites by analysing the geometry and physico-chemical properties of the predicted conformations. Indeed, the machine may not only be able to predict plausible bound conformations but also to identify the location of the "missing" co-factor(s), ligand(s) or partner(s). Next, it can be asked – is an accurate physical description necessary for other tasks such as binding free energy estimation and mutant stability prediction? The main limitation might be the amount of training data available to develop such methods, but we certainly do foresee many attempts to transfer these ideas to other areas of computational structural biology.

6.3 | Protein disorder, flexibility and dynamics

Beyond 3D structure, proteins' dynamical behaviour is important for their functions. Flexibility is necessary for binding, enzymatic reactions, transport, and many more [149, 150]. Many proteins adopt two or more stable conformations and the equilibrium between these states has a direct implication on their functioning. For instance, protein kinases, representing about 2% of the human proteome, adopt two distinct forms, one inactive and the other active. These two states are clearly distinguishable and can be captured by X-ray crystallography. As an extreme version of flexibility, intrinsic disorder is commonly observed in eukaryotic proteins and plays crucial roles in transient protein-protein interactions as well as in linkers between domains. Some intrinsically disordered regions (IDRs) form a stable structure upon binding to their partners but it is difficult to experimentally identify them.

Co-evolutionary patterns extracted from related sequences have proven useful to predict some IDR bound forms [151] and, in some cases, to untangle a protein's multiple functional states [152, 153, 154]. However, systematically training DL models to predict protein flexibility, either as a probabilistic structural profile or as conformational *multi-modalities*, remains very challenging. Experimental measurements are very scarce and/or probe conformational states only indirectly. For instance, crystallographic temperature (B-) factors, although abundant, are not reliable proxies of internal molecular fluctuations. Indeed, at cryogenic temperatures, the main contribution to B-factors will be crystal lattice disorder. Another option is to use nuclear magnetic resonance (NMR) data as the ground truth for structure prediction architectures. However, we have relatively few NMR structures (6K that contain at least one protein chain as of 2021), and even fewer collected raw NMR observations. In principle, one can train deep models on NMR-inferred 3D reconstructions, often given as multiple models in the PDB, instead of the raw NMR data, but this most likely does not reflect the

true flexibility of a model, as it is also dependent on the number of nuclear Overhauser effect (NOE) constraints obtained in the NMR experiments. It is also possible to obtain direct measurements of flexibility by studying amide-protein exchange rates by NMR, but this does not provide detailed structural information on different structural states.

Other experimental techniques can also provide information of flexibility. Small-angle X-ray scattering (SAXS) can determine a rough low-resolution shape of the molecules, but it is limited to a few hundred collected datasets. Different structures of (related) proteins solved by X-ray crystallography can shed light into the different conformational states of some proteins. However, only states that form stable crystal forms can be measured, limiting the types of flexibility that can be detected. Moreover, there is an imbalance in the PDB related to the abundance of pharmaceutically important proteins in complex with different ligands, or other factors. For example, the inactive state of kinases is largely underrepresented in the PDB compared to the active state [155]. This may bias data-driven approaches while, in principle, without any extra-information about the context (post-translational modifications, bound ligand...etc), there is no reason why one state should be favoured over the other one. Cryo-EM can also provide information about multiple structural states as well as flexibility. Here, current methods are often limited to a fixed set of clearly distinguishable shapes/conformations present in the sample and selected during refinement. Most of the flexibility information comes in the form of missing density, without any details about the flexible regions beyond the fact that they are flexible. However, we see the community moving toward the reconstruction of continuous structural heterogeneity, also using DL techniques [156, 157, 158]. Similar architectures, i.e. generative adversarial networks and variational auto-encoders, have also been used to generate protein backbones and produce smooth motions through linear interpolations in the latent space [159, 160].

Finally, large collections of molecular dynamics (MD) trajectories [161, 162] may be exploited toward protein flexibility learning. However, today unbiased MD simulations are still too short (and likely too inaccurate) to sample large conformational changes. Therefore, learning from these simulations would be limited to small fluctuations around the starting structure. In fact, as we have seen in the recent CASP structure refinement studies, MD is only practical if additional restraints are applied to keep the structure near the initial conformation. Alternatively, instead of learning from MD trajectories, deep learning can be used to generate conformational ensembles obeying the Boltzmann distribution [163].

6.4 | Protein complexes and interactions

Most proteins do not act alone. They function by interacting with other proteins and molecules. Protein complexes come in different forms and shapes. A complex can consist of one or several types of molecules, contain anything from two to hundreds of different protein chains (as well as other macromolecules), and can have different degrees of symmetries. Experimentally, the study of stable protein interactions can be carried out using various techniques. While many of them only provide an estimate of the strength or probability of the interaction, structure determination methods, including crystallography and Cryo-EM electron microscopy, unveil the atomic-resolution details of the assemblies.

Co-evolutionary information can, in principle, also be used to extract information about protein-protein interactions. Such strategy has been employed to predict bacterial complexes [164, 165, 166], to single out pairs of interacting paralogs [167, 23] and to gain insight into the interactomes of viruses, like HCV [168], or bacteria, like *E. coli* [24]. We do believe that this is the next area where end-to-end learning methods will make an impact. In contrast to the prediction of a single structure, one limitation here might be to detect a strong enough signal, since interactions across protein (and domain) interfaces are less conserved than intra-domain contacts [169]. Noticeably, the different types of assemblies have different specific properties and may require the development of different strategies.

Homodimeric complexes are special as the co-evolutionary signals from a single protein describe both inter- and intra-protein residue-residue interactions. Current methods assume that only predicted interactions not satisfied within a single protein (given some error margin) are potential inter-unit connections although this is not always the case [170]. It is also common that homomeric protein complexes can adopt different quaternary forms, further complicating the prediction, but we would expect that extending AlphaFold2 to predict the structure of homodimers (and even homomultimers) should not be too difficult, at least as long as the multimeric formation is conserved within a family. What might prove to be more difficult is to identify the multimeric state of a protein without some type of experimental information. To the best of our knowledge, this problem is not yet studied.

Heterodimeric complexes create a different challenge for multiple sequence alignments. In short, here it is necessary to match the exact pairs of interacting proteins from two lists of homologs. In rare cases, where there exists exactly one homolog to each of the proteins in a genome, this is trivial. However, many proteins have paralogs that might not all interact with each other. Some paralogs might interact with the same protein and some might not. One common approach is to identify the top hit in each proteome — but this is not always correct and it significantly reduces the number

of sequences in the MSA. In a small benchmark of 215 proteins [171], the structure for only a handful (5-10%) of the complexes could be predicted correctly using a naive approach matching top hits from all genomes [20]. Other methods trying to identify the pairs might work better but are computationally expensive [167]. It is also possible that methods using unaligned sequences will provide a solution to this problem [172]. Assuming that this problem can be solved, we do not see that there should exist any major obstacle to develop an end-to-end solution for the modeling of heterodimeric complexes.

Large molecular machines, such as the ribosome, may represent the most challenging case. They typically perform very important functions in a cell. Their interaction networks may comprise very dense and stable subnetworks, and also parts where binary or ternary interactions are established at a given time-point. Recent Cryo-EM structures of large complexes have revealed that these machines often are dynamic with subunits coming on and off. Clearly, we are still far away from being able to fully predict their structure and dynamics.

In the few weeks since the release of the AlphaFold2 source code, various groups have shown that the program in many cases can produce accurate assemblies, for most types of complexes. Different strategies seem to work for different types of complexes. First, many groups simply added a poly-G linker to join two (or more) chains and successfully docked both protein-protein and protein-peptide complexes. Later, it was discovered that the poly-G linker was not needed, as it was sufficient to just change the residue numbering. Interestingly, in some cases it seems not to be necessary to "merge" the alignments by matching orthologous pairs, instead, the alignments can just be added with gaps at the end - this strategy seems to be the best for homomeric complexes. Exact limits and optimal strategies for using AlphaFold2 for docking will certainly be known in a short time - but we will have to wait for larger benchmarks and not just rely on anecdotal stories.

6.5 | Protein mutations and design

Even one single-point mutation can have a dramatic effect on a protein's ability to fold and/or perform its function(s). In parallel to the evolution of CASP, the past few years have seen a significant improvement in the field of mutational outcome prediction. By leveraging the large amounts of available sequence data, several recent methods have achieved much higher accuracy than established popular approaches relying on a variety of sequence and structure-based features [173, 174, 175, 176, 177, 178, 179, 180, 181]. These approaches make the estimation of the impact of every possible substitution at every position in a protein-coding genome computationally feasible [182]. They also hold great potential for guiding protein design and engineering [183,

184]. The success of these methods lies in their ability to capture dependencies between protein residues either by explicitly estimating inter-residue (pairwise) couplings [178, 179] or by implicitly accounting for global sequence contexts [174, 176]. In essence, the concepts at play are no different from those implemented for protein contact prediction, suggesting that mutational outcome prediction, protein structure prediction and protein design can be unified in a common theoretical framework extracting information from protein sequences [173, 172, 176]. Along this line, recent works have shown that NLP models pre-trained on millions of unlabelled protein sequences can be effectively fine-tuned with small amount of labelled data toward accurately predicting mutational effects as well as 3D contacts [185, 184, 186]. Additionally, fully trained DNNs designed to predict inter-residue distances can be repurposed to estimate the impact of mutations on the 3D structure toward guiding the generation of new sequences predicted to fold to new structures [187].

Acknowledgements

SE was supported by a Stanford Bio-X Bowes Fellowship. AE was funded by grants from the Swedish Natural Science Research Council No. VR-NT 2016-03798, Swedish E-science Research Center and Swedish National Infrastructure for Computing. EL acknowledges the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-17-CE12-0009. The founder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The authors thank Kliment Olechnovič from Vilnius University for his help with illustrating Voronoi cells and proof-reading the manuscript, and Bowen Jing for his feedback on the manuscript.

Conflict of interest

The authors declare no conflict of interest.

references

- [1] Benner SA, Gerloff D. Patterns of divergence in homologous proteins as indicators of secondary and tertiary structure: a prediction of the structure of the catalytic domain of protein kinases. *Adv Enzyme Regul* 1991;31:121–181.
- [2] Gobel U, Sander C, Schneider R, Valencia A. Correlated mutations and residue contacts in proteins. *Proteins* 1994 Apr;18(4):309–317.
- [3] Ortiz AR, Kolinski A, Rotkiewicz P, Ilkowsk B, Skolnick J. Ab initio folding of proteins using restraints derived from evolutionary information. *Proteins* 1999;Suppl 3:177–185.
- [4] Lapedes AS, Giraud BG, Liu L, Stormo GD. Correlated Mutations in Models of Protein Sequences: Phylogenetic and Structural Effects. *Statistics in Molecular Biology and Genetics* 1999;33:236–256.
- [5] Giraud BG, Heumann JM, Lapedes AS. Superadditive correlation. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdisc Topics* 1999 May;59(5 Pt A):4983–4991.
- [6] Thomas J, Ramakrishnan N, Bailey-Kellogg C. Graphical Models of Residue Coupling in Protein Families. In: 5th ACM SIGKDD Workshop on Data Mining in Bioinformatics (BIOKDD); 2005. .
- [7] Weigt M, White RA, Szurmant H, Hoch JA, Hwa T. Identification of direct residue contacts in protein-protein interaction by message passing. *Proc Natl Acad Sci U S A* 2009 Jan;106(1):67–72.
- [8] Balakrishnan S, Kamisetty H, Carbonell JG, Lee SI, Langmead CJ. Learning generative models for protein fold families. *Proteins* 2011 Apr;79(4):1061–1078.
- [9] Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, Sander C, et al. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proc Natl Acad Sci U S A* 2011 Dec;108(49):1293–301.
- [10] Taylor WR, Jones DT, Sadowski MI. Protein topology from predicted residue contacts. *Protein Sci* 2012 Feb;21(2):299–305.
- [11] Sadowski MI, Maksimiak K, Taylor WR. Direct correlation analysis improves fold recognition. *Comput Biol Chem* 2011 Oct;35(5):323–332.
- [12] Marks DS, Colwell LJ, Sheridan R, Hopf TA, Pagnani A, Zecchina R, et al. Protein 3D structure computed from evolutionary sequence variation. *PLoS One* 2011;6(12):e28766.
- [13] Sułkowska JJ, Morcos F, Weigt M, Hwa T, Onuchic JN. Genomics-aided structure prediction. *Proceedings of the National Academy of Sciences* 2012;109(26):10340–10345. <http://www.pnas.org/content/109/26/10340.abstract>.
- [14] Marks DS, Hopf TA, Sander C. Protein structure prediction from sequence variation. *Nature biotechnology* 2012;30(11):1072–1080.
- [15] Skwark MJ, Raimondi D, Michel M, Elofsson A. Improved contact predictions using the recognition of protein like contact patterns. *PLoS Comput Biol* 2014 Nov;10(11):e1003889.
- [16] Wang S, Sun S, Li Z, Zhang R, Xu J. Accurate de novo prediction of protein contact map by ultra-deep learning model. *PLoS computational biology* 2017;13(1):e1005324.
- [17] Xu J. Distance-based protein folding powered by deep learning. *Proc Natl Acad Sci U S A* 2019 Aug;116(34):16856–16865.
- [18] Senior AW, Evans R, Jumper J, Kirkpatrick J, Sifre L, Green T, et al. Improved protein structure prediction using potentials from deep learning. *Nature* 2020 Jan;577(7792):706–710.

- [19] Ovchinnikov S, Kinch L, Park H, Liao Y, Pei J, Kim DE, et al. Large-scale determination of previously unsolved protein structures using evolutionary information. *Elife* 2015;4:e09248.
- [20] Lamb J, Jarmolinska AI, Michel M, Menendez-Hurtado D, Sulkowska JI, Elofsson A. PconsFam: An Interactive Database of Structure Predictions of Pfam Families. *J Mol Biol* 2019 Jun;431(13):2442–2448.
- [21] Greener JG, Kandathil SM, Jones DT. Deep learning extends de novo protein modelling coverage of genomes using iteratively predicted structural constraints. *Nat Commun* 2019 Sep;10(1):3977.
- [22] Hopf TA, Schärfe CPI, Rodrigues JPGLM, Green AG, Kohlbacher O, Sander C, et al. Sequence co-evolution gives 3D contacts and structures of protein complexes. *eLife* 2014 sep;3:e03430. <https://dx.doi.org/10.7554/eLife.03430>.
- [23] Green AG, Elhabashy H, Brock KP, Maddamsetti R, Kohlbacher O, Marks DS. Large-scale discovery of protein interactions at residue resolution using co-evolution calculated from genomic sequences. *Nat Commun* 2021 Mar;12(1):1396.
- [24] Cong Q, Anishchenko I, Ovchinnikov S, Baker D. Protein interaction networks revealed by proteome co-evolution. *Science* 2019 Jul;365(6449):185–189.
- [25] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al., Attention Is All You Need; 2017.
- [26] Bartók AP, Kondor R, Csányi G. On representing chemical environments. *Phys Rev B* 2013 May;87:184115. <https://link.aps.org/doi/10.1103/PhysRevB.87.184115>.
- [27] Cohen TS, Welling M. Steerable cnns. *arXiv preprint arXiv:161208498* 2016;.
- [28] Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural Message Passing for Quantum Chemistry. In: Precup D, Teh YW, editors. *Proceedings of the 34th International Conference on Machine Learning*, vol. 70 of *Proceedings of Machine Learning Research PMLR*; 2017. p. 1263–1272. <http://proceedings.mlr.press/v70/gilmer17a.html>.
- [29] Thomas N, Smidt T, Kearnes S, Yang L, Li L, Kohlhoff K, et al. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:180208219* 2018;.
- [30] Bronstein MM, Bruna J, Cohen T, Veličković P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:210413478* 2021;.
- [31] Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, et al. Highly accurate protein structure prediction with AlphaFold. *Nature* 2021;p. 1–11.
- [32] Kandathil SM, Greener JG, Jones DT. Recent developments in deep learning applied to protein structure prediction. *Proteins: Structure, Function, and Bioinformatics* 2019;87(12):1179–1189.
- [33] Elofsson A, LeGrand S, Eisenberg D. Local moves, an efficient method for protein folding simulations. *Proteins* 1995;23:73–82.
- [34] Billings W, Hedelius B, Millecam T, Wingate D, Corte D, ProSPR: democratized implementation of alphaFold protein distance prediction network. *Biorxiv*, 2019;.
- [35] Subramaniya SRMV, Terashi G, Jain A, Kagaya Y, Kihara D. Protein contact map denoising using generative adversarial networks. *bioRxiv* 2020;.
- [36] Xu J, Mcpartlon M, Li J. Improved protein structure prediction by deep learning irrespective of co-evolution information. *bioRxiv* 2020;.
- [37] Li Y, Zhang C, Bell EW, Zheng W, Zhou X, Yu DJ, et al. Deducing high-accuracy protein contact-maps from a triplet of coevolutionary matrices through deep residual convolutional networks. *PLOS Computational Biology* 2021;17(3):e1008865.
- [38] Gao M, Zhou H, Skolnick J. DESTINI: A deep-learning approach to contact-driven protein structure prediction. *Scientific reports* 2019;9(1):1–13.
- [39] Sun J, Frishman D. DeepHelicon: Accurate prediction of inter-helical residue contacts in transmembrane proteins by residual neural networks. *Journal of Structural Biology* 2020;212(1):107574.
- [40] Yan Y, Huang SY. Accurate prediction of residue-residue contacts across homo-oligomeric protein interfaces through deep leaning. *bioRxiv* 2020;.
- [41] Adhikari B. A fully open-source framework for deep learning protein real-valued distances. *Scientific reports* 2020;10(1):1–10.
- [42] Adhikari B. DEEPCON: protein contact prediction using dilated convolutional neural networks with dropout. *Bioinformatics* 2020;36(2):470–477.
- [43] Stahl K, Schneider M, Brock O. EPSILON-CP: using deep learning to combine information from multiple sources for protein contact prediction. *BMC bioinformatics* 2017;18(1):1–11.
- [44] Margelevičius M. COMER2: GPU-accelerated sensitive and specific homology searches. *Bioinformatics* 2020;36(11):3570–3572.
- [45] Wu T, Guo Z, Hou J, Cheng J. DeepDist: real-value inter-residue distance prediction with deep residual convolutional network. *BMC bioinformatics* 2021;22(1):1–17.
- [46] Jumper JM, Faruk NF, Freed KF, Sosnick TR. Trajectory-based training enables protein simulations with accurate folding and Boltzmann ensembles in cpu-hours. *PLoS computational biology* 2018;14(12):e1006578.
- [47] Ingraham J, Riesselman A, Sander C, Marks D. Learning protein structure with a differentiable simulator. In: *7th International Conference on Learning Representations, ICLR New Orleans, LA, USA, May 6–9, 2019*; 2019. .

- [48] Mirabello C, Wallner B. rawMSA: End-to-end Deep Learning using raw Multiple Sequence Alignments. *PLoS One* 2019;14(8):e0220182.
- [49] Kandathil SM, Greener JG, Lau AM, Jones DT. Deep learning-based prediction of protein structure using learned representations of multiple sequence alignments. *bioRxiv* 2020;.
- [50] Kandathil SM, Greener JG, Lau AM, Jones DT. Ultra-fast end-to-end protein structure prediction enables high-throughput exploration of uncharacterised proteins. *bioRxiv* 2021;p. 2020–11.
- [51] Ju F, Zhu J, Shao B, Kong L, Liu TY, Zheng WM, et al. CopulaNet: Learning residue co-evolution directly from multiple sequence alignment for protein structure prediction. *Nature communications* 2021;12(1):1–9.
- [52] Yang J, Anishchenko I, Park H, Peng Z, Ovchinnikov S, Baker D. Improved protein structure prediction using predicted interresidue orientations. *Proc Natl Acad Sci U S A* 2020 Jan;117(3):1496–1503.
- [53] Norn C, Wicky BI, Juergens D, Liu S, Kim D, Tischer D, et al. Protein sequence design by conformational landscape optimization. *Proceedings of the National Academy of Sciences* 2021;118(11).
- [54] AlQuraishi M. End-to-end differentiable learning of protein structure. *Cell systems* 2019;8(4):292–301.
- [55] Drori I, Thaker D, Srivatsa A, Jeong D, Wang Y, Nan L, et al. Accurate protein structure prediction by embeddings and deep learning representations. *arXiv preprint arXiv:191105531* 2019;.
- [56] Adhikari B, Bhattacharya D, Cao R, Cheng J. CON-FOLD: Residue-residue contact-guided ab initio protein folding. *Proteins* 2015 Aug;83(8):1436–1449.
- [57] Mariani V, Biasini M, Barbato A, Schwede T. IDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics* 2013;29(21):2722–2728.
- [58] Shuvo MH, Bhattacharya S, Bhattacharya D. QDeep: distance-based protein model quality estimation by residue-level ensemble error classifications using stacked deep residual neural networks. *Bioinformatics* 2020;36(Supplement_1):i285–i291.
- [59] Lee GR, Won J, Heo L, Seok C. GalaxyRefine2: simultaneous refinement of inaccurate local regions and overall protein structure. *Nucleic acids research* 2019;47(W1):W451–W455.
- [60] Bhattacharya D. refineD: improved protein structure refinement using machine learning based restrained relaxation. *Bioinformatics* 2019;35(18):3320–3328.
- [61] Hiranuma N, Park H, Baek M, Anishchenko I, Dau-paras J, Baker D. Improved protein structure refinement guided by deep learning based accuracy estimation. *Nature communications* 2021;12(1):1–11.
- [62] Rao R, Liu J, Verkuil R, Meier J, Canny JF, Abbeel P, et al. MSA transformer. *bioRxiv* 2021;.
- [63] Wang T, Qiao Y, Ding W, Mao W, Zhou Y, Gong H. Improved fragment sampling for ab initio protein structure prediction using deep neural networks. *Nature Machine Intelligence* 2019;1(8):347–355.
- [64] Stephens ZD, Lee SY, Faghri F, Campbell RH, Zhai C, Efron MJ, et al. Big data: astronomical or genomic? *PLoS biology* 2015;13(7):e1002195.
- [65] Navarro FC, Mohsen H, Yan C, Li S, Gu M, Meyerson W, et al. Genomics and data science: an application within an umbrella. *Genome biology* 2019;20(1):1–11.
- [66] Benson DA, Cavanaugh M, Clark K, Karsch-Mizrachi I, Ostell J, Pruitt KD, et al. GenBank. *Nucleic acids research* 2018;46(D1):D41–D47.
- [67] Steinegger M, Söding J. Clustering huge protein sequence sets in linear time. *Nature communications* 2018;9(1):1–8.
- [68] Steinegger M, Mirdita M, Söding J. Protein-level assembly increases protein sequence recovery from metagenomic samples manyfold. *Nature methods* 2019;16(7):603–606.
- [69] Levy Karin E, Mirdita M, Söding J. MetaEuk-Sensitive, high-throughput gene discovery, and annotation for large-scale eukaryotic metagenomics. *Microbiome* 2020;8:1–15.
- [70] Mitchell AL, Almeida A, Beracochea M, Boland M, Burgin J, Cochrane G, et al. MGnify: the microbiome analysis resource in 2020. *Nucleic acids research* 2020;48(D1):D570–D578.
- [71] Chen IMA, Chu K, Palaniappan K, Pillay M, Ratner A, Huang J, et al. IMG/M v. 5.0: an integrated data management and comparative analysis system for microbial genomes and microbiomes. *Nucleic acids research* 2019;47(D1):D666–D677.
- [72] Ovchinnikov S, Park H, Varghese N, Huang PS, Pavlopoulos GA, Kim DE, et al. Protein structure determination using metagenome sequence data. *Science* 2017 Jan;355(6322):294–298.
- [73] Wu Q, Peng Z, Anishchenko I, Cong Q, Baker D, Yang J. Protein contact prediction using metagenome sequence data and residual neural networks. *Bioinformatics* 2020;36(1):41–48.
- [74] AlQuraishi M. ProteinNet: a standardized data set for machine learning of protein structure. *BMC bioinformatics* 2019;20(1):1–10.
- [75] Zhang C, Zheng W, Mortuza S, Li Y, Zhang Y. DeepMSA: constructing deep multiple sequence alignment to improve contact prediction and fold-recognition for distant-homology proteins. *Bioinformatics* 2020;36(7):2105–2112.
- [76] Gao M, Skolnick J. A novel sequence alignment algorithm based on deep learning of the protein folding code. *Bioinformatics* 2020;.
- [77] Wu F, Xu J. Deep Template-based Protein Structure Prediction. *bioRxiv* 2020;.

- [78] Kong L, Ju F, Zheng WM, Sun S, Xu J, Bu D. ProALIGN: Directly learning alignments for protein structure prediction via exploiting context-specific alignment motifs. *bioRxiv* 2020;.
- [79] Buchan DW, Jones DT. EigenTHREADER: analogous protein fold recognition by efficient contact map threading. *Bioinformatics* 2017;33(17):2684–2690.
- [80] Zheng W, Wuyun Q, Li Y, Mortuza S, Zhang C, Pearce R, et al. Detecting distant-homology protein structures by aligning deep neural-network based contact maps. *PLoS computational biology* 2019;15(10):e1007411.
- [81] Kamisetty H, Ovchinnikov S, Baker D. Assessing the utility of coevolution-based residue–residue contact predictions in a sequence- and structure-rich era. *Proceedings of the National Academy of Sciences* 2013;110(39):15674–15679. <http://www.pnas.org/content/110/39/15674.abstract>.
- [82] Ekeberg M, Hartonen T, Aurell E. Fast pseudo-likelihood maximization for direct-coupling analysis of protein structure from many homologous amino acid sequences. *Journal of Computational Physics* 2014;276:341–356.
- [83] Seemayer S, Gruber M, Söding J. CCMpred – fast and precise prediction of protein residue-residue contacts from correlated mutations. *Bioinformatics* 2014;30(21):3128–3130.
- [84] Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Bio-statistics* 2008;9(3):432–441.
- [85] Jones DT, Buchan DW, Cozzetto D, Pontil M. PSICOV: precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. *Bioinformatics* 2012 Jan;28(2):184–190.
- [86] Ma J, Wang S, Wang Z, Xu J. Protein contact prediction by integrating joint evolutionary coupling analysis and supervised learning. *Bioinformatics* 2015;31(21):3506–3513.
- [87] Jones DT, Singh T, Kosciółek T, Tetchner S. MetaPSICOV: combining coevolution methods for accurate prediction of contacts and long range hydrogen bonding in proteins. *Bioinformatics* 2015 Apr;31(7):999–1006.
- [88] Skwark MJ, Abdel-Rehim A, Elofsson A. PconsC: Combination of direct information methods and alignments improves contact prediction. *Bioinformatics* 2013 May;29.
- [89] Li Y, Hu J, Zhang C, Yu DJ, Zhang Y. ResPRE: high-accuracy protein contact prediction by coupling precision matrix with deep residual neural networks. *Bioinformatics* 2019;35(22):4647–4655.
- [90] Fukuda H, Tomii K. DeepECA: an end-to-end learning framework for protein contact prediction from a multiple sequence alignment. *BMC bioinformatics* 2020;21(1):1–15.
- [91] Golkov V, Skwark MJ, Golkov A, Dosovitskiy A, Brox T, Meiler J, et al. Protein contact prediction from amino acid co-evolution using convolutional networks for graph-valued images. In: *NIPS*; 2016. p. 4215–4223.
- [92] Jones DT, Kandathil SM. High precision in protein contact prediction using fully convolutional neural networks and minimal sequence features. *Bioinformatics* 2018 Oct;34(19):3308–3315.
- [93] Kandathil SM, Greener JG, Jones DT. Prediction of interresidue contacts with DeepMetaPSICOV in CASP13. *Proteins* 2019 Dec;87(12):1092–1099.
- [94] Yang H, Wang M, Yu Z, Zhao XM, Li A. GANcon: protein contact map prediction with deep generative adversarial network. *IEEE Access* 2020;8:80899–80907.
- [95] Devlin J, Chang MW, Lee K, Toutanova K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* 2018;.
- [96] Peters ME, Neumann M, Iyyer M, Gardner M, Clark C, Lee K, et al. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365* 2018;.
- [97] Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I. Language models are unsupervised multitask learners. *OpenAI blog* 2019;1(8):9.
- [98] Rao R, Bhattacharya N, Thomas N, Duan Y, Chen X, Canny J, et al. Evaluating protein transfer learning with tape. *Advances in Neural Information Processing Systems* 2019;32:9689.
- [99] Zhang Q, Zhu J, Ju F, Kong L, Sun S, Zheng WM, et al. ISSEC: inferring contacts among protein secondary structure elements using deep object detection. *BMC bioinformatics* 2020;21(1):1–13.
- [100] Derevyanko G, Grudin S, Bengio Y, Lamoureux G. Deep convolutional networks for quality assessment of protein folds. *Bioinformatics* 2018 Dec;34(23):4046–4053.
- [101] Derevyanko G, Lamoureux G. Protein-protein docking using learned three-dimensional representations. *bioRxiv* 2019;p. 738690.
- [102] Weiler M, Geiger M, Welling M, Boomsma W, Cohen T. 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data. In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, editors. *Advances in Neural Information Processing Systems 31* Curran Associates, Inc.; 2018.p. 10381–10392.
- [103] Pagès G, Charmettant B, Grudin S. Protein model quality assessment using 3D oriented convolutional neural networks. *Bioinformatics* 2019 Feb;.
- [104] Sato R, Ishida T. Protein model accuracy estimation based on local structure quality assessment using 3D convolutional neural network. *PloS one* 2019;14(9):e0221347.

- [105] Takei Y, Ishida T. P3CMQA: Single-Model Quality Assessment Using 3DCNN with Profile-Based Features. *Bioengineering* 2021;8(3):40.
- [106] Baldassarre F, Hurtado DM, Elofsson A, Azizpour H. GraphQA: Protein Model Quality Assessment using Graph Convolutional Networks. *Bioinformatics* 2020 Aug;.
- [107] Igashov I, Pavlichenko N, Grudinin S. Spherical convolutions on molecular graphs for protein model quality assessment. *arXiv preprint arXiv:201107980* 2020;.
- [108] Jing X, Xu J. Fast and effective protein model refinement by deep graph neural networks. *bioRxiv* 2020;.
- [109] Jing B, Eismann S, Suriana P, Townshend RJL, Dror R. Learning from Protein Structure with Geometric Vector Perceptrons. In: *International Conference on Learning Representations*; 2021. <https://openreview.net/forum?id=1YLJDvSx6J4>.
- [110] Eismann S, Townshend RJ, Thomas N, Jagota M, Jing B, Dror R. Hierarchical, rotation-equivariant neural networks to predict the structure of protein complexes. *arXiv preprint arXiv:200609275* 2020;.
- [111] Eismann S, Suriana P, Jing B, Townshend RJ, Dror RO. Protein model quality assessment using rotation-equivariant, hierarchical neural networks. *arXiv preprint arXiv:201113557* 2020;.
- [112] Pontius J, Richelle J, Wodak SJ. Deviations from standard atomic volumes as a quality measure for protein crystal structures. *Journal of molecular biology* 1996;264(1):121–136.
- [113] Zimmer R, Wöhler M, Thiele R. New scoring schemes for protein fold recognition based on Voronoi contacts. *Bioinformatics (Oxford, England)* 1998;14(3):295–308.
- [114] Poupon A. Voronoi and Voronoi-related tessellations in studies of protein structure and interaction. *Current Opinion in Structural Biology* 2004;14(2):233–241.
- [115] Cazals F, Proust F, Bahadur RP, Janin J. Revisiting the Voronoi description of protein–protein interfaces. *Protein Science* 2006;15(9):2082–2092.
- [116] Olechnovič K, Venclovas Č. Voronota: A fast and reliable tool for computing the vertices of the Voronoi diagram of atomic balls. *Journal of computational chemistry* 2014;35(8):672–681.
- [117] Igashov I, Olechnovic K, Kadukova M, Venclovas C, Grudinin S. VoroCNN: Deep convolutional neural network built on 3D Voronoi tessellation of protein structures. *Bioinformatics* 2021;p. btab118.
- [118] Dapkunas J, Olechnovič K, Venclovas Č. Modeling of protein complexes in CASP14 with emphasis on the interaction interface prediction. *Authorea*, Submitted to *Proteins* 2021;.
- [119] Olechnovic K, Venclovas C. VoroMQA: Assessment of protein structure quality using interatomic contact areas. *Proteins* 2017 Jun;85(6):1131–1145.
- [120] Dapkūnas J, Olechnovič K, Venclovas Č. Structural modeling of protein complexes: Current capabilities and challenges. *Proteins: Structure, Function, and Bioinformatics* 2019;87(12):1222–1232.
- [121] Olechnovič K, Venclovas Č. VoroMQA web server for assessing three-dimensional structures of proteins and protein complexes. *Nucleic acids research* 2019;47(W1):W437–W442.
- [122] Monti F, Boscaini D, Masci J, Rodola E, Svoboda J, Bronstein MM. Geometric deep learning on graphs and manifolds using mixture model cnns. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*; 2017. p. 5115–5124.
- [123] Gainza P, Sverrisson F, Monti F, Rodola E, Boscaini D, Bronstein M, et al. Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nature Methods* 2020;17(2):184–192.
- [124] Sverrisson F, Feydy J, Correia B, Bronstein M. Fast end-to-end learning on protein surfaces. *bioRxiv* 2020;.
- [125] Smith J, Conover M, Stephenson N, Eickholt J, Si D, Sun M, et al. TopQA: a topological representation for single-model protein quality assessment with machine learning. *International Journal of Computational Biology and Drug Design* 2020;13(1):144–153.
- [126] Yu F, Koltun V. Multi-Scale Context Aggregation by Dilated Convolutions; 2016.
- [127] Gupta A, Rush AM. Dilated Convolutions for Modeling Long-Distance Genomic Dependencies; 2017.
- [128] Bahdanau D, Cho K, Bengio Y. Neural Machine Translation by Jointly Learning to Align and Translate; 2016.
- [129] Gehring J, Auli M, Grangier D, Dauphin YN. A Convolutional Encoder Model for Neural Machine Translation; 2017.
- [130] Choromanski K, Likhoshesterov V, Dohan D, Song X, Gane A, Sarlos T, et al. Rethinking attention with performers. *arXiv preprint arXiv:200914794* 2020;.
- [131] Kipf TN, Welling M. Semi-Supervised Classification with Graph Convolutional Networks; 2017.
- [132] Veličković P, Cucurull G, Casanova A, Romero A, Liò P, Bengio Y. Graph Attention Networks. *International Conference on Learning Representations* 2018;<https://openreview.net/forum?id=rJXMpikCZ>.
- [133] Bhattacharya N, Thomas N, Rao R, Daupras J, Koo P, Baker D, et al. Single Layers of Attention Suffice to Predict Protein Contacts. *bioRxiv* 2020;.
- [134] Rao R, Ovchinnikov S, Meier J, Rives A, Sercu T. Transformer protein language models are unsupervised structure learners. *bioRxiv* 2020;.
- [135] Wallner B, Elofsson A. Can correct protein models be identified? *Protein Sci* 2003;12(5):1073–1086.

- [136] Karasikov M, Pagès G, Grudinin S. Smooth orientation-dependent scoring function for coarse-grained protein quality assessment. *Bioinformatics* 2019;35(16):2801–2808.
- [137] Miyazawa S, Jernigan RL. Estimation of effective interresidue contact energies from protein crystal structures: quasi-chemical approximation. *Macromolecules* 1985;18(3):534–552.
- [138] Sippl M. Calculation of conformational ensembles from potentials of mean force. An approach to the knowledge-based prediction of local structures in globular proteins. *J Mol Biol* 1990;213(4):859–883.
- [139] Kondor R. N-body networks: a covariant hierarchical neural network architecture for learning atomic potentials. *arXiv preprint arXiv:180301588* 2018;.
- [140] Anderson B, Hy TS, Kondor R. Cormorant: Covariant molecular neural networks. *arXiv preprint arXiv:190604015* 2019;.
- [141] Behler J, Parrinello M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical Review Letters* 2007;98(14):1–4.
- [142] Schütt KT, Kindermans PJ, Sauceda HE, Chmiela S, Tkatchenko A, Müller KR. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. *arXiv:170608566* 2017;.
- [143] Smith JS, Isayev O, Roitberg AE. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chemical Science* 2017;8(4):3192–3203.
- [144] Max NL, Getzoff ED. Spherical harmonic molecular surfaces. *IEEE Computer Graphics and Applications* 1988;8(4):42–50.
- [145] Ritchie DW, Kemp GJ. Fast computation, rotation, and comparison of low resolution spherical harmonic molecular surfaces. *Journal of Computational Chemistry* 1999;20(4):383–395.
- [146] Greengard L, Rokhlin V. A fast algorithm for particle simulations. *Journal of computational physics* 1987;73(2):325–348.
- [147] Fuchs FB, Worrall DE, Fischer V, Welling M, SE(3)-Transformers: 3D Roto-Translation Equivariant Attention Networks; 2020.
- [148] Baek M, DiMaio F, Anishchenko I, Dauparas J, Ovchinnikov S, Lee GR, et al. Accurate prediction of protein structures and interactions using a 3-track network. *Science* 2021;p. eabj8754.
- [149] Frauenfelder H, Sligar SG, Wolynes PG. The energy landscapes and motions of proteins. *Science* 1991;254(5038):1598–1603.
- [150] Henzler-Wildman K, Kern D. Dynamic personalities of proteins. *Nature* 2007;450(7172):964–972.
- [151] Toth-Petroczy A, Palmedo P, Ingraham J, Hopf TA, Berger B, Sander C, et al. Structured States of Disordered Proteins from Genomic Sequences. *Cell* 2016 Sep;167(1):158–170.e12.
- [152] Morcos F, Onuchic JN. The role of coevolutionary signatures in protein interaction dynamics, complex inference, molecular recognition, and mutational landscapes. *Current opinion in structural biology* 2019;56:179–186.
- [153] Sfriso P, Duran-Frigola M, Mosca R, Emperador A, Aloy P, Orozco M. Residues coevolution guides the systematic identification of alternative functional conformations in proteins. *Structure* 2016;24(1):116–126.
- [154] Morcos F, Jana B, Hwa T, Onuchic JN. Coevolutionary signals across protein lineages help capture multiple protein conformations. *Proceedings of the National Academy of Sciences* 2013;110(51):20533–20538.
- [155] Schwarz D, Merget B, Deane C, Fulle S. Modeling conformational flexibility of kinases in inactive states. *Proteins: Structure, Function, and Bioinformatics* 2019;87(11):943–951.
- [156] Zhong ED, Bepler T, Berger B, Davis JH. Cryo-DRGN: reconstruction of heterogeneous cryo-EM structures using neural networks. *Nature Methods* 2021;18(2):176–185.
- [157] Punjani A, Fleet DJ. 3D Flexible Refinement: Structure and Motion of Flexible Proteins from Cryo-EM. *bioRxiv* 2021;.
- [158] Rosenbaum D, Garnelo M, Zielinski M, Beattie C, Clancy E, Huber A, et al. Inferring a Continuous Distribution of Atom Coordinates from Cryo-EM Images using VAEs. *arXiv preprint arXiv:210614108* 2021;.
- [159] Anand N, Eguchi R, Huang PS. Fully differentiable full-atom protein backbone generation. *ICLR* 2019;.
- [160] Eguchi RR, Anand N, Choe CA, Huang PS. Ig-VAE: generative modeling of immunoglobulin proteins by direct 3D coordinate generation. *bioRxiv* 2020;.
- [161] Meyer T, D'Abramo M, Hospital A, Rueda M, Ferrer-Costa C, Pérez A, et al. MoDEL (Molecular Dynamics Extended Library): a database of atomistic molecular dynamics trajectories. *Structure* 2010;18(11):1399–1409.
- [162] Rodríguez-Espigares I, Torrens-Fontanals M, Tiemann JK, Aranda-García D, Ramírez-Anguita JM, Stepniowski TM, et al. GPCRmd uncovers the dynamics of the 3D-GPCRome. *Nature Methods* 2020;17(8):777–787.
- [163] Noé F, Olsson S, Köhler J, Wu H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science* 2019;365(6457).
- [164] Schug A, Weigt M, Onuchic JN, Hwa T, Szurmant H. High-resolution protein complexes from integrating genomic information with molecular simulation. *Proc Natl Acad Sci U S A* 2009 Dec;106(52):22124–22129.
- [165] Hopf TA, Scharfe CP, Rodrigues JP, Green AG, Kohlbacher O, Sander C, et al. Sequence coevolution gives 3D contacts and structures of protein complexes. *Elife* 2014;3.

- [166] Ovchinnikov S, Kamisetty H, Baker D. Robust and accurate prediction of residue-residue interactions across protein interfaces using evolutionary information. *Elife* 2014;3:e02030.
- [167] Marmier G, Weigt M, Bitbol AF. Phylogenetic correlations can suffice to infer protein partners from sequences. *PLoS Comput Biol* 2019 Oct;15(10):e1007179.
- [168] Champeimont R, Laine E, Hu SW, Penin F, Carbone A. Coevolution analysis of Hepatitis C virus genome to identify the structural and functional dependency network of viral proteins. *Scientific reports* 2016;6(1):1–20.
- [169] Ghadie MA, Coulombe-Huntington J, Xia Y. Interaction evolution: insights from genome-wide analyses of protein-protein interactions. *Curr Opin Struct Biol* 2018 Jun;50:42–48.
- [170] Quadir F, Roy R, Halfmann R, Cheng J. Predicting interchain contacts for homodimeric and homomultimeric protein complexes using multiple sequence alignments of monomers and deep learning. *bioRxiv* 2020;<https://www.biorxiv.org/content/early/2020/11/10/2020.11.09.373878>.
- [171] Kundrotas PJ, Anishchenko I, Dauzhenka T, Kotthoff I, Mnevets D, Copeland MM, et al. Dockground: A comprehensive data resource for modeling of protein complexes. *Protein Sci* 2018 Jan;27(1):172–181.
- [172] Weinstein EN, Marks DS. A structured observation distribution for generative biological sequence prediction and forecasting. *bioRxiv* 2021;<https://www.biorxiv.org/content/early/2021/02/24/2020.07.31.231381>.
- [173] Trinquier J, Uguzzoni G, Pagnani A, Zamponi F, Weigt M. Efficient generative modeling of protein sequences using simple autoregressive models. *arXiv preprint arXiv:210303292* 2021;.
- [174] Laine E, Karami Y, Carbone A. GEMME: a simple and fast global epistatic model predicting mutational effects. *Molecular biology and evolution* 2019;36(11):2604–2619.
- [175] Louie RH, Kaczorowski KJ, Barton JP, Chakraborty AK, McKay MR. Fitness landscape of the human immunodeficiency virus envelope protein that is targeted by antibodies. *Proceedings of the National Academy of Sciences* 2018;115(4):E564–E573.
- [176] Riesselman AJ, Ingraham JB, Marks DS. Deep generative models of genetic variation capture the effects of mutations. *Nature methods* 2018;15(10):816–822.
- [177] Flynn WF, Haldane A, Torbett BE, Levy RM. Inference of epistatic effects leading to entrenchment and drug resistance in HIV-1 protease. *Molecular biology and evolution* 2017;34(6):1291–1306.
- [178] Hopf TA, Ingraham JB, Poelwijk FJ, Schärfe CP, Springer M, Sander C, et al. Mutation effects predicted from sequence co-variation. *Nature biotechnology* 2017;35(2):128–135.
- [179] Figliuzzi M, Jacquier H, Schug A, Tenaillon O, Weigt M. Coevolutionary landscape inference and the context-dependence of mutations in beta-lactamase TEM-1. *Molecular biology and evolution* 2016;33(1):268–280.
- [180] Mann JK, Barton JP, Ferguson AL, Omarjee S, Walker BD, Chakraborty A, et al. The fitness landscape of HIV-1 gag: advanced modeling approaches and validation of model predictions by in vitro testing. *PLoS Comput Biol* 2014;10(8):e1003776.
- [181] Ferguson AL, Mann JK, Omarjee S, Ndung'u T, Walker BD, Chakraborty AK. Translating HIV sequences into quantitative fitness landscapes predicts viral vulnerabilities for rational immunogen design. *Immunity* 2013;38(3):606–617.
- [182] Frazer J, Notin P, Dias M, Gomez A, Brock K, Gal Y, et al. Large-scale clinical interpretation of genetic variants using evolutionary data and deep learning. *bioRxiv* 2020;.
- [183] Russ WP, Figliuzzi M, Stocker C, Barrat-Charlaix P, Socolich M, Kast P, et al. An evolution-based model for designing chorisate mutase enzymes. *Science* 2020;369(6502):440–445.
- [184] Biswas S, Khimulya G, Alley EC, Esvelt KM, Church GM. Low-N protein engineering with data-efficient deep learning. *Nature Methods* 2021;18(4):389–396.
- [185] Rives A, Meier J, Sercu T, Goyal S, Lin Z, Liu J, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences* 2021;118(15).
- [186] Madani A, McCann B, Naik N, Keskar NS, Anand N, Eguchi RR, et al. Progen: Language modeling for protein generation. *arXiv preprint arXiv:200403497* 2020;.
- [187] Anishchenko I, Chidyausiku TM, Ovchinnikov S, Pellock SJ, Baker D. De novo protein design by deep network hallucination. *bioRxiv* 2020;.