Generative power of a protein language model trained on multiple sequence alignments

Damiano Sgarbossa[1,2], Umberto Lupo[1,2,*], Anne-Florence Bitbol[1,2,*]

1 Institute of Bioengineering, School of Life Sciences, École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
2 SIB Swiss Institute of Bioinformatics, CH-1015 Lausanne, Switzerland
* Email: umberto.lupo@epfl.ch, anne-florence.bitbol@epfl.ch

Abstract

Computational models starting from large ensembles of evolutionarily related protein sequences capture a representation of protein families and learn constraints associated to protein structure and function. They thus open the possibility for generating novel sequences belonging to protein families. Protein language models trained on multiple sequence alignments, such as MSA Transformer, are highly attractive candidates to this end. We propose and test an iterative method that directly employs the masked language modeling objective to generate sequences using MSA Transformer. We demonstrate that the resulting sequences score as well as natural sequences, for homology, coevolution and structure-based measures. For large protein families, our synthetic sequences have similar or better properties compared to sequences generated by Potts models, including experimentally-validated ones. Moreover, for small protein families, our generation method based on MSA Transformer outperforms Potts models. Our method also more accurately reproduces the higher-order statistics and the distribution of sequences in sequence space of natural data than Potts models. MSA Transformer is thus a strong candidate for protein sequence generation and protein design.

Introduction

Designing new proteins with specific structure and function is a highly important goal of bioengineering. Indeed, it can allow to tune their stability or their biochemical properties, including their enzymatic activities, enabling important medical applications. The search for novel proteins is difficult due to the huge size of protein sequence space: for instance, there are 20^{100} different possible sequences for a short protein domain with 100 amino acids. Furthermore, only a small fraction of this space comprises sequences that do fold, as demonstrated by experiments studying random sequences [1], and by theoretical arguments based on the physics of disordered systems [2]. *De novo* or rational protein design, which starts with target three-dimensional structures and physico-chemical potentials, can generate proteins which are not in a known protein family [3–5], but is generally restricted to small proteins [6]. Conversely, directed evolution allows to perform a local search of sequence space, but generally remains limited to the vicinity of a natural sequence [7].

Generative computational models that build on the breadth of available natural protein sequence data, and capture a representation of protein families, now offer great alternatives that

can allow to sample novel sequences belonging to protein families. In particular, Potts models, or DCA models [8–11], which are pairwise maximum entropy models trained to reproduce the one- and two-body statistics of the sequences of a family, allow direct sampling from a probability distribution modeling this family [12], and have been used successfully for protein design [13]. Variational autoencoders are deep learning models which also allow sampling, and they have been shown to successfully produce functional proteins [14], although their statistical properties appear to have a lower quality than with Potts models [15].

Protein language models are deep learning models based on natural language processing methods, especially attention [16] and transformers [17]. They are trained on large ensembles of protein sequences, and capture long-range dependencies within a protein sequence [18–25]. These pre-trained models are able to predict structure in an unsupervised way [21], either taking as input a single sequence [20] or an MSA [25], potentially by transferring knowledge from their large training set [26, 27]. The great success of supervised protein structure prediction by AlphaFold [28] is partly based on the use of transformers. It is therefore of strong interest to assess the generative ability of protein language models, and recent works show that this has high potential [24, 29–32].

Correlations in amino-acid usage that can be observed between the columns of multiple sequence alignments (MSAs) of homologous proteins [33–35] were experimentally demonstrated to be highly important to generate functional synthetic proteins [1, 36]. The importance of pairwise coevolution signals was then corroborated by the success of Potts models at predicting structural contacts [8–11, 37], analyzing mutational effects [38–41], protein evolution [42] and conformational changes [43, 44], designing proteins [13], and at predicting protein-protein interaction partners [45–48]. Protein language models that take MSAs as input [25, 28] are able to directly exploit this covariation signal, and are thus particularly interesting candidates for protein design. Thus motivated, we focus on MSA Transformer [25], a protein language model which was trained on MSAs using the masked language modeling objective, without additional supervised training – by contrast to AlphaFold [28]. We ask how the generative properties of MSA Transformer compare to those of Boltzmann machine DCA (bmDCA) [12, 13], a state-of-the-art generative Potts model.

We propose and test a generating method that directly uses the masked language modeling objective in an iterative way to generate sequences using MSA Transformer. Using homology, coevolution and structural scores, we demonstrate that the sequences generated by this method score as well as natural sequences. We further show that this good performance is not restricted to synthetic sequences that are very similar to natural sequences. For large protein families, our synthetic sequences have homology and structure-based scores as good as or better than sequences generated by bmDCA, and have similar properties to experimentally-validated bmDCA-generated sequences. Moreover, for small protein families, our generation method based on MSA Transformer outperforms bmDCA, by providing synthetic sequences that score well without being extremely similar to natural ones. However, we find that bmDCA better reproduces the one- and two-body statistics of the natural MSAs than MSA Transformer when used with default parameters, consistently with its training objective. Interestingly, the opposite generally holds for higher-order statistics. MSA-Transformer–generated sequences also better reproduce the distribution of sequences in sequence space than bmDCA-generated ones. Our conclusion is that MSA Transformer is a strong candidate for protein sequence generation and protein design.

Results

An iterative masking procedure allows MSA Transformer to generate novel sequences with high scores

Can the protein language model MSA Transformer [25] be used to generate sequences that are credible members of protein families? How do its generative abilities compare to Potts models inferred by Boltzmann machine DCA (bmDCA) [12], a state-of-the-art generative DCA method which has been experimentally shown to generate functional proteins [13]? To address these questions, we developed and employed an iterative masking procedure to generate synthetic MSAs from natural MSAs of 14 different large Pfam protein families (see Table 1) and 7 small ones (see Table S1) with MSA Transformer, as described in "Using MSA Transformer to generate sequences via an iterative masking procedure". We also generated synthetic sequences by Markov Chain Monte Carlo (MCMC) sampling from Potts models inferred from these MSAs by bmDCA, using two variants that differ by sampling temperature T and regularization strength λ, respectively matching the default parameters employed in [12], and some of those used in [13], see "Sampling sequences from Potts models" for details. For each protein family, we thus obtained four different MSAs of the same depth: the natural one, the one generated by our iterative masking procedure using MSA Transformer, and the two MSAs sampled from the inferred Potts model. To characterize each sequence, we consider four different scores (see "Scoring individual sequences"). First, we assess the quality of the generated sequences as homologs of the protein family of interest; we do this via the HMMER (http://hmmer.org) score of the hidden Markov model employed by Pfam to retrieve natural homologs. Second, we consider a score that accounts for coevolution between amino-acid sites, namely the statistical energy score from the Potts model fitted on the natural MSA. Third, we determine AlphaFold's confidence in its determination of the three-dimensional structure of these sequences, via the predicted local-distance difference test (pLDDT) score. Fourth, to assess whether the predicted structures are similar to the native ones, we compute the root-mean-squared deviation (RMSD) between a reference experimental structure and the AlphaFold predicted structures. The first three scores are such that higher values are better, while smaller RMSD values indicate that predicted structures are similar to the native ones. Together, these scores account for very different aspects of proteins, namely homology, coevolution and structure.

Let us first consider the 14 large protein families in Table 1, where MSAs are deep enough to accurately fit Potts models using bmDCA [12]. Fig. 1 shows that, for all these protein families, and for these four different scores, the sequences generated by MSA Transformer using our iterative masking procedure have scores that are at least as good as those of natural sequences, and better than those of sequences generated by bmDCA with default parameters [12], as confirmed by the Kolmogorov–Smirnov test (see Table S2). Decreasing the sampling temperature and the regularization strength used with bmDCA improves the statistical energy score as expected [13], but also other scores. These other scores, and most importantly our two structural scores, are similar or better for MSA-Transformer–generated sequences compared to those generated by bmDCA with non-default parameters. In particular, the median pLDDT score is larger for the former than for the latter in 11 protein families out of 14, by a margin larger than the standard deviation in 4 of them (see Table S3). These results demonstrate that MSA Transformer is a good candidate to generate synthetic sequences from protein families, and that our iterative masking procedure allows to perform such generation.

How different are these synthetic sequences from the natural ones? In particular, are the best-scoring sequences novel, or are they almost copies of natural sequences? In Fig. 2 we show, for two example protein families (PF00072 and PF00153), the HMMER score and the DCA

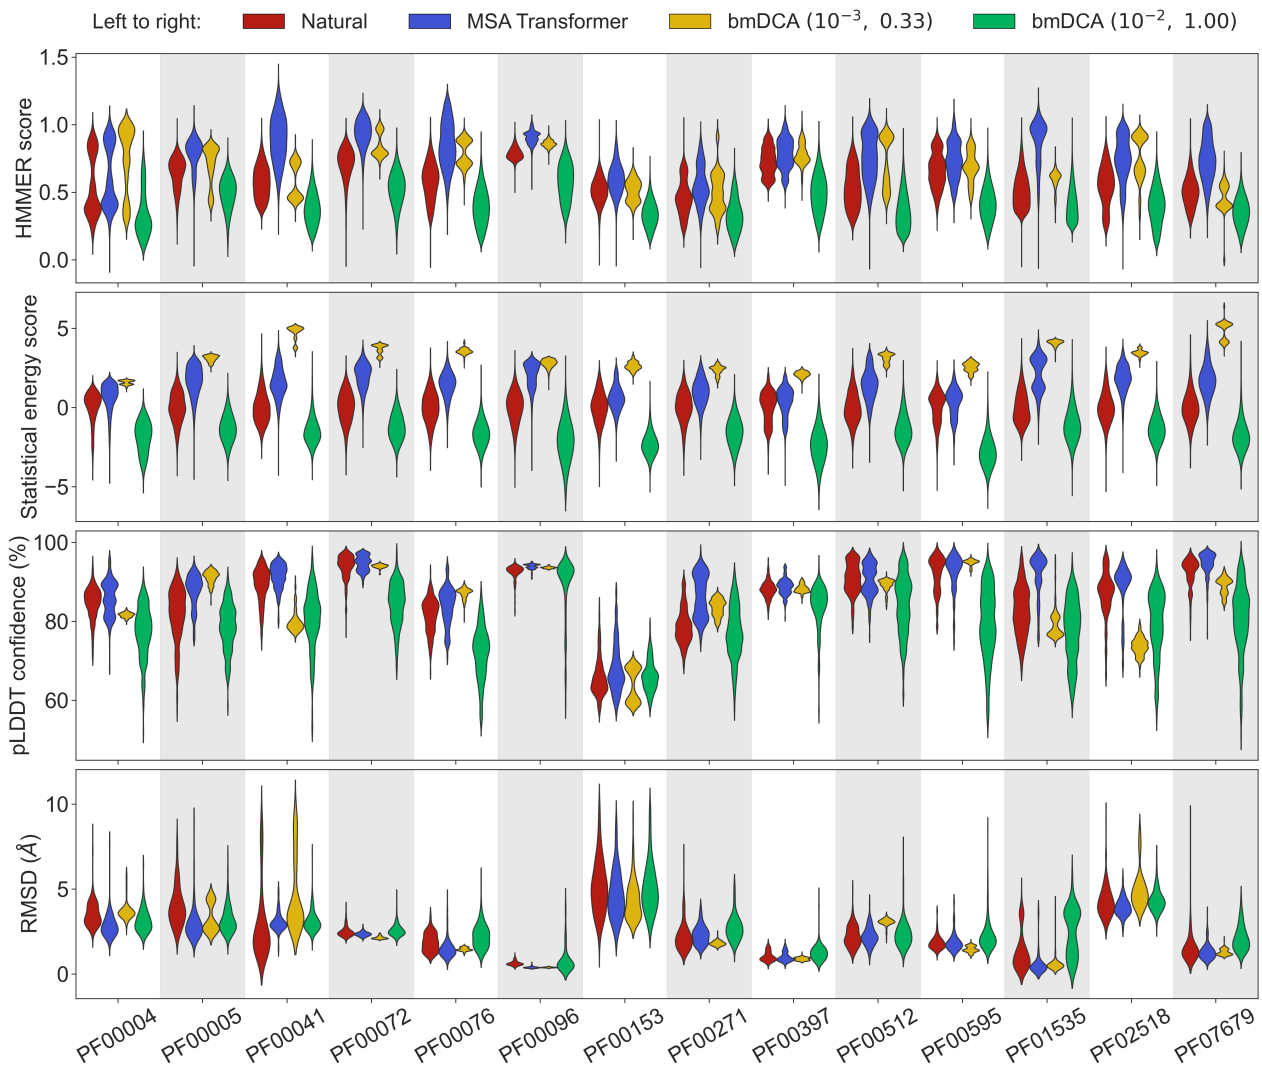

Figure 1: **Comparison of homology, coevolution, and structure-based scores between natural sequences and sequences generated by MSA Transformer or bmDCA.** For each Pfam family in Table 1, we compare a natural MSA from Pfam and three synthetic MSAs of the same depth. The first synthetic MSA was obtained using MSA Transformer via our iterative masking procedure, and the second and third ones were generated by a Potts model inferred from the natural MSA using bmDCA with two different pairs (λ, T) of regularization strength λ and sampling temperature T. For each of the four scores described in "Scoring individual sequences", we show the distributions of score values among sequences in each MSA as a violin plot. Higher score values are better for all scores except RMSD (bottom panel), where smaller values indicate a closer match to an experimental structure. **Top panel:** For each Pfam family, HMMER scores are divided by the highest score found in the natural MSA. Note that sequences below HMMER's default homology detection score (E-value larger than 10), and whose HMMER score is thus 0, are not shown (the median over families of the fraction of such sequences is 2% for bmDCA$(10^{-2}, 1.00)$-generated MSAs, while there are no such sequences among the MSA-Transformer–generated ones). **Second panel:** Statistical energy scores are defined as minus the bmDCA statistical energies. To accommodate the highly family-dependent ranges of these scores, for each Pfam family we show their values after shifting by the mean score in the natural MSA, and normalizing by the standard deviation of natural MSA scores. **Third panel:** AlphaFold's pLDDT confidence scores. **Bottom panel:** RMSD of predicted structures with respect to the experimental structures in Table 1. Structural scores (pLDDT and RMSD) were computed on 200 randomly chosen sequences from each MSA. All kernel-smoothed histograms are normalized such that all violins have the same maximal width. Outliers (less than 1% in all cases) were discarded for legibility.

statistical energy score versus the sequence's Hamming distance to its closest natural sequence in the natural MSA.

From the marginal distributions of the Hamming distances in Fig. 2, we observe that MSA Transformer generates sequences with variable distances to their closest natural sequences, and that these distances are overall larger than those between natural sequences and their closest neighbor (excluding themselves). With default parameters, bmDCA generates sequences which are generally very different from the natural ones, but decreasing sampling temperature makes bmDCA–generated sequences more similar to natural ones and to each other, see Fig. S1. Besides, the marginal distributions of scores illustrate the general observation made on Fig. 1 and in Table S2 that MSA-Transformer–generated sequences have good scores. Moreover, the plots in Fig. 2 reveal that the MSA-Transformer–generated sequences featuring the highest HMMER scores tend to have large Hamming distances to natural sequences, i.e. to be truly novel (see also "Choosing parameters in the iterative masking procedure"). We observe these trends for most large protein families studied, and they are robust to using BLOSUM similarity scores [49] instead of Hamming distances. Therefore, our sequence generation method based on MSA Transformer is not reaching good scores by just reproducing natural sequences. Besides, the diversity of MSA-Transformer–generated MSAs, as measured by their effective depth (Eq. (8)), is only slightly smaller than that of natural MSAs (see Fig. S1). Conversely, bmDCA at low temperature produces highly redundant sequences (Fig. S1), which are concentrated in specific regions of the scores vs. distance space in Fig. 2. Indeed, sequence generation by bmDCA is then constrained to exploring the local minima of the Potts model energy landscapes.

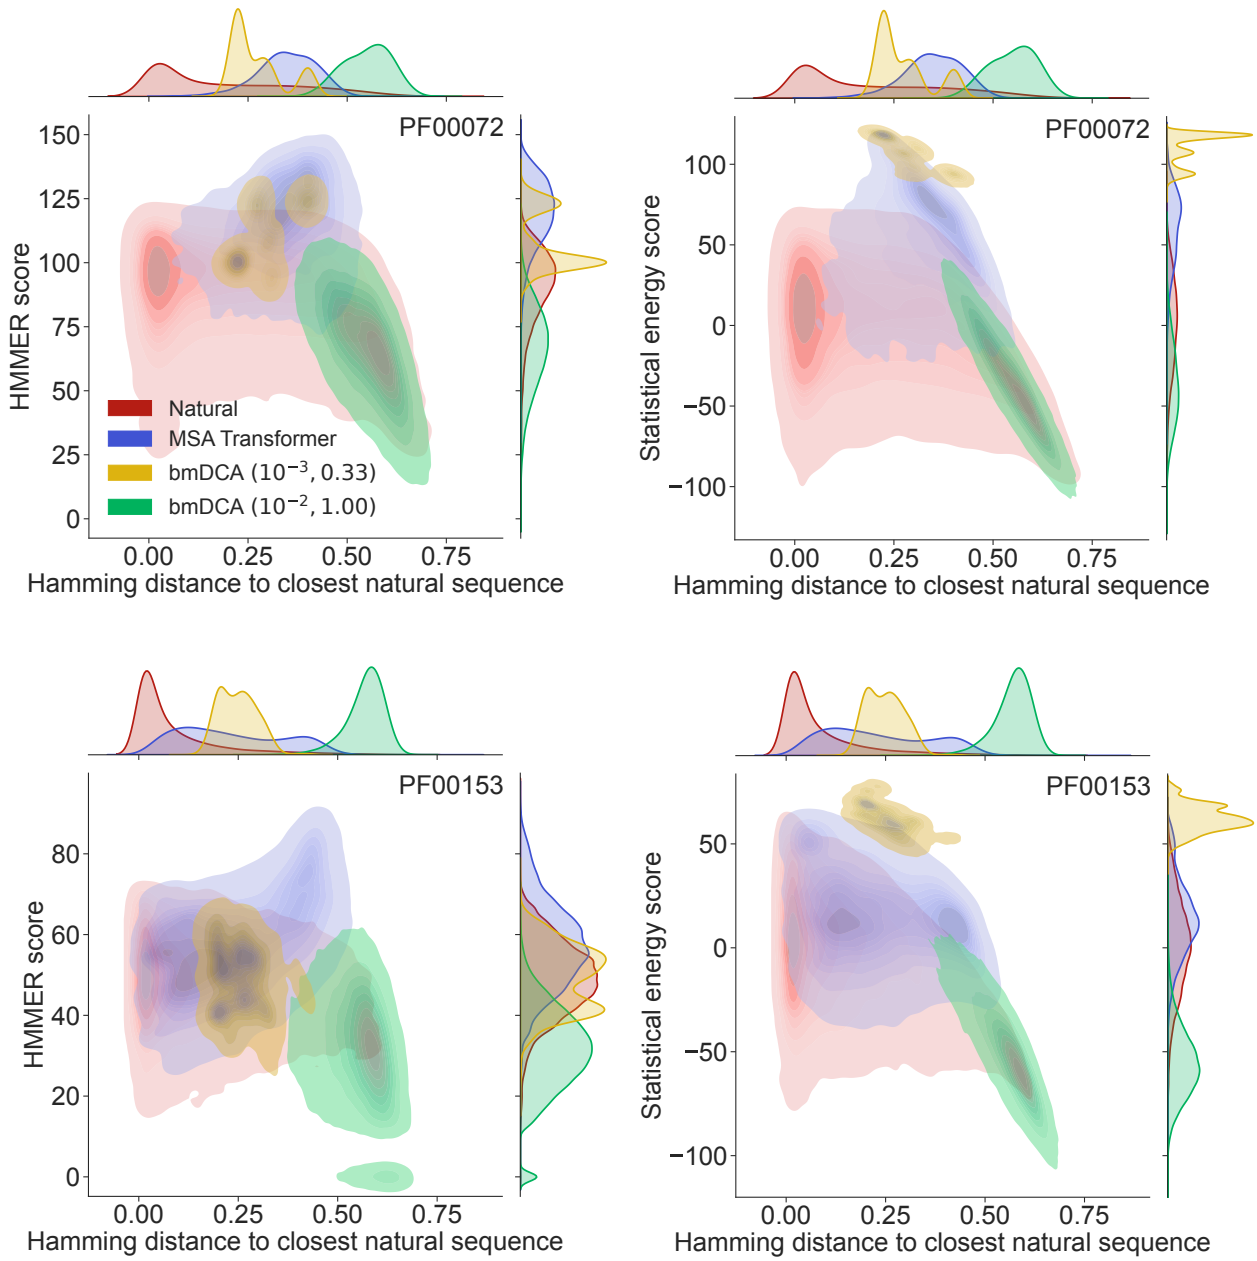

Figure 2: **Homology and coevolution scores vs. distance to the natural MSA, for protein families PF00072 and PF00153.** We show contour plots of the HMMER score and the statistical energy score (defined as minus the DCA statistical energy, shifted by its mean value in the natural MSA) versus the Hamming distance of each sequence to the closest natural sequence (which is not itself, in the case of natural sequences). Results are shown for natural sequences and for sequences generated using MSA Transformer and bmDCA (the same two (λ, T) pairs as in Fig. 1 are used for bmDCA). The lightest contours shown include 99% of the cumulative probability mass.

Sequence generation by the iterative masking procedure is successful for small protein families

Accurately fitting Potts models requires deep and diverse MSAs, as evidenced by the strong dependence of structural contact prediction by Potts models on MSA depth [9, 10]. By contrast, MSA Transformer was trained on many MSAs, and is able to transfer knowledge across protein families. It outperforms Potts models at unsupervised contact prediction most strongly for shallow MSAs [25]. How does sequence generation using our iterative masking procedure based on MSA Transformer compare to bmDCA in the case of small protein families?

To address this question, we generated synthetic MSAs starting from 7 small families, using both our iterative masking procedure based on MSA Transformer and bmDCA with default parameters and with low sampling temperature. Fig. 3 reports all four scores discussed above in the case of these 7 small families, listed in Table S1. We observe that MSA-Transformer–generated sequences have similar HMMER scores and structural scores to natural sequences. MSA-Transformer–generated sequences also generally have better HMMER scores and structural scores than those generated by bmDCA with default parameters. While low-temperature bmDCA yields better statistical energy scores (as expected), and also gives HMMER scores and structural scores comparable to natural sequences, it in fact generates sequences that are almost exact copies of natural ones (see Fig. 3, bottom row). By contrast, MSA Transformer produces sequences that are quite different from natural ones, and have very good scores. Thus, our method based on MSA Transformer is particularly promising in the tricky case of small protein families.

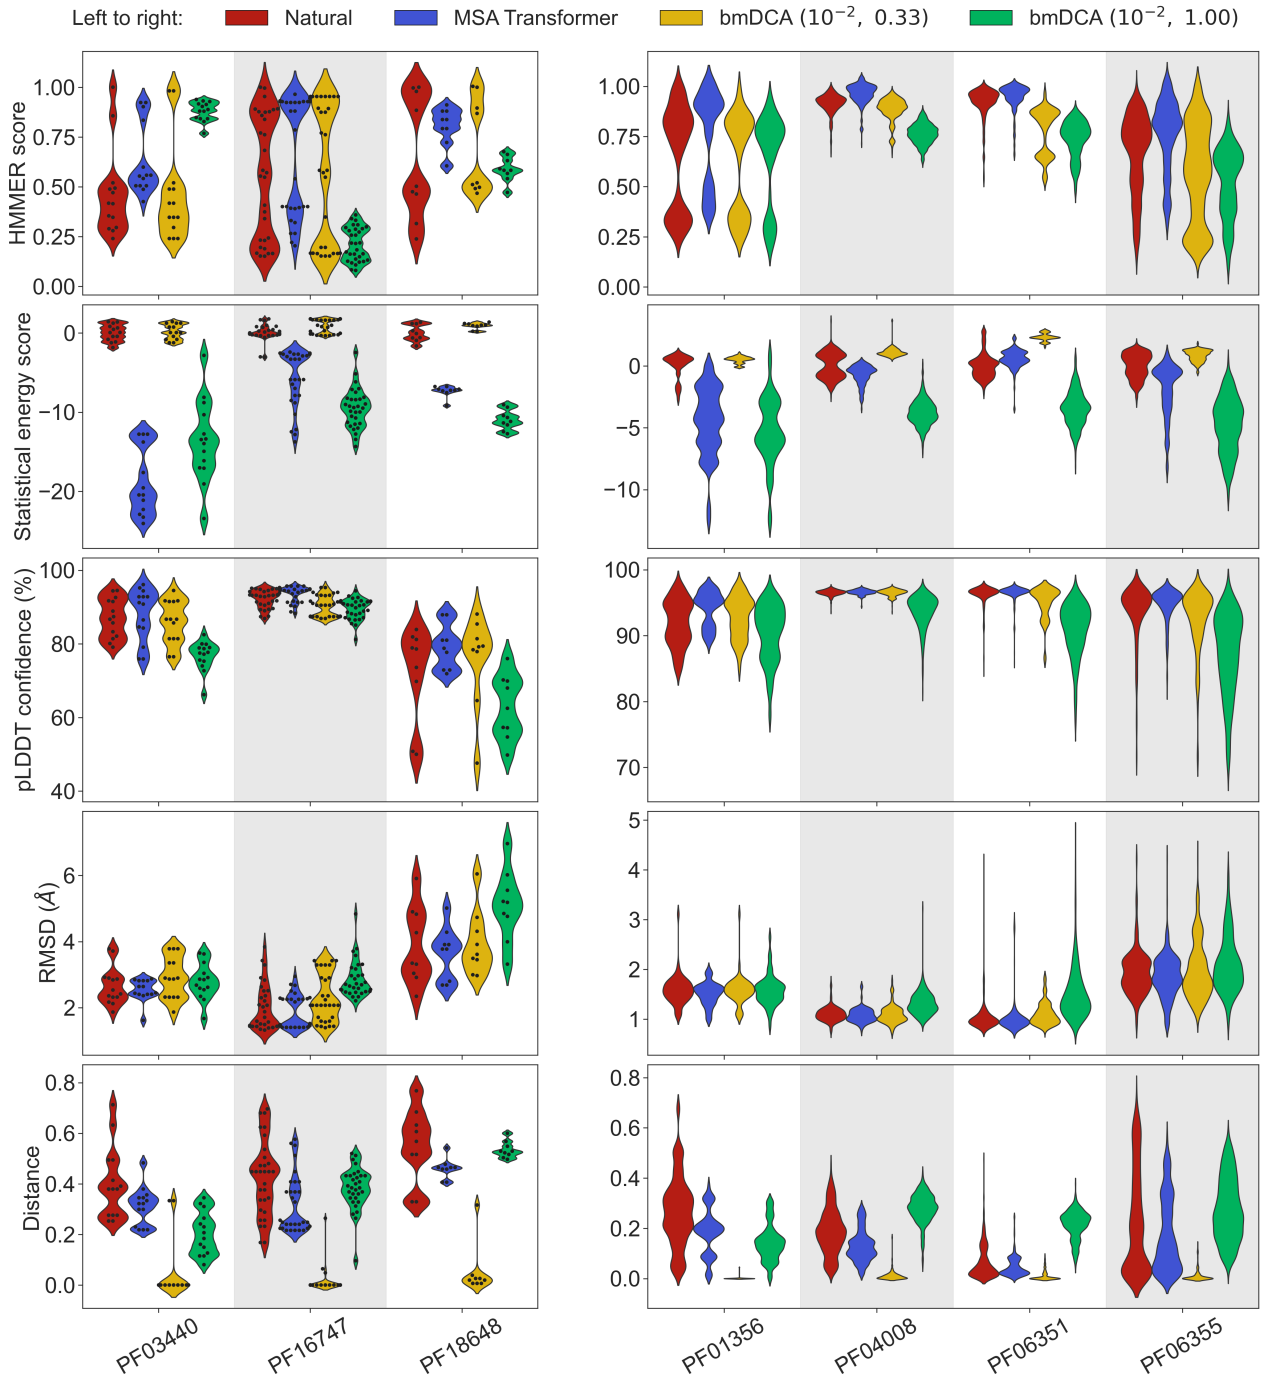

Figure 3: **Application of our sequence generation method based on MSA Transformer to small protein families.** We consider 7 small protein families, with natural MSAs that comprise from 9 to a few hundreds of sequences, see Table S1. As in Fig. 1, for each family, we compare the natural MSA and three synthetic MSAs of the same depth. In all cases, we show violin plots of the same four scores as for large families in Fig. 1, as well as of the Hamming distance to the closest natural sequence, which is not itself in the case of natural sequences ("Distance"). For the three smallest families (left panel; fewer than 40 sequences), we also show the score of each individual sequence as a swarm plot. Note that while we employ the same sampling temperatures T as in Fig. 1 for bmDCA, here, we use regularization strength $\lambda = 10^{-2}$ throughout, due to MSA shallowness (see "Sampling sequences from Potts models").

Higher-order statistics are better reproduced by MSA Transformer, while lower-order statistics are better reproduced by bmDCA

How well do synthetic MSAs generated by our method based on MSA Transformer, and by bmDCA, reproduce the statistics of amino-acid usage observed in natural MSAs? To address this question, we consider the r20 score [15, 50], which quantifies the statistical similarity of two datasets at various orders (see "Analyzing the statistics of MSAs"). We compute it between each of our synthetic MSAs and the corresponding natural one, for the 14 large protein families in Table 1. We also present as reference an assumption-free null model, namely the r20 score between two subsets of each natural MSA. Fig. 4 shows that bmDCA with default parameters is most often the best method at reproducing lower-order statistics, while MSA Transformer is the best at reproducing higher-order statistics, in all families considered. bmDCA at lower temperature performs more poorly at reproducing the statistics of natural MSAs than other methods, because low temperature biases the sampling (bmDCA models are effectively learned at temperature $T = 1$).

To have a more detailed insight into lower-order correlations, we estimate frequencies and information theory measures, at the one-, two- and three-body level, from our natural and synthetic MSAs, and compare them (see "Analyzing the statistics of MSAs"). Fig. S2 shows that one- and two-body statistics are generally better reproduced by bmDCA with default parameters than by MSA Transformer, while results are more mixed for three-body statistics. Figs. S3 and S4 show a comparison of second- and third-order connected correlations for PF00072 and PF00153. For PF00072, bmDCA reproduces better the second- but also third-order connected correlations of the natural data than MSA Transformer, while for PF00153, MSA Transformer reproduces the third-order connected correlations better than bmDCA, consistently with Fig. 4. Potts models are pairwise maximum entropy models constrained to match the one- and two-body frequencies from natural MSAs. Thus, bmDCA is trained to reproduce these frequencies, and achieves these objectives quite well, although the comparison to the null model in Figs. S3 and S4 hints that further improvements remain possible, see [51]. MSA Transformer has entirely different training objectives, but, interestingly, it performs comparably at reproducing three-body statistics and is better at reproducing even higher-order statistics than bmDCA.

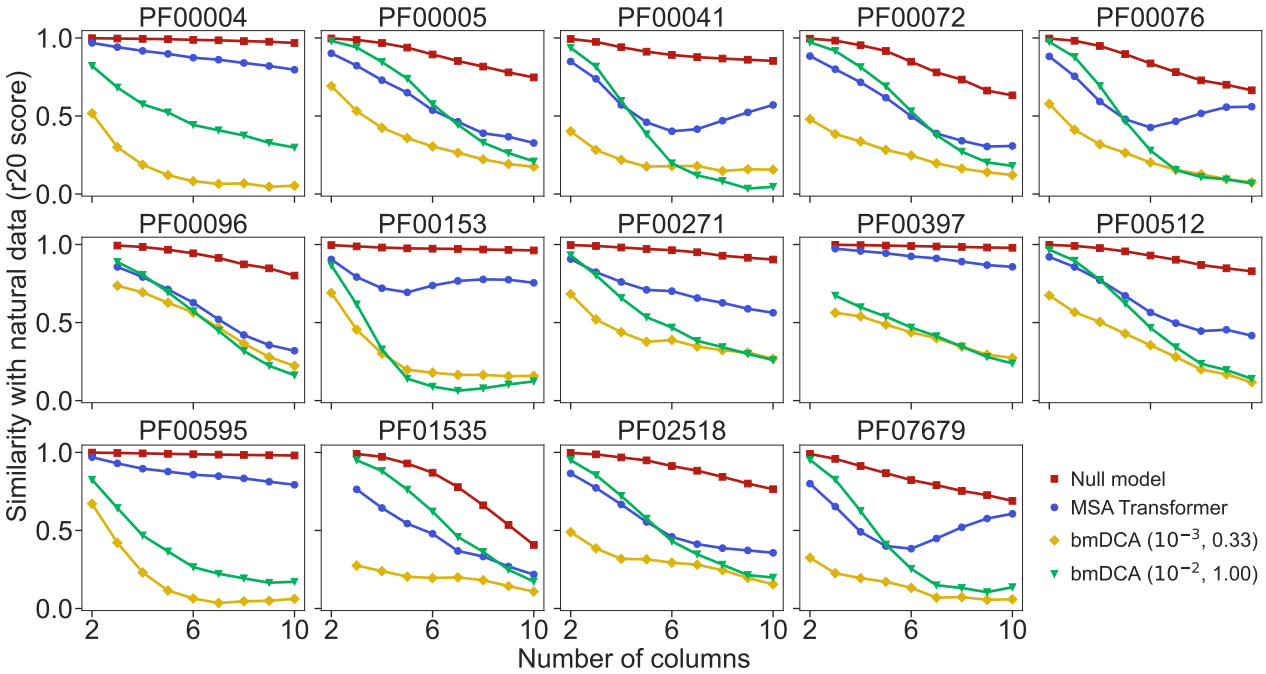

Figure 4: **Similarity of statistics between synthetic and natural MSAs.** To compare the statistics of synthetic and natural MSAs at various orders, we compute r20 scores [15, 50], and plot them versus the number of different MSA columns that are considered (see "Analyzing the statistics of MSAs" for details). All families in Table 1 are considered. For each of them, the reference MSA comprises either half of the natural MSA (with sequences selected uniformly at random), or 30,000 sequences from it if the natural MSA depth is larger than 60,000. The null model compares the other half of the natural MSA to this reference MSA. It yields an estimate of the expected r20 scores due only to finite-size effects in a model-free, purely data-driven way.

MSA Transformer captures well the distribution of sequences in sequence space

How are synthetic MSAs generated by MSA Transformer and bmDCA impacted by the heterogeneous repartition of natural sequences in sequence space? While natural protein sequences in a family have evolved from a common ancestor along a phylogeny, synthetic sequences do not have a real evolutionary history. However, as bmDCA and MSA Transformer are trained on natural data, they can capture phylogenetic correlations [52]. Besides, inferred Potts models are known to be impacted by phylogenetic correlations [8, 10, 53–59].

First, to assess whether generated sequences most resemble natural ones that are well represented in their family or, rather, rare ones, we consider the closest natural sequence to each synthetic sequence, and count the neighbors of this natural sequence in the natural MSA (see "Characterizing the distribution of sequences in MSAs"). Fig. 5 compares the distribution of these numbers of neighbors for natural sequences and for the closest natural sequences to generated sequences, in the cases of PF00072 and PF00153. It shows that bmDCA generates sequences similar to natural sequences with fewer neighbors than typical in the natural data. Conversely, MSA Transformer generates sequences whose closest natural sequences have a distribution of number of neighbors similar to that of the natural MSA. This suggests that our generation method based on MSA Transformer tends to sample from denser regions of the sequence space than bmDCA, while not reproducing natural sequences (see also Fig. 2 and "Choosing parameters in the iterative masking procedure").

To analyze the distribution of MSA sequences in sequence space in more detail, we perform a principal component analysis of one-hot encoded MSAs, and focus on the top two principal components (PCs) [12] (see "Characterizing the distribution of sequences in MSAs"). Fig. 6 shows the distribution of sequences in the space spanned by these top two PCs, for natural and synthetic MSAs, in the cases of PF00072 and PF00153. We observe that MSA Transformer is able to generate sequences with a distribution in sequence space that is very similar to that of the natural MSA. Conversely, bmDCA captures the overall shape of this distribution, but appears to smooth it compared to the natural data with default parameters and to restrict to sparse regions of the sequence space at low temperature, consistently with our previous results. These observations are general across all the deep MSAs we considered (see Fig. S5). Note that a limitation of this analysis is that the top two PCs explain a small fraction of the variance in all cases (see Fig. 6).

Finally, to analyze in more detail the apparent relatedness of generated sequences, and compare it to real phylogenetic relationships in natural sequences, we infer phylogenetic trees from each synthetic and natural MSA, and analyze the eigenvalue spectrum of their modified graph Laplacian (MGL) to compare them [60] (see "Characterizing the distribution of sequences in MSAs"). Fig. S6 compares the density of these eigenvalue spectra for natural and synthetic MSAs regarding families PF00072 and PF00153. The skewness and the position of such distributions are indicators of the topology of the tree. In particular, distributions with negative skewness (right unbalanced) or which are shifted to the right, correspond to "tippy" trees, while the opposite case corresponds to "stemmy" trees [60], which feature an accumulation of recent speciation events (short leaves length) [61]. In this light, Fig. S6 shows that both MSA Transformer and low-temperature bmDCA generate sequences with an apparent phylogeny that is more stemmy than the natural one, while bmDCA with default parameters yields a slightly more tippy tree. This is consistent with our observations regarding sequence diversity, which is larger than in natural data for bmDCA with default parameters, slightly smaller than in natural data using MSA Transformer and much lower using low-temperature bmDCA (see Fig. S1).

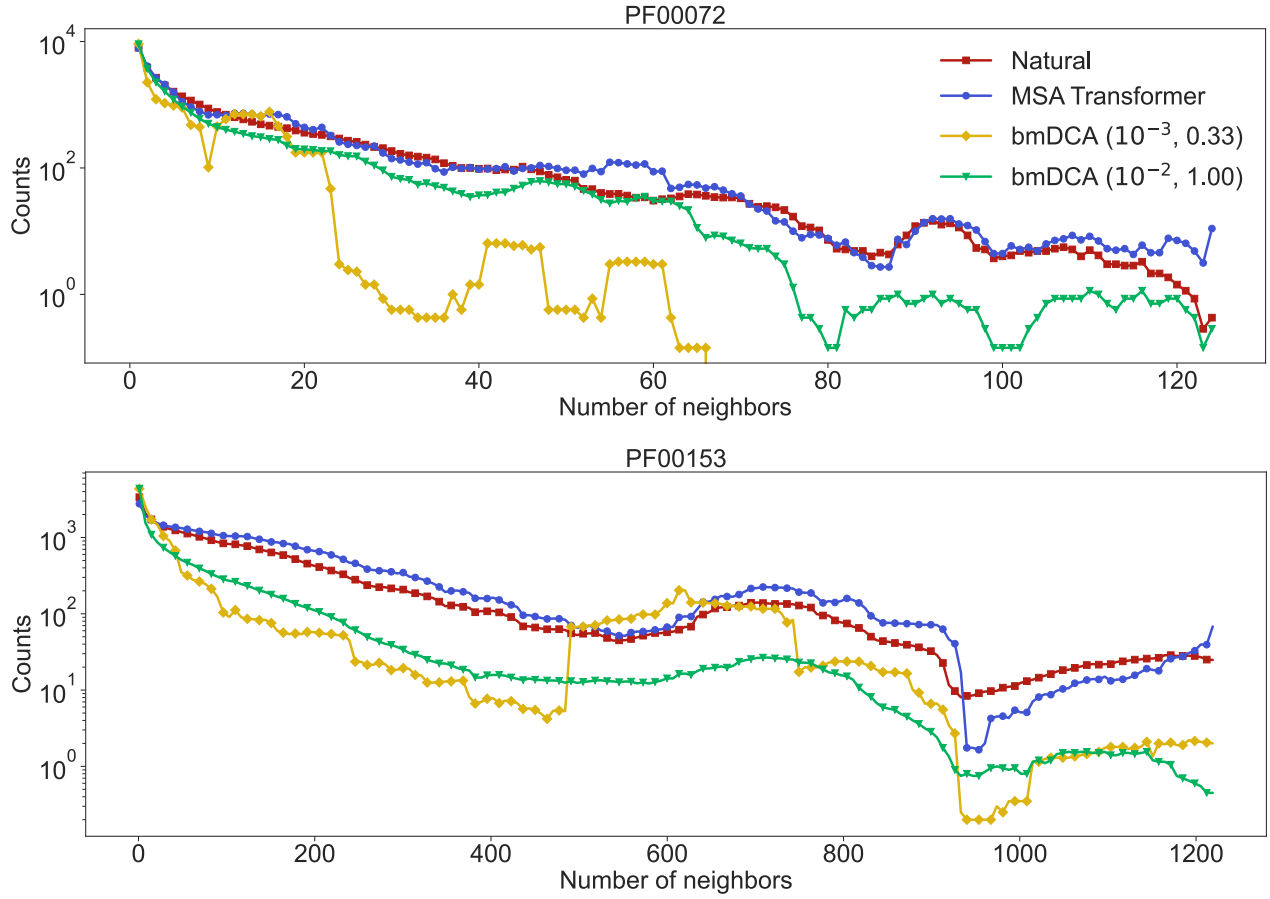

Figure 5: **Neighbors of natural and synthetic sequences, for families PF00072 and PF00153.** We show the distribution of the number of neighbors of sequences in the natural MSA, and the distribution of the number of neighbors of the closest natural sequence to each of our generated sequences. Given a sequence in a natural MSA, its number of neighbors is the number of natural sequences that are within a (normalized) Hamming distance $\delta = 0.2$ from it. The moving average of the results is shown, using a window representing 5% of the total number of points.

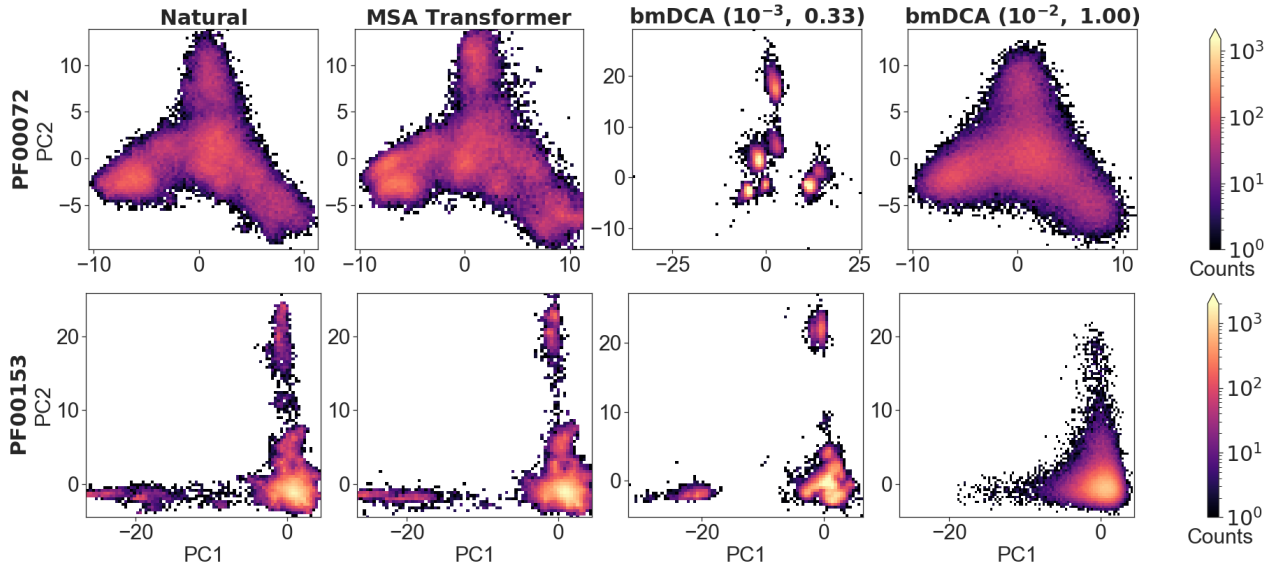

Figure 6: **Distribution of sequences in sequence space, for families PF00072 and PF00153.** We show the distribution of one-hot encoded natural and synthetic sequences projected in the subspace of the first two principal components of the natural MSA. The same axis limits are used within one family, except for bmDCA (10^{-3}, 0.33) in the case of PF00072. Note that the fraction of the total variance explained by the first two principal components of each MSA is less than 4% for all families and all generation methods.

Comparison with published experimental datasets

How do the sequences generated by our method based on MSA Transformer compare to published protein design experimental datasets? Recently, sequences sampled from a bmDCA Potts model of the chorismate mutase protein family were experimentally demonstrated to be functional [13]. In Fig. S7, we show plots analogous to those in Fig. 2, plus additional ones for our two structural scores (pLDDT and RMSD), in the case of chorismate mutase. This allows a detailed comparison between the sequences we generate using MSA Transformer and the sequences designed in [13] using bmDCA with a combination of different temperatures and regularization strengths. We find that our method based on MSA Transformer produces sequences that score as well as artificial sequences which have been tested experimentally. Besides, we obtained these results without fine-tuning the parameters of our generative procedure to this family, while several specific combinations of parameters were used in [13].

To further compare our generated sequences to those tested experimentally in [13], we focus on the top third of sequences in terms of pLDDT scores, as it was shown in [62] that good structural scores help to select functional sequences. Relative enrichment, which is the experimental score used in [13] to assess the function of chorismate mutase enzymes, was measured in [13] for all sequences in the natural MSA and for sequences generated with bmDCA in that work. We estimate the expected relative enrichment of our generated sequences as the relative enrichment of the closest natural sequence. To test our estimation procedure, we estimate the relative enrichments of the bmDCA-generated sequences from [13], and we compare them to the experimentally measured values. Fig. S8 shows that sequences with a high (resp. low) estimated score have a high (resp. low) experimental score too. Next, we compare the distributions of estimated relative enrichment for sequences generated using our method based on MSA-Transformer and for the bmDCA-generated sequences from [13]. Fig. S8 shows that they are quite similar to each other and to the distribution of measured relative enrichment for bmDCA-generated sequences. In addition, a slightly larger fraction of MSA-Transformer–generated sequences than of bmDCA–generated sequences have a large estimated

relative enrichment. This suggests that our method based on MSA Transformer should be able to generate functional sequences.

While the data from [13] is particularly well-suited to retrospectively evaluate our sequence generation method, we also propose a comparison of the distributions of scores based on experimental deep mutational scans (DMS) for protein families PF00595 [63] and PF13354 [64]. We compute these DMS scores for each natural and synthetic sequence, by summing the experimentally-measured effects of the relevant single-point mutations with respect to the reference sequence of the experimental studies. Fig. S9 shows the distribution of the DMS scores of natural and generated sequences for these two families. Our sequence generation method based on MSA Transformer better reproduces the DMS score distribution of natural sequences than bmDCA, and generates sequences with better average DMS scores. Despite the potential limitations of our DMS scores, e.g. their additivity, these results corroborate our other findings and provide further encouragement for our sequence generation based on MSA Transformer.

Discussion

In this work, we proposed an iterative masking procedure which directly exploits the masked language modeling objective of protein language models to generate sequences using the MSA-based neural language model MSA Transformer. We found that these sequences score as well as natural ones on three very different aspects, namely homology, coevolution and structure-based scores. For large protein families, our synthetic sequences have homology and structure-based scores at least as good as bmDCA-generated sequences, and have similar properties to experimentally-validated ones. Moreover, our generation method based on MSA Transformer is less limited by shallow MSAs than bmDCA, and is thus particularly promising for small protein families. Besides, MSA-Transformer–generated sequences better reproduce the higher-order statistics and the distribution of sequences in sequence space of natural data than bmDCA-generated ones. Conversely, bmDCA, with default parameters, better reproduces first- and second-order statistics, consistently with its training objective.

Our results are highly promising for sequence generation by MSA-based protein language models, and we hope that they will motivate further studies, especially experimental tests. They also show that protein deep learning models based on the masked language modeling objective have great generative potential, despite not being obvious generative models. More generally, our results reinforce the new promising "coevolution-driven" protein design approach of learning from sequences of evolutionarily related proteins the constraints associated to protein structure and function. This concept differs from structure- and physics-based *de novo* design [3–5], and from the new possibility to use supervised deep learning models able to accurately predict protein structures [28, 65, 66] for structure-driven sequence generation [67]. One can view the coevolution-driven approach as intermediate between structure-based approaches and directed evolution ones [7]. The coevolution-driven approach was recently experimentally validated in the case of bmDCA Potts models, which capture pairwise coevolution patterns in MSAs [13], and for variational autoencoders [14, 15]. Protein language models trained on multiple sequence alignments provide state-of-the-art unsupervised contact prediction and are able to capture coevolutionary patterns in their tied row attentions [25], and capture phylogenetic relationships in column attentions [52]. This makes them ideal candidates to generate new protein sequences from given families. However, contrary to Potts models and variational autoencoders [15], they do not allow direct sampling from a probability distribution over sequences [68]. Here, we demonstrated the power of a simple generation method directly based on the masked language modeling objective used for the training of MSA-based protein language models. It differs from using a decoder, which, though designed to perform autoregressive generation of amino acids to form a new sequence, requires training a full encoder-decoder model and

learning a parametric function mapping an MSA to a distribution over its sequences [30]. We instead directly employed the representation of protein families captured by the self-supervised model MSA Transformer to generate sequences. More sophisticated sampling methods could be considered along this line [68], but our minimal approach already gives very promising results.

We have focused on a large protein language model and compared it to the simplest model capturing coevolution, namely the Potts model, but we note that interpretable models of intermediate complexity such as restricted Boltzmann machines [69] could also be explored for coevolution-driven protein design. All these methods rely on MSAs; this is very useful to capture coevolution, but also means that one has to rely on potentially imperfect alignments. Thus, starting from alignment-free methods [24, 70, 71] also constitutes a promising direction.

Methods

Using MSA Transformer to generate sequences via an iterative masking procedure

Iterative masking procedure. In order to generate new sequences using MSA Transformer, we directly leverage the model's ability to assign, to arbitrary masked residue positions, a probability for each of the possible amino-acid tokens, given by the softmax of the model's output logits [21, 68, 72]. Indeed, in its pre-training, MSA Transformer applies the masked language modeling (MLM) objective to a training set of 26 million MSAs [21]. For this, it minimizes a pseudolikelihood loss, which reads, for an MSA $\mathcal{M}$, and a version $\widetilde{\mathcal{M}}$ of $\mathcal{M}$ in which some amino acids (those in a "mask") are masked:

$$\mathcal{L}_{\text{MLM}}(\mathcal{M}, \widetilde{\mathcal{M}}; \theta) = -\sum_{(m,i)\,\in\,\text{mask}} \log p(x_{m,i} \,|\, \widetilde{\mathcal{M}}; \theta)\,. \tag{1}$$

Here, $x_{m,i}$ denotes the amino acid at the i-th residue position in the m-th sequence of $\mathcal{M}$, and θ denotes all model parameters. For each position i in each sequence m, the model outputs one value ("logit") per amino-acid/gap symbol, and softmax-normalizing all values from this location in the MSA yields the conditional probabilities $p(x_{m,i} \,|\, \widetilde{\mathcal{M}}; \theta)$ in Eq. (1), which are then summed over the subset of masked MSA locations.

We propose an iterative masking procedure (see Fig. 7) which, given an arbitrary MSA $\mathcal{M}$ of natural sequences, proceeds as follows:

1. If necessary, subsample $\mathcal{M}$ to obtain an input MSA $\mathcal{M}'$ for MSA Transformer. The depth of $\mathcal{M}'$ is chosen given the memory footprint of MSA Transformer. In practice, we use input MSAs containing 600 sequences,[1] picked uniformly at random from our natural MSAs. Note that, for large protein families, multiple 600-sequence MSAs obtained using the procedure presented here are then combined into a single MSA of the same depth as the natural one (see below).

2. Randomly mask each residue of $\mathcal{M}'$ with a masking probability p, otherwise leave it unchanged. In practice, we choose $p = 0.1$ (see "Choosing parameters in the iterative masking procedure").

3. Feed the masked MSA to the model, and fill each masked entry with the token with highest probability (obtained from the model's output logits).

[1] During training, the authors of Ref. [25] kept $LM < 2^{14}$, where L is sequence length and M is MSA depth. However, we found that during inference we can use 2^{17} tokens on an Nvidia V100 32 GB GPU.

4. Repeat Steps 2 and 3 a number of times. In practice, we stop the algorithm after $I = 200$ iterations.

As natural MSAs, we use Pfam full MSAs for 14 protein families, described in "Datasets". For each natural MSA $\mathcal{M}$, we repeat the procedure above multiple times, sampling sequences each time from $\mathcal{M}$ without replacement to obtain a different input MSA $\mathcal{M}'$ in Step 1, until all the sequences in $\mathcal{M}$ are used. Note that sequences remain aligned at all times during the procedure. Combining the MSAs resulting from all these batches then yields a synthetic MSA with the same depth as the natural one, which ensures that the statistical properties of the synthetic MSA are subject to the same magnitude of finite-size errors as those of the natural MSA.

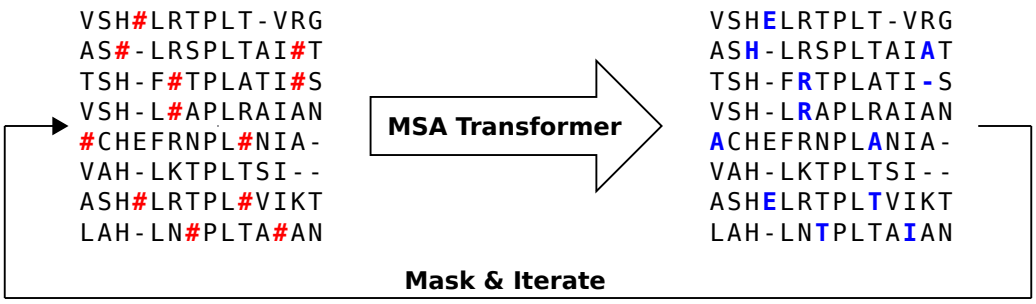

Figure 7: Iterative masking procedure to generate sequences using MSA Transformer. Here, the red symbol # stands for a masked amino acid, while blue uppercase letters (e.g. E) stand for predicted amino acids at the masked positions.

Choosing parameters in the iterative masking procedure. Fig. S10 illustrates, in the case of Pfam family PF00153, for different values of the masking probability p, how different properties of the generated MSAs evolve with the number I of iterations in the iterative masking procedure. For $p < 0.5$, we observe a gradual divergence from the initial natural sequences (Fig. S10A-B) and a simultaneous increase of scores (Fig. S10C-D, see "Scoring individual sequences" for definitions) and decrease of MSA diversity (Fig. S10E), and then a saturation of these various measures, as I increases. Our choice $I = 200$ is motivated by the fact that plateaus are reached at this point. However, the final values of all scores depend on p. Fig. S13 shows the contact maps inferred by MSA Transformer (using the logistic regression on tied row attentions trained in Ref. [25]) from generated sequences, for various values of I and p, in the case of family PF00153. We observe that the contact map characteristic of the protein family of interest gets gradually lost as I is increased for larger values of p (see Fig. S10F and Fig. S13)). These issues when p is increased are understandable, given that the pseudolikelihood loss used for the MLM objective in MSA Transformer ignores dependencies between masked entries. We note that despite this, larger values of p yield overall better average HMMER scores [73] and statistical energy scores (for $p < 0.5$). Our choice of $p = 0.1$ is motivated by the fact that this value is close to that employed in the training of the model ($p \approx 0.12$) [25], and that it better preserves contact maps. The product pI gives the average number of times that each amino acid of the MSA is changed during the generation process. With our choices, each amino acid is masked 20 times on average.

The behaviors observed in Fig. S10 for PF00153 are generic across the protein families we studied, as can be seen in Figs. S11 and S12, which show the same data as in Fig. S10 for Pfam families PF00096 and PF13354 (which have different sequence lengths). This demonstrates that our sequence generation method is robust. In particular, as the parameters $p = 0.1$ and $I = 200$ yield satisfactory convergence of MSA properties and preservation of contact maps in all cases, we used these parameters throughout, without any family-specific fine-tuning.

The sequences thus generated by our method do not coincide with natural ones. The fraction of MSA-Transformer–generated sequences which are identical to sequences in the input natural MSAs is below 5×10^{-4} for all large families considered, except three families with low diversity and/or very short sequence length (PF00096, PF00397 and PF00595).

Variants of the iterative masking procedure. In our algorithm, we mask tokens randomly throughout the input MSA. We also explored an alternative procedure where masking is restricted to the first sequence of the input MSA. Thus, all other sequences act as a context for the first sequence which is gradually modified. This can be done either with a fixed context, or by sampling different sequences from the natural MSA at each iteration to form a variable context. Note that the procedure with fixed context is reminiscent of the non-iterative one used in Ref. [74] to compute deep mutational scanning scores from MSA Transformer. For the same masking probability $p = 0.1$ as in our standard procedure (note that fewer iterations are needed for convergence, in practice $I = 20$ suffices), the alternative procedure with fixed context yields sequences that are overall slightly less different from natural ones than the standard iterative masking procedure, while the opposite holds with variable context. Besides, both alternative procedures yield sequences with better HMMER scores, but worse statistical energy scores, than natural ones – see Table S4. Finally, the two- and three-body statistics (defined in "Analyzing the statistics of MSAs") of the natural MSA are less well reproduced using these alternative procedures than the standard one – see Table S4. We also note that these variants are computationally more demanding. In this context, we decided to focus on the standard iterative masking procedure.

There are also different ways of selecting the token to fill each masked position. We have chosen a greedy sampling method where the token with highest probability is selected. We also explored an alternative method where the new token to fill the masked position is chosen by sampling the probability distribution given by the softmax of the logits, see Eq. (1). This method allows to introduce a sampling temperature T into the softmax operation and compute the probability as $p = \mathrm{softmax}(\boldsymbol{\xi}/T)$, where $\boldsymbol{\xi}$ is the logit vector. Note that the greedy method that we employ corresponds to sampling at $T = 0$. We found that MSAs generated with higher values of T are farther from the corresponding natural MSAs, showing that increasing this sampling temperature promotes originality. However, they are of lower quality according to our HMMER and statistical energy scores, and reproduce the statistics of the natural data less well. These results, summarized in Table S4, motivated us to mainly consider greedy sampling.

Finally, in our iterative masking procedure, we subsample the initial natural MSAs uniformly at random. We also tried diversity maximizing sampling [25], but we found that random sampling gives slightly better results.

Sampling sequences from Potts models

To sample independent equilibrium sequences from Potts models, we used the strategy described in Ref. [52]. Specifically, we fitted Potts models on each of our natural MSAs using bmDCA [12] (https://github.com/ranganathanlab/bmDCA). Using bmDCA is known to yield Potts models with good generative power [12, 13].

Consider a sequence of L amino-acid sites. We denote by $x_i \in \{1, \ldots, q\}$ the state of site $i \in \{1, \ldots, L\}$, where $q = 21$ is the number of possible states, namely the 20 natural amino acids and the alignment gap. The Potts model Hamiltonian of a sequence $\boldsymbol{x} = (x_1, \ldots, x_L)$ reads [8, 75]:

$$H(\boldsymbol{x}) = -\sum_{i=1}^{L} h_i(x_i) - \sum_{j=1}^{L} \sum_{i=1}^{j-1} e_{ij}(x_i, x_j)\,. \tag{2}$$

For each MSA $\mathcal{M}$ in Table 1, we inferred parameters $h_i(x_i)$ and $e_{ij}(x_i, x_j)$ by bmDCA [12, 13]. The Potts model probability distribution is then given by the Boltzmann distribution associated to the Hamiltonian H in Eq. (2):

$$P(\boldsymbol{x}) = \frac{e^{-H(\boldsymbol{x})/T}}{Z}\,, \tag{3}$$

where Z is a constant ensuring normalization and T is a parameter whose default value is 1. To generate a synthetic MSA from $\mathcal{M}$, we performed equilibrium Markov Chain Monte Carlo (MCMC) sampling from the Potts model with Hamiltonian H in Eq. (2). Specifically, we used the implementation in Ref. [52] of the Metropolis–Hastings algorithm, in which each step is a proposed mutation at a single amino-acid site. We started from a set of M randomly and independently initialized sequences, where M is the depth of $\mathcal{M}$, and made a total number N of Monte Carlo steps on each sequence. For each $\mathcal{M}$, suitable values for N are estimated by bmDCA during its training, to ensure that Metropolis–Hastings sampling reaches thermal equilibrium after N steps when starting from a randomly initialized sequence [12]. We thus used the value of N estimated by bmDCA at the end of training. This yielded, for each MSA in Table 1, a synthetic MSA of the same depth, composed of independent equilibrium sequences.

This procedure allows to tune the sampling temperature T, in a similar spirit as for MSA Transformer, c.f. "Variants of the iterative masking procedure". This amounts to tuning the selection strength. Recall that Potts models are inferred at $T = 1$, which is thus the default value. Using MCMC sampling as described above, we first generated synthetic MSAs at $T = 1$, and using regularization strength $\lambda = 10^{-2}$. These correspond to the default parameters (λ, T), matching those employed in [12], and allowing direct comparison with those results. Importantly, using sampling temperature $T = 1$ means that the distribution learnt from natural data is directly sampled. However, it was found in [13] that sequences generated at $T = 1$ have worse statistical energy scores than natural sequences, due at least in part to high regularization, and that this can be corrected by lower-temperature sampling. Therefore, for completeness, we also considered all parameter combinations (λ, T) used in [13] for PF00072. Table S5 shows that decreasing sampling temperature strongly improves the mean statistical energy score, as it should, and somewhat improves HMMER scores and structural scores. However, this comes at the cost of decreasing MSA diversity and getting sequences substantially more similar to natural ones. It also strongly impairs the fitting of the one- and two-body statistics. The effect of changing regularization strength (at inference) appears to be more minor, but decreasing it allows to somewhat mitigate the loss of diversity associated to lowering temperature. In light of these results, and to make our comparison to bmDCA comprehensive, we used $(\lambda, T) = (10^{-3}, 0.33)$ [13] in addition to $(\lambda, T) = (10^{-2}, 1)$ [12] throughout our analysis of deep MSAs. In the case of shallow MSAs (Fig. 3), we employed $(\lambda, T) = (10^{-2}, 0.33)$ instead of $(\lambda, T) = (10^{-3}, 0.33)$ because shallow MSAs require stronger regularization strengths.

Scoring individual sequences

We use different scores to compare natural and generated sequences.

First, HMMER scores [73] are computed, for each sequence, from the Pfam profile Hidden Markov Models (HMM), employing the function `hmmsearch` from the HMMER Suite version 3.3.2 (`http://hmmer.org`). HMMER scores are homology scores, which are in particular used in Pfam to search sequence databases for sequence homologs and to construct full MSAs starting from curated seed MSAs. Higher HMMER scores indicate better matches to the Pfam HMM.

Second, DCA statistical energy scores are computed for each sequence using the Potts model Hamiltonian H in Eq. (2) with the couplings and the fields inferred by bmDCA on the natural MSA of the family of interest (see "Sampling sequences from Potts models"). The statistical

energy score is then defined as the opposite of the statistical energy, i.e. $-H(\boldsymbol{x})$ for a sequence $\boldsymbol{x}$, so that, here too, higher values mean better scores.

We also compute AlphaFold [28] structural prediction confidence scores, i.e. predicted local-distance difference test (pLDDT) values. Given the computational cost, for each natural or generated MSA, we evaluate pLDDT values for a subset of 200 randomly sampled sequences.

Finally, we compute the root-mean-square deviation (RMSD) between a reference experimental structure of the family of focus (see list in Table 1) and the AlphaFold predicted structures, also for a subset of 200 randomly sampled sequences in each MSA.

Because AlphaFold takes MSAs as input, we compute these two structural scores using the whole natural MSA of the family of interest as context in all cases. In addition, for the protein family PF00072, we also used fully synthetic MSAs as input to AlphaFold. Structural scores are then very similar to those obtained using natural context (see Table S5).

Analyzing the statistics of MSAs

To compare the generated MSAs to the natural ones, we consider different statistical measures.

First, to analyze how faithfully the generated MSAs reproduce the statistics of the natural ones at various orders, we compute the r20 score [15, 50]. Specifically, to obtain Fig. 4, we analyze the frequency of sub-sequences spanning 2 to 10 non-contiguous columns. In each of 1000 randomly sampled sets of columns for each sub-sequence length, we compute the frequency of the 20 most frequent words in natural and synthetic MSAs of the family considered, and evaluate the Pearson correlation between these top 20 frequencies in the MSA of focus and those in a reference MSA. We then average these Pearson correlation values over all sets of 1000 columns, yielding the r20 score.

To further inspect low-order statistics, in each MSA, we compute the one-body frequencies of occurrence of each amino acid at each site, the two-body frequencies of each pair of amino acids at each pair of sites, and the three-body frequencies associated to triplets. We denote them by $f_i(x)$, $f_{ij}(x,y)$, $f_{ijk}(x,y,z)$, where i, j and k denote sites, while x, y and z represent amino acids (see "Sampling sequences from Potts models"). We then estimate the second and third order connected correlations as:

$$C_{ij}(x,y) = f_{ij}(x,y) - f_i(x)f_j(y)\,; \tag{4}$$

$$\begin{aligned} C_{ijk}(x,y,z) = {} & f_{ijk}(x,y,z) - f_{ij}(x,y)f_k(z) - f_{ik}(x,z)f_j(y) - f_{jk}(y,z)f_i(x) \\ & + 2f_i(x)f_j(y)f_k(z)\,. \end{aligned} \tag{5}$$

We also compute the "plug-in" estimates of the Shannon entropy of each site H_i, and of the two- and three-body joint entropies H_{ij} and H_{ijk}, from the frequencies. They yield the plug-in estimates of the mutual information I_{ij} between two columns, and of the *co-information* I_{ijk} between three columns:

$$I_{ij} = H_i + H_j - H_{ij}\,; \tag{6}$$

$$I_{ijk} = H_i + H_j + H_k - H_{ij} - H_{ik} - H_{jk} + H_{ijk}\,. \tag{7}$$

Co-information is a measure of higher-order statistical dependencies [76–79], which generalizes mutual information to triplets of random variables, vanishes for independent variables, and reflects the balance between redundancy and synergy in these triplets [80, 81]. A systematic finite-size error occurs when estimating entropies using the plug-in estimate from frequencies measured in finite datasets [2], and it affects entropy-derived quantities such as mutual information and co-information. Here, we do not attempt to correct it. Rather, we only make comparisons between MSAs of the same length and depth, which are affected by the same finite-size errors.

Characterizing the distribution of sequences in MSAs

Another way of studying the properties of generated MSAs is to analyze the distribution of their sequences in sequence space, and to compare it to that of natural sequences in the same family.

First, to assess whether generated sequences most resemble natural ones that are well-represented in their family or, rather, rare ones, we consider for each synthetic sequence its closest natural sequence. We then count the number of neighbors of this natural sequence in the natural MSA, i.e. the number of natural sequences that have (normalized) Hamming distance below $\delta = 0.2$ with the sequence of interest.[2]

Second, to explore the distributions in sequence space of sequences within each MSA, and compare synthetic and natural MSAs, we associate to each sequence the concatenation of the one-hot encodings of each of its amino acids [12]. We perform a principal component analysis of the matrix corresponding to the natural MSA in this representation. We can then represent natural and synthetic sequences as points projected in the space defined by the first two principal components of the natural MSA.

Third, to analyze in more detail the apparent relatedness of generated sequences, and compare it to real phylogenetic relationships in natural sequences, we infer phylogenetic trees from each MSA using `FastTree 2` [82]. To quantitatively compare the topologies of these trees, which do not have the same leaves, we analyze the eigenvalue spectrum of their modified graph Laplacian (MGL) [60]. The MGL of a phylogenetic tree is defined as the difference between its degree matrix (a diagonal matrix whose i-th diagonal entry is the sum of the branch lengths from node i to all other nodes in the tree) and the matrix of patristic distances (whose (i, j)-th entry is the branch length between nodes i and j). Given the computational cost of running such an analysis on our deep MSAs, we use a bootstrap-aggregating strategy in the spirit of Ref. [83]. Namely, for each MSA we compute 200 different trees, each one inferred from a different sub-MSA of 500 sequences, itself randomly sampled from the whole MSA. Then, for each of these trees, we compute the eigenvalue spectrum of the MGL. Next, we merge all these spectra together to obtain a single eigenvalue spectral density. Note that this method has the advantage of not depending on the details of the topology of one large inferred tree, which are known to be sensitive to the choice of phylogeny reconstruction algorithm.

Datasets

To generate synthetic MSAs with MSA Transformer and bmDCA and compare them to their natural counterparts, we consider the deep Pfam "full" alignments [84] associated to 14 different protein domains (Table 1). Each MSA is a matrix $\mathcal{M}$ with L columns, representing the different amino-acid sites, and M rows. Each row i, denoted by $\boldsymbol{x}^{(i)}$, represents one sequence of the alignment. We refer to L as the MSA length, and to M as its depth. For all our MSAs, $M > 36000$. These alignments are the same as in Ref. [52], except that we removed PF13354 (Beta-lactamase2) from this set of deep MSAs because of its smaller depth. However, this family is included in our additional analyses (see Table S1).

Deep MSAs generally include some highly similar sequences due to phylogenetic relatedness. This can be characterized via the effective depth [8]

$$M_{\text{eff}}^{(\delta)} := \sum_{i=1}^{M} w_i, \quad \text{with} \quad w_i := |\{i' : d_{\text{H}}(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(i')}) < \delta\}|^{-1}, \tag{8}$$

where $d_{\text{H}}(\boldsymbol{x}, \boldsymbol{y})$ is the (normalized) Hamming distance between two sequences $\boldsymbol{x}$ and $\boldsymbol{y}$, i.e. the fraction of sites where the amino acids differ, and we set $\delta = 0.2$. Note that the inverse of the

[2]Note that the inverse of this number of neighbors gives the sequence weight w_i introduced in Eq. (8).

Pfam ID	Family name	L	M	$M_{\text{eff}}^{(0.2)}$	PDB ID	Resol.
PF00004	AAA	132	39277	9049	4D81	2.40 Å
PF00005	ABC_tran	137	68891	43881	1L7V	3.20 Å
PF00041	fn3	85	42721	17782	3UP1	2.15 Å
PF00072	Response_reg	112	73063	40180	3ILH	2.59 Å
PF00076	RRM_1	69	51964	20273	3NNH	2.75 Å
PF00096	zf-C2H2	23	38996	12581	4R2A	1.59 Å
PF00153	Mito_carr	94	93776	17859	1OCK	2.20 Å
PF00271	Helicase_C	111	66809	25017	3EX7	2.30 Å
PF00397	WW	31	39045	3361	4REX	1.60 Å
PF00512	HisKA	66	154998	67303	3DGE	2.80 Å
PF00595	PDZ	82	71303	4053	1BE9	1.82 Å
PF01535	PPR	31	109064	37514	4M57	2.86 Å
PF02518	HATPase_c	111	80714	59189	3G7E	2.20 Å
PF07679	I-set	90	36141	14611	1FHG	2.00 Å

Table 1: **Pfam families and natural MSAs used in our analysis.** L denotes the length of an MSA, M its depth, and $M_{\text{eff}}^{(0.2)}$ its effective depth with distance threshold $\delta = 0.2$, see Eq. (8). The reference experimental PDB structures used for our RMSD calculations, and their resolutions ("Resol."), are also reported.

sequence weight w_i in Eq. (8) is the number of neighbors in "Characterizing the distribution of sequences in MSAs", and that $M_{\text{eff}}^{(0.2)}/M$ can be as low as 0.06 for our natural MSAs.

All these families were previously shown to be well fitted by Potts models inferred by bmDCA [12], making our results on sequence generation by bmDCA readily comparable with previous results. Our domains' short lengths are convenient because bmDCA is computationally demanding, and also in view of MSA Transformer's large memory footprint, which is $O(LM^2)+O(L^2)$. Furthermore, their large depth is crucial to our comparisons, as it allows Potts model to be accurately fitted [12].

We extended our study to small protein families by considering 7 additional families, listed in Table S1, for which we also started from Pfam "full" MSAs. These families comprise from 9 to a few hundreds of sequences. We also considered two additional protein families, also listed in Table S1, for our comparison with published experimental datasets.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 851173, to A.-F. B.).

Data availability statement

Python code for generating sequences using the iterative masking procedure is available in our GitHub repository: https://github.com/Bitbol-Lab/Iterative_masking.

Raw data was collected from two public sources: 1) MSAs from the Pfam database (https://pfam.xfam.org/); 2) further MSAs from https://github.com/matteofigliuzzi/bmDCA. We generated sequences with bmDCA using code publicly available at https://github.com/ranganathanlab/bmDCA.

References

[1] Michael Socolich, Steve W. Lockless, William P. Russ, Heather Lee, Kevin H. Gardner, and Rama Ranganathan. Evolutionary information for specifying a protein fold. *Nature*, 437(7058):512–518, 2005.

[2] William Bialek. *Biophysics: Searching for Principles*. Princeton University Press, 2012.

[3] B. I. Dahiyat and S. L. Mayo. De novo protein design: fully automated sequence selection. *Science*, 278(5335):82–87, Oct 1997.

[4] B. Kuhlman, G. Dantas, G. C. Ireton, G. Varani, B. L. Stoddard, and D. Baker. Design of a novel globular protein fold with atomic-level accuracy. *Science*, 302(5649):1364–1368, Nov 2003.

[5] Huanhuan Liang, Hao Chen, Keqiang Fan, Ping Wei, Xianrong Guo, Changwen Jin, Chen Zeng, Chao Tang, and Luhua Lai. De novo design of a $\beta\alpha\beta$ motif. *Angewandte Chemie International Edition*, 48(18):3301–3303, 2009.

[6] G. J. Rocklin, T. M. Chidyausiku, I. Goreshnik, A. Ford, S. Houliston, A. Lemak, L. Carter, R. Ravichandran, V. K. Mulligan, A. Chevalier, C. H. Arrowsmith, and D. Baker. Global analysis of protein folding using massively parallel design, synthesis, and testing. *Science*, 357(6347):168–175, 07 2017.

[7] Frances H. Arnold. Directed evolution: Bringing new chemistry to life. *Angewandte Chemie International Edition*, 57(16):4106–4106, 2018.

[8] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa. Identification of direct residue contacts in protein-protein interaction by message passing. *Proc. Natl. Acad. Sci. U.S.A.*, 106(1):67–72, Jan 2009.

[9] Faruck Morcos, Andrea Pagnani, Bryan Lunt, Arianna Bertolino, Debora S. Marks, Chris Sander, Riccardo Zecchina, José N. Onuchic, Terence Hwa, and Martin Weigt. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proc. Natl. Acad. Sci. U.S.A.*, 108(49):E1293–E1301, 2011.

[10] D. S. Marks, L. J. Colwell, R. Sheridan, T. A. Hopf, A. Pagnani, R. Zecchina, and C. Sander. Protein 3D structure computed from evolutionary sequence variation. *PLoS ONE*, 6(12):1–20, 2011.

[11] Magnus Ekeberg, Cecilia Lövkvist, Yueheng Lan, Martin Weigt, and Erik Aurell. Improved contact prediction in proteins: Using pseudolikelihoods to infer Potts models. *Phys. Rev. E*, 87(1):1–16, 2013.

[12] Matteo Figliuzzi, Pierre Barrat-Charlaix, and Martin Weigt. How pairwise coevolutionary models capture the collective residue variability in proteins? *Mol. Biol. Evol.*, 35(4):1018–1027, 2018.

[13] William P. Russ, Matteo Figliuzzi, Christian Stocker, Pierre Barrat-Charlaix, Michael Socolich, Peter Kast, Donald Hilvert, Remi Monasson, Simona Cocco, Martin Weigt, and Rama Ranganathan. An evolution-based model for designing chorismate mutase enzymes. *Science*, 369(6502):440–445, 2020.

[14] A. Hawkins-Hooker, F. Depardieu, S. Baur, G. Couairon, A. Chen, and D. Bikard. Generating functional protein variants with variational autoencoders. *PLoS Comput. Biol.*, 17(2):1–23, 02 2021.

[15] Francisco McGee, Sandro Hauri, Quentin Novinger, Slobodan Vucetic, Ronald M. Levy, Vincenzo Carnevale, and Allan Haldane. The generative capacity of probabilistic protein sequence models. *Nature Commun.*, 12(1):6302, 11 2021.

[16] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. In *International Conference on Learning Representations*, 2015.

[17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30:5998–6008, 2017.

[18] E. C. Alley, G. Khimulya, S. Biswas, M. AlQuraishi, and G. M. Church. Unified rational protein engineering with sequence-based deep representation learning. *Nature Methods*, 16(12):1315–1322, 12 2019.

[19] Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, Debsindhu Bhowmik, and Burkhard Rost. ProtTrans: Towards cracking the language of life's code through self-supervised deep learning and high performance computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2021.

[20] Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C. Lawrence Zitnick, Jerry Ma, and Rob Fergus. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proc. Natl. Acad. Sci. U.S.A.*, 118(15), 2021.

[21] Roshan Rao, Joshua Meier, Tom Sercu, Sergey Ovchinnikov, and Alexander Rives. Transformer protein language models are unsupervised structure learners. In *International Conference on Learning Representations*, 2021.

[22] Jesse Vig, Ali Madani, Lav R. Varshney, Caiming Xiong, Richard Socher, and Nazneen Rajani. BERTology meets biology: Interpreting attention in protein language models. In *International Conference on Learning Representations*, 2021.

[23] Ali Madani, Bryan McCann, Nikhil Naik, Nitish Shirish Keskar, Namrata Anand, Raphael R. Eguchi, Po-Ssu Huang, and Richard Socher. ProGen: Language Modeling for Protein Generation. *bioRxiv*, 10.1101/2020.03.07.982272, 2020.

[24] Ali Madani, Ben Krause, Eric R. Greene, Subu Subramanian, Benjamin P. Mohr, James M. Holton, Jose Luis Olmos, Caiming Xiong, Zachary Z. Sun, Richard Socher, James S. Fraser, and Nikhil Naik. Deep neural language modeling enables functional protein generation across families. *bioRxiv*, 10.1101/2021.07.18.452833, 2021.

[25] Roshan M Rao, Jason Liu, Robert Verkuil, Joshua Meier, John Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. MSA Transformer. *Proceedings of the 38th International Conference on Machine Learning*, 139:8844–8856, 18–24 Jul 2021.

[26] Nicholas Bhattacharya, Neil Thomas, Roshan Rao, Justas Dauparas, Peter K. Koo, David Baker, Yun S. Song, and Sergey Ovchinnikov. Single layers of attention suffice to predict protein contacts. *bioRxiv*, 10.1101/2020.12.21.423882, 2020.

[27] Nicholas Bhattacharya, Neil Thomas, Roshan Rao, Justas Dauparas, Peter K Koo, David Baker, Yun S Song, and Sergey Ovchinnikov. Interpreting Potts and Transformer Protein Models Through the Lens of Simplified Attention. *Pacific Symposium on Biocomputing.*, pages 34–45, 2022.

[28] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A A Kohl, Andrew J Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596:583–589, 2021.

[29] Sean R Johnson, Kenneth Massie, Sarah Monaco, and Zaid Syed. Generating novel protein sequences using Gibbs sampling of masked language models. *bioRxiv*, 2021.

[30] Alex Hawkins-Hooker, David T. Jones, and Brooks Paige. MSA-conditioned generative protein language models for fitness landscape modelling and design. In *Machine Learning for Structural Biology Workshop, NeurIPS*, 2021.

[31] Noelia Ferruz, Steffen Schmidt, and Birte Höcker. A deep unsupervised language model for protein design. *bioRxiv*, 10.1101/2022.03.09.483666, 2022.

[32] Brian L Hie, Duo Xu, Varun R Shanker, Theodora UJ Bruun, Payton A Weidenbacher, Shaogeng Tang, and Peter S Kim. Efficient evolution of human antibodies from general protein language models and sequence information alone. *bioRxiv*, 10.1101/2022.04.10.487811, 2022.

[33] Georg Casari, Chris Sander, and Alfonso Valencia. A method to predict functional residues in proteins. *Nature*, 2(2):171–178, 1995.

[34] A. S. Lapedes, B. G. Giraud, L. Liu, and G. D. Stormo. Correlated mutations in models of protein sequences: phylogenetic and structural effects. In *Statistics in molecular biology and genetics - IMS Lecture Notes - Monograph Series*, volume 33, pages 236–256. 1999.

[35] S. D. Dunn, L. M. Wahl, and G. B. Gloor. Mutual information without the influence of phylogeny or entropy dramatically improves residue contact prediction. *Bioinformatics*, 24(3):333–340, Feb 2008.

[36] W. Bialek and R. Ranganathan. Rediscovering the power of pairwise interactions. *ArXiv*, 10.48550/arxiv.0712.4397, 2007.

[37] J. I. Sułkowska, F. Morcos, M. Weigt, T. Hwa, and J. N. Onuchic. Genomics-aided structure prediction. *Proc. Natl. Acad. Sci. U.S.A.*, 109(26):10340–10345, Jun 2012.

[38] R. S. Dwyer, D. P. Ricci, L. J. Colwell, T. J. Silhavy, and N. S. Wingreen. Predicting functionally informative mutations in Escherichia coli BamA using evolutionary covariance analysis. *Genetics*, 195(2):443–455, Oct 2013.

[39] R. R. Cheng, F. Morcos, H. Levine, and J. N. Onuchic. Toward rationally redesigning bacterial two-component signaling systems using coevolutionary information. *Proc. Natl. Acad. Sci. U.S.A.*, 111(5):E563–571, Feb 2014.

[40] R. R. Cheng, O. Nordesjo, R. L. Hayes, H. Levine, S. C. Flores, J. N. Onuchic, and F. Morcos. Connecting the Sequence-Space of Bacterial Signaling Proteins to Phenotypes Using Coevolutionary Landscapes. *Mol. Biol. Evol.*, 33(12):3054–3064, 12 2016.

[41] M. Figliuzzi, H. Jacquier, A. Schug, O. Tenaillon, and M. Weigt. Coevolutionary Landscape Inference and the Context-Dependence of Mutations in Beta-Lactamase TEM-1. *Mol. Biol. Evol.*, 33(1):268–280, Jan 2016.

[42] J. A. de la Paz, C. M. Nartey, M. Yuvaraj, and F. Morcos. Epistatic contributions promote the unification of incompatible models of neutral molecular evolution. *Proc. Natl. Acad. Sci. U.S.A.*, 117(11):5873–5882, 03 2020.

[43] F. Morcos, B. Jana, T. Hwa, and J. N. Onuchic. Coevolutionary signals across protein lineages help capture multiple protein conformations. *Proc. Natl. Acad. Sci. U.S.A.*, 110(51):20533–20538, Dec 2013.

[44] D. Malinverni, S. Marsili, A. Barducci, and P. De Los Rios. Large-Scale Conformational Transitions and Dimerization Are Encoded in the Amino-Acid Sequences of Hsp70 Chaperones. *PLoS Comput. Biol.*, 11(6):1–15, Jun 2015.

[45] Anne-Florence Bitbol, Robert S Dwyer, Lucy J Colwell, and Ned S Wingreen. Inferring interaction partners from protein sequences. *Proc. Natl. Acad. Sci. U.S.A.*, 113(43):12180–12185, 2016.

[46] T. Gueudre, C. Baldassi, M. Zamparo, M. Weigt, and A. Pagnani. Simultaneous identification of specifically interacting paralogs and interprotein contacts by direct coupling analysis. *Proc. Natl. Acad. Sci. U.S.A.*, 113(43):12186–12191, 10 2016.

[47] Qian Cong, Ivan Anishchenko, Sergey Ovchinnikov, and David Baker. Protein interaction networks revealed by proteome coevolution. *Science*, 365(6449):185–189, 2019.

[48] Anna G. Green, Hadeer Elhabashy, Kelly P. Brock, Rohan Maddamsetti, Oliver Kohlbacher, and Debora S. Marks. Large-scale discovery of protein interactions at residue resolution using co-evolution calculated from genomic sequences. *Nature Commun.*, 12(1):1–12, 2021.

[49] https://www.ncbi.nlm.nih.gov/Class/FieldGuide/BLOSUM62.txt.

[50] Allan Haldane, William F. Flynn, Peng He, and Ronald M. Levy. Coevolutionary Landscape of Kinase Family Proteins: Sequence Probabilities and Functional Motifs. *Biophysical Journal*, 114(1):21–31, 2018.

[51] Leenoy Meshulam, Jeffrey L. Gauthier, Carlos D. Brody, David W. Tank, and William Bialek. Successes and failures of simplified models for a network of real neurons. *arXiv*, 10.48550/arxiv.2112.14735, 2021.

[52] Umberto Lupo, Damiano Sgarbossa, and Anne-Florence Bitbol. Protein language models trained on multiple sequence alignments learn phylogenetic relationships. *Nature Commun.*, 13(6298), 2022.

[53] C. Qin and L. J. Colwell. Power law tails in phylogenetic systems. *Proc. Natl. Acad. Sci. U.S.A.*, 115(4):690–695, Jan 2018.

[54] Susann Vorberg, Stefan Seemayer, and Johannes Söding. Synthetic protein alignments by ccmgen quantify noise in residue-residue contact prediction. *PLoS Comput. Biol.*, 14(11):1–25, 11 2018.

[55] Edwin Rodriguez Horta, Pierre Barrat-Charlaix, and Martin Weigt. Toward inferring potts models for phylogenetically correlated sequence data. *Entropy*, 21(11), 2019.

[56] Edwin Rodriguez Horta and Martin Weigt. On the effect of phylogenetic correlations in coevolution-based contact prediction in proteins. *PLoS Comput. Biol.*, 17(5), 2021.

[57] G. Marmier, M. Weigt, and A.-F. Bitbol. Phylogenetic correlations can suffice to infer protein partners from sequences. *PLoS Comput. Biol.*, 15(10):e1007179, Oct 2019.

[58] A. Colavin, E. Atolia, A.-F. Bitbol, and K. C. Huang. Extracting phylogenetic dimensions of coevolution reveals hidden functional signals. *Scientific Reports*, 12:820, 2022.

[59] Andonis Gerardos, Nicola Dietler, and Anne-Florence Bitbol. Correlations from structure and phylogeny combine constructively in the inference of protein partners from sequences. *bioRxiv*, 10.1101/2021.11.22.469574, 2021.

[60] Eric Lewitus and Helene Morlon. Characterizing and comparing phylogenies from their laplacian spectrum. *Systematic Biology*, 65(3):495–507, 2016.

[61] Rafael Molina-Venegas. What are "tippy" and "stemmy" phylogenies? Resolving a phylogenetic terminological tangle. *Journal of Systematics and Evolution*, 59(2):403–404, mar 2021.

[62] C. Malbranke, D. Bikard, S. Cocco, and R. Monasson. Improving sequence-based modeling of protein families using secondary structure quality assessment. *Bioinformatics*, 37(22):4083–4090, Jun 2021.

[63] Richard N McLaughlin Jr, Frank J Poelwijk, Arjun Raman, Walraj S Gosal, and Rama Ranganathan. The spatial architecture of protein function and adaptation. *Nature*, 491(7422):138–142, 2012.

[64] Michael A. Stiffler, Doeke R. Hekstra, and Rama Ranganathan. Evolvability as a function of purifying selection in TEM-1 β-lactamase. *Cell*, 160(5):882–892, 2015.

[65] Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N. Kinch, R. Dustin Schaeffer, Claudia Millán, Hahnbeom Park, Carson Adams, Caleb R. Glassman, Andy DeGiovanni, Jose H. Pereira, Andria V. Rodrigues, Alberdina A. van Dijk, Ana C. Ebrecht, Diederik J. Opperman, Theo Sagmeister, Christoph Buhlheller, Tea Pavkov-Keller, Manoj K. Rathinaswamy, Udit Dalwadi, Calvin K. Yip, John E. Burke, K. Christopher Garcia, Nick V. Grishin, Paul D. Adams, Randy J. Read, and David Baker. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021.

[66] Ratul Chowdhury, Nazim Bouatta, Surojit Biswas, Charlotte Rochereau, George M. Church, Peter K. Sorger, and Mohammed AlQuraishi. Single-sequence protein structure prediction using language models from deep learning. *bioRxiv*, 10.1101/2021.08.02.454840, 2021.

[67] I. Anishchenko, S. J. Pellock, T. M. Chidyausiku, T. A. Ramelot, S. Ovchinnikov, J. Hao, K. Bafna, C. Norn, A. Kang, A. K. Bera, F. DiMaio, L. Carter, C. M. Chow, G. T. Montelione, and D. Baker. De novo protein design by deep network hallucination. *Nature*, 600:547–552, Dec 2021.

[68] Kartik Goyal, Chris Dyer, and Taylor Berg-Kirkpatrick. Exposing the implicit energy networks behind masked language models via Metropolis–Hastings. *arXiv*, 10.48550/arxiv.2106.02736, 2021.

[69] Jérôme Tubiana, Simona Cocco, and Rémi Monasson. Learning protein constitutive motifs from sequence data. *eLife*, 8, 2019.

[70] Maxwell L. Bileschi, David Belanger, Drew Bryant, Theo Sanderson, Brandon Carter, D. Sculley, Mark A. DePristo, and Lucy J. Colwell. Using deep learning to annotate the protein universe. *Nature Biotechnol.*, pages 1–6, 2022.

[71] J. E. Shin, A. J. Riesselman, A. W. Kollasch, C. McMahon, E. Simon, C. Sander, A. Manglik, A. C. Kruse, and D. S. Marks. Protein design and variant prediction using autoregressive generative models. *Nature Commun.*, 12(1):2403, 04 2021.

[72] Alex Wang and Kyunghyun Cho. BERT has a mouth, and it must speak: BERT as a Markov random field language model. *CoRR*, abs/1902.04094, 2019.

[73] Sean R. Eddy. Profile hidden Markov models. *Bioinformatics*, 14(9):755–763, 1998.

[74] Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alex Rives. Language models enable zero-shot prediction of the effects of mutations on protein function. *Advances in Neural Information Processing Systems*, 34:29287–29303, 2021.

[75] S. Cocco, C. Feinauer, M. Figliuzzi, R. Monasson, and M. Weigt. Inverse statistical physics of protein sequences: a key issues review. *Rep Prog Phys*, 81(3):032601, 03 2018.

[76] William J. McGill. Multivariate information transmission. *Psychometrika*, 19(2):97–116, jun 1954.

[77] N. Timme, W. Alford, B. Flecker, and J. M. Beggs. Synergy, redundancy, and multivariate information measures: an experimentalist's perspective. *J Comput Neurosci*, 36(2):119–140, Apr 2014.

[78] R. Quax, O. Har-Shemesh, and P. M. A. Sloot. Quantifying Synergistic Information Using Intermediate Stochastic Variables. *Entropy*, 19:85, 2017.

[79] F. E. Rosas, P. A. M. Mediano, M. Gastpar, and H. J. Jensen. Quantifying high-order interdependencies via multivariate extensions of the mutual information. *Phys. Rev. E*, 100(3-1):032305, Sep 2019.

[80] Paul L. Williams and Randall D. Beer. Nonnegative decomposition of multivariate information. *arXiv*, 10.48550/arxiv.1004.2515, 2010.

[81] Fernando Rosas, Vasilis Ntranos, Christopher J. Ellison, Sofie Pollin, and Marian Verhelst. Understanding interdependency through complex information sharing. *Entropy*, 18(2), 2016.

[82] Morgan N. Price, Paramvir S. Dehal, and Adam P. Arkin. Fasttree 2 – approximately maximum-likelihood trees for large alignments. *PLoS ONE*, 5(3):1–10, 03 2010.

[83] C. Colijn and G. Plazzotta. A Metric on Phylogenetic Tree Shapes. *Systematic Biology*, 67(1):113–126, 2018.

[84] Jaina Mistry, Sara Chuguransky, Lowri Williams, Matloob Qureshi, Gustavo A Salazar, Erik L L Sonnhammer, Silvio C E Tosatto, Lisanna Paladin, Shriya Raj, Lorna J Richardson, Robert D Finn, and Alex Bateman. Pfam: The protein families database in 2021. *Nucleic Acids Research*, 49(D1):D412–D419, 10 2020.

Supplementary material

Supplementary figures

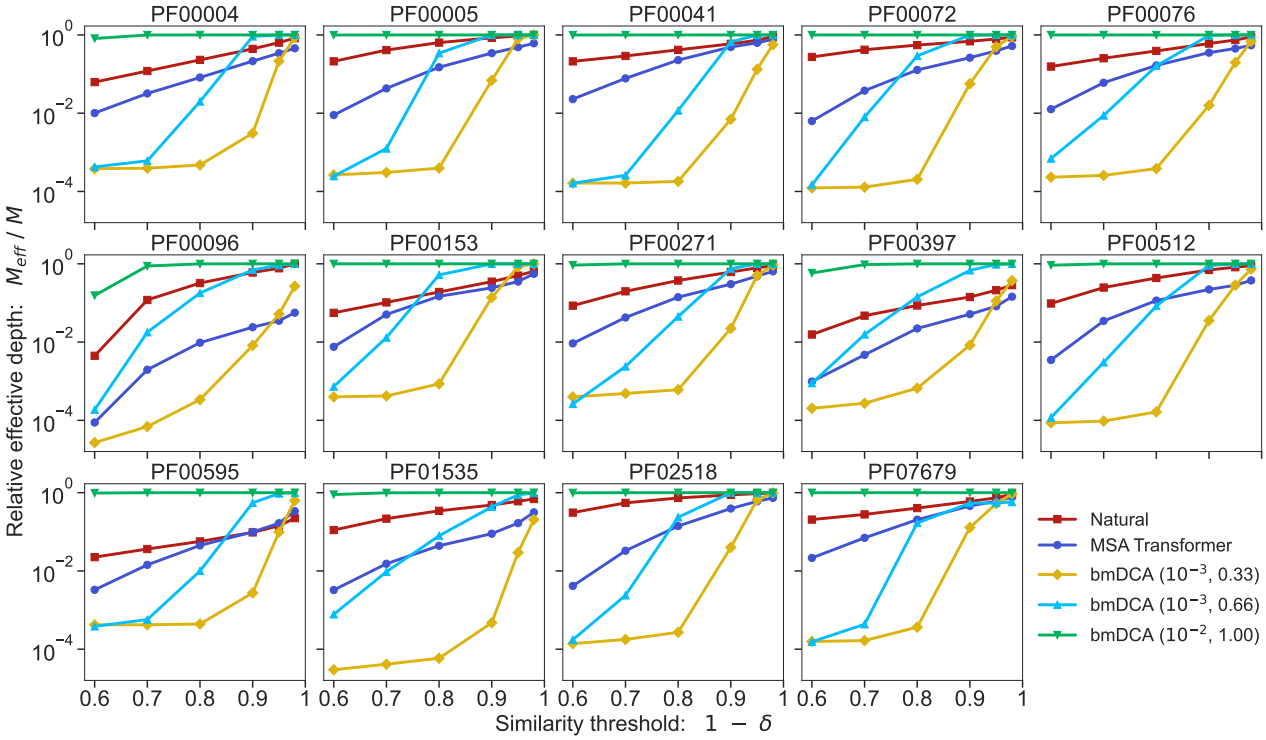

Figure S1: **MSA diversity for each protein family and each generation method.** We show the relative effective depth of each MSA (natural or generated), i.e. M_{eff}/M, where M_{eff} is effective MSA depth and M is actual MSA depth, versus the similarity threshold $1 - \delta$ used to define the effective depth (see Eq. (8)).

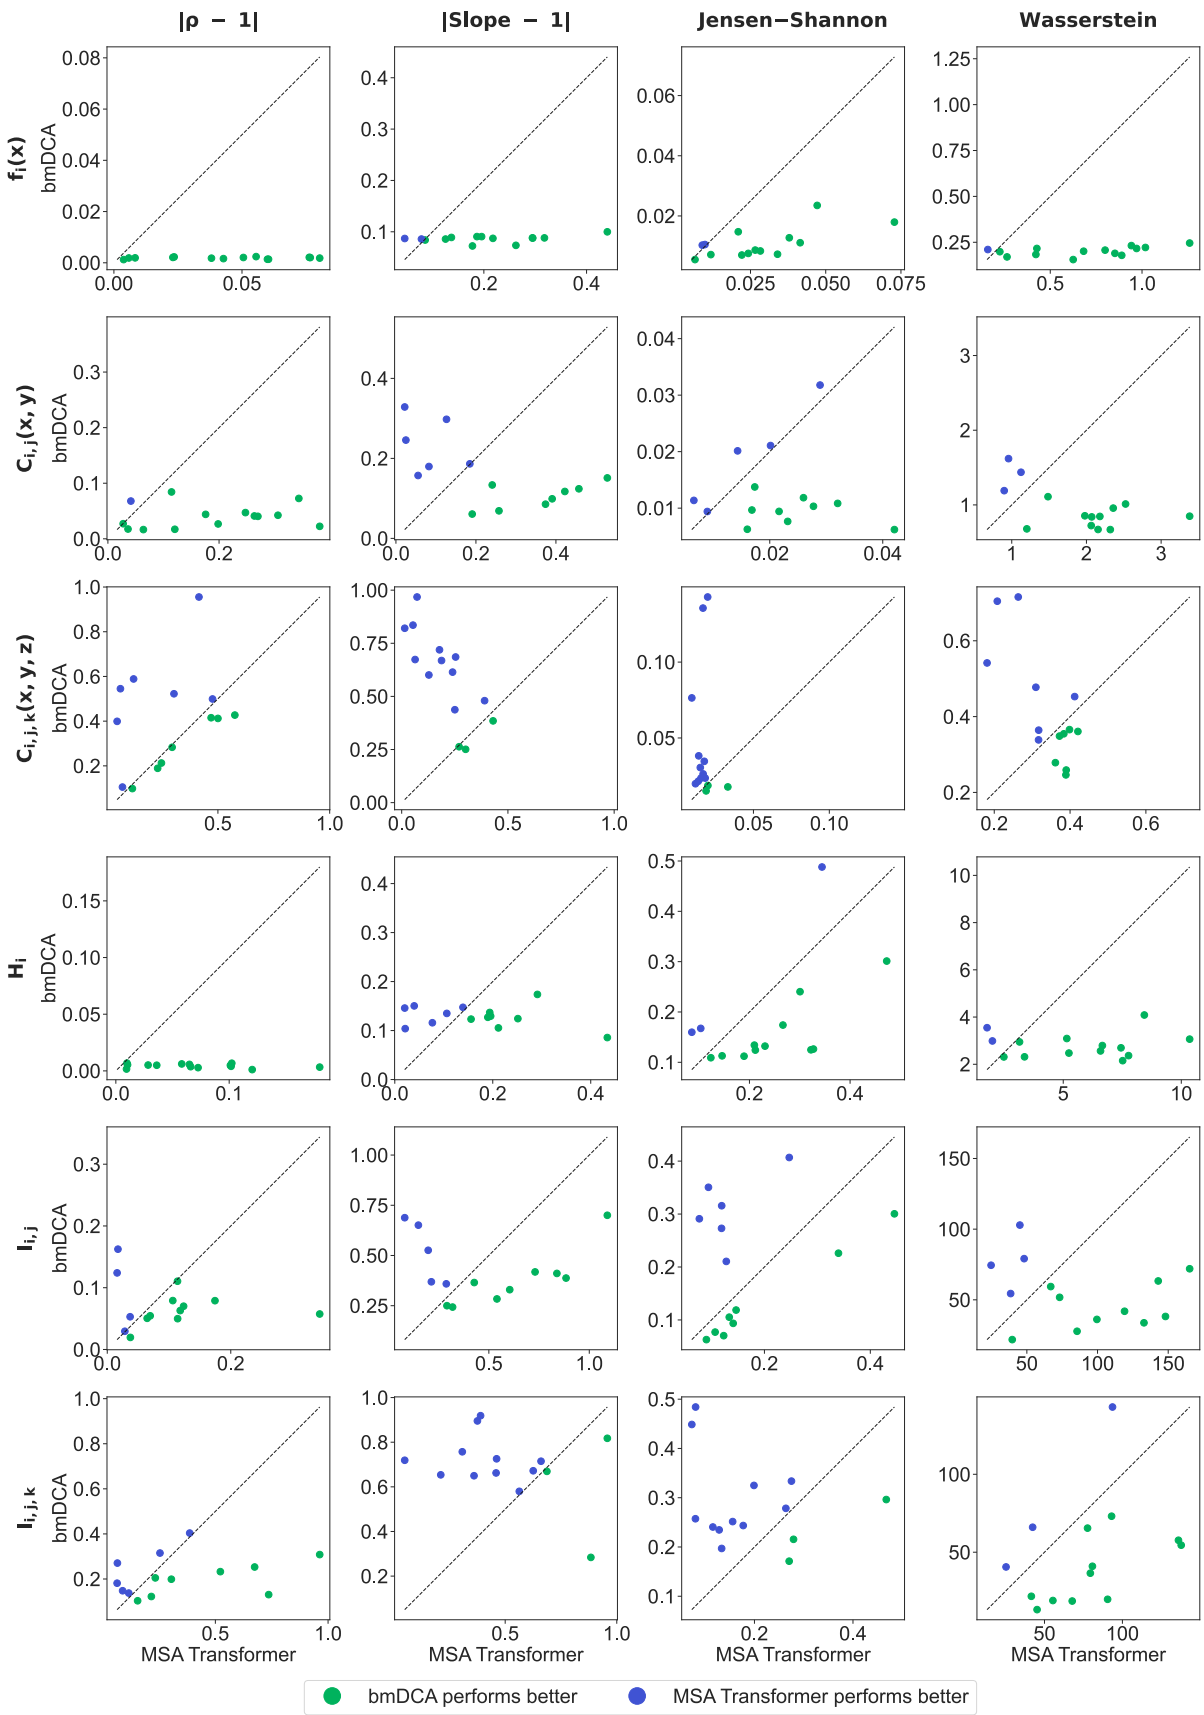

Figure S2: **Ability of generated sequences to reproduce one-, two- and three-body statistics (details on next page).** In all panels, each marker represents a family in Table 1. For each statistical or information measure (rows), different scores (columns) comparing generated and natural sequences are the coordinates of these points. All scores are such that being close to 0 is better, and their value for MSAs generated by bmDCA with default parameters is shown versus that for MSAs generated by MSA Transformer.

Figure S2: **Ability of generated sequences to reproduce one-, two- and three-body statistics (continued).** The statistical or information measures considered in each row are defined in "Analyzing the statistics of MSAs" – from the top: one-body frequency, two- and three-body connected correlations, entropy, mutual information and co-information. For each of them, we consider its values over all MSA columns (or pairs or triplets of columns), and all amino acids if appropriate, for both natural and synthetic MSAs. To obtain the vertical and horizontal coordinates (respectively) of the markers in each panel, we compare these values for each natural MSA with the values from the corresponding synthetic MSAs generated by bmDCA with default parameters or by our method based on MSA Transformer (respectively). We use four different scores for this comparison, and devote each column of the figure to one of these scores – from the left: $|\rho - 1|$, where ρ denotes the Pearson correlation; $|\text{Slope} - 1|$, where "Slope" means the slope of best linear fit (see Figs. S3 and S4 for illustrations of these first two quantities in the case of two- and three-body connected correlations for families PF00072 and PF00153); the Jensen–Shannon divergence between the distributions of values; the Wasserstein distance between these distributions. For each statistical or information measure (row) and each score (column), and for each family in Table 1, we have one value of the score comparing the natural and bmDCA-generated MSAs and another one comparing the natural and MSA-Transformer–generated MSAs. We plot the former value versus the latter, yielding one marker per protein family in each plot. Thus, each plot compares the ability of bmDCA and MSA Transformer to reproduce the statistics of the natural data. Blue markers (above the diagonal) mean that the scores for MSA-Transformer–generated MSAs are better, while green markers (below the diagonal) mean the opposite.

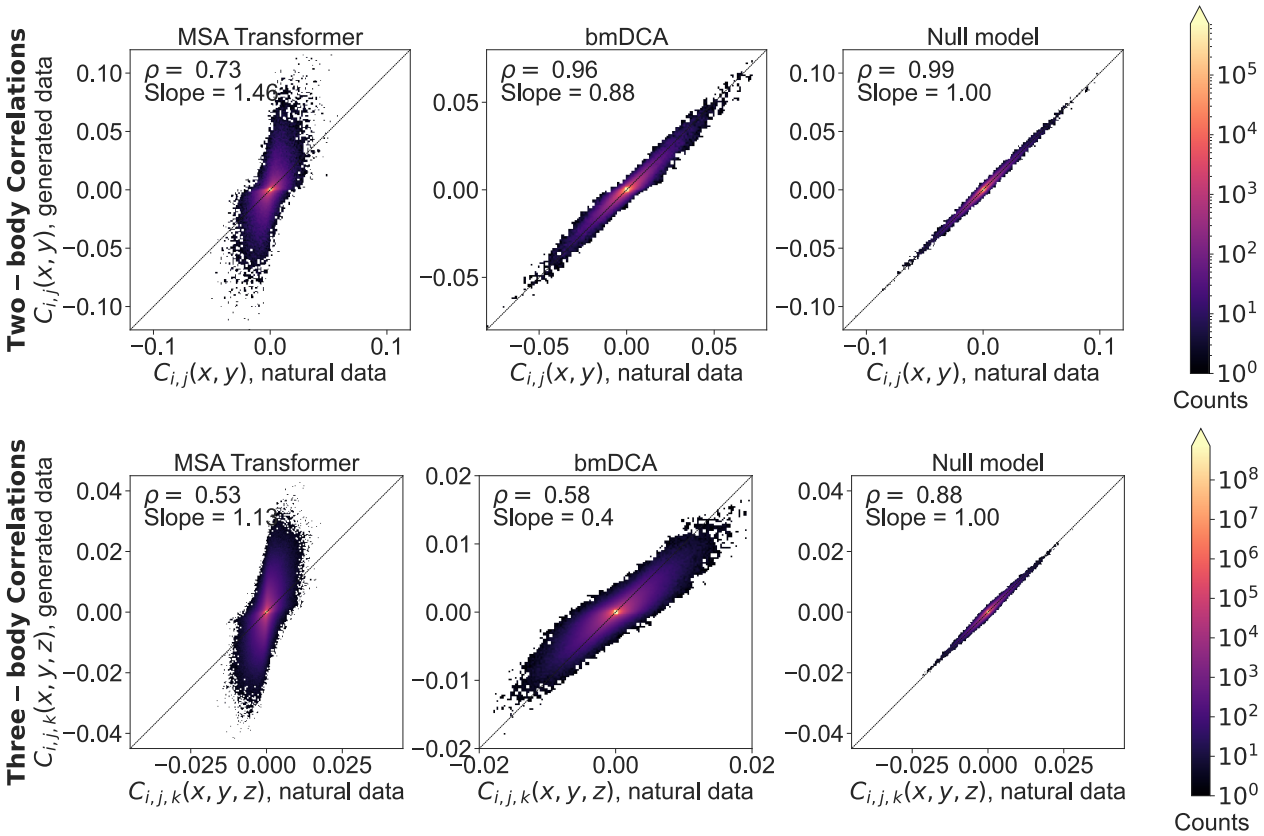

Figure S3: **Two- and three-body connected correlations estimated from generated MSAs versus the natural one, for family PF00072.** Relationships between connected correlations estimated from the MSA generated by MSA Transformer or bmDCA, and those estimated from the natural MSA, are shown as binned scatter plots both for two-body (top row) and three-body (bottom row) statistics. We include a null model (third column) obtained by splitting the natural MSA in half and comparing the statistics of one half with those of the other. It yields an estimate of the expected dispersion in these plots due only to finite size effects in a model-free, purely data-driven way. Pearson correlation coefficients ρ, and slopes of lines of best fit, are reported in each case. For the comparisons involving MSA Transformer, we used Eqs. (4) and (5) to estimate correlations. On the other hand, since bmDCA is trained to reproduce frequencies rescaled with phylogenetic weights (w_i in Eq. (8)), for the comparisons involving bmDCA we rescaled the natural frequencies before using Eqs. (4) and (5) to estimate correlations.

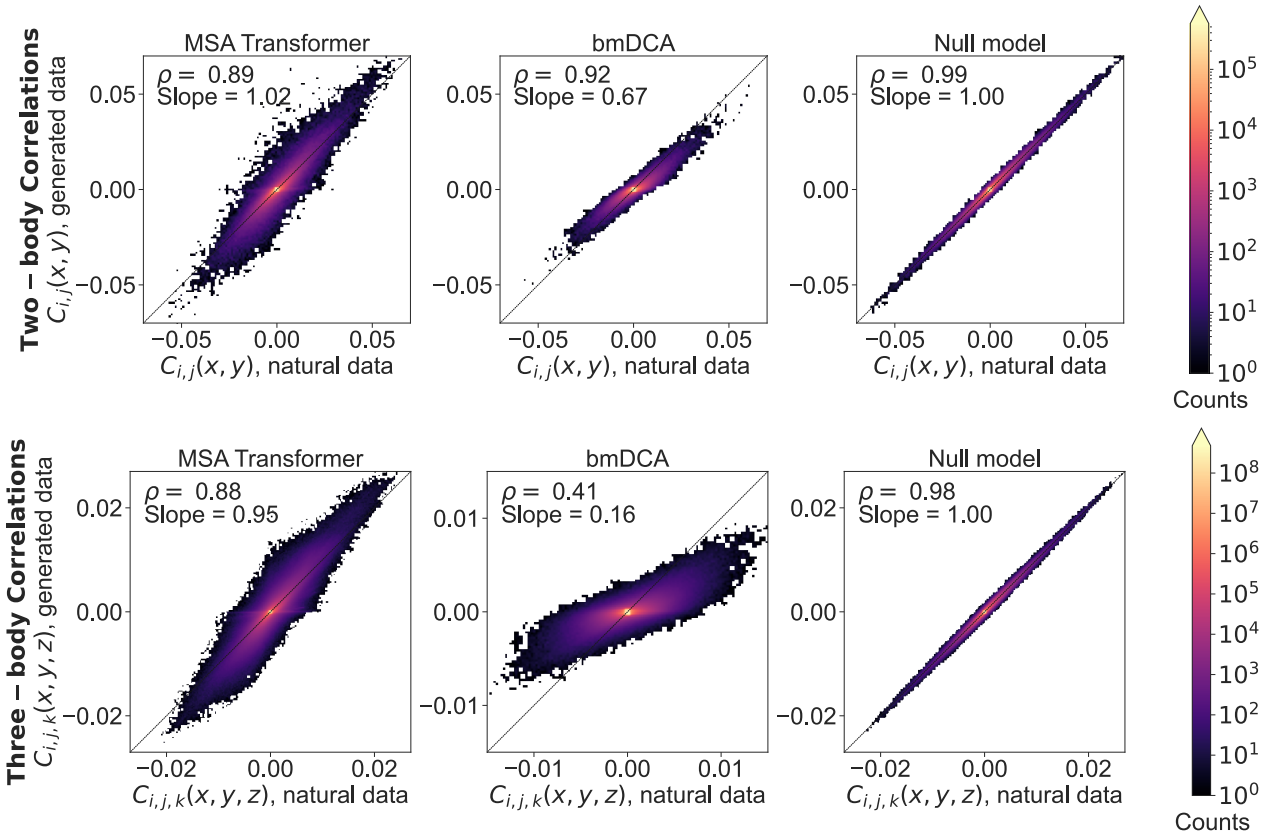

Figure S4: **Two- and three-body connected correlations estimated from generated MSAs versus the natural one, for family PF00153.** Same as Fig. S3, but for family PF00153.

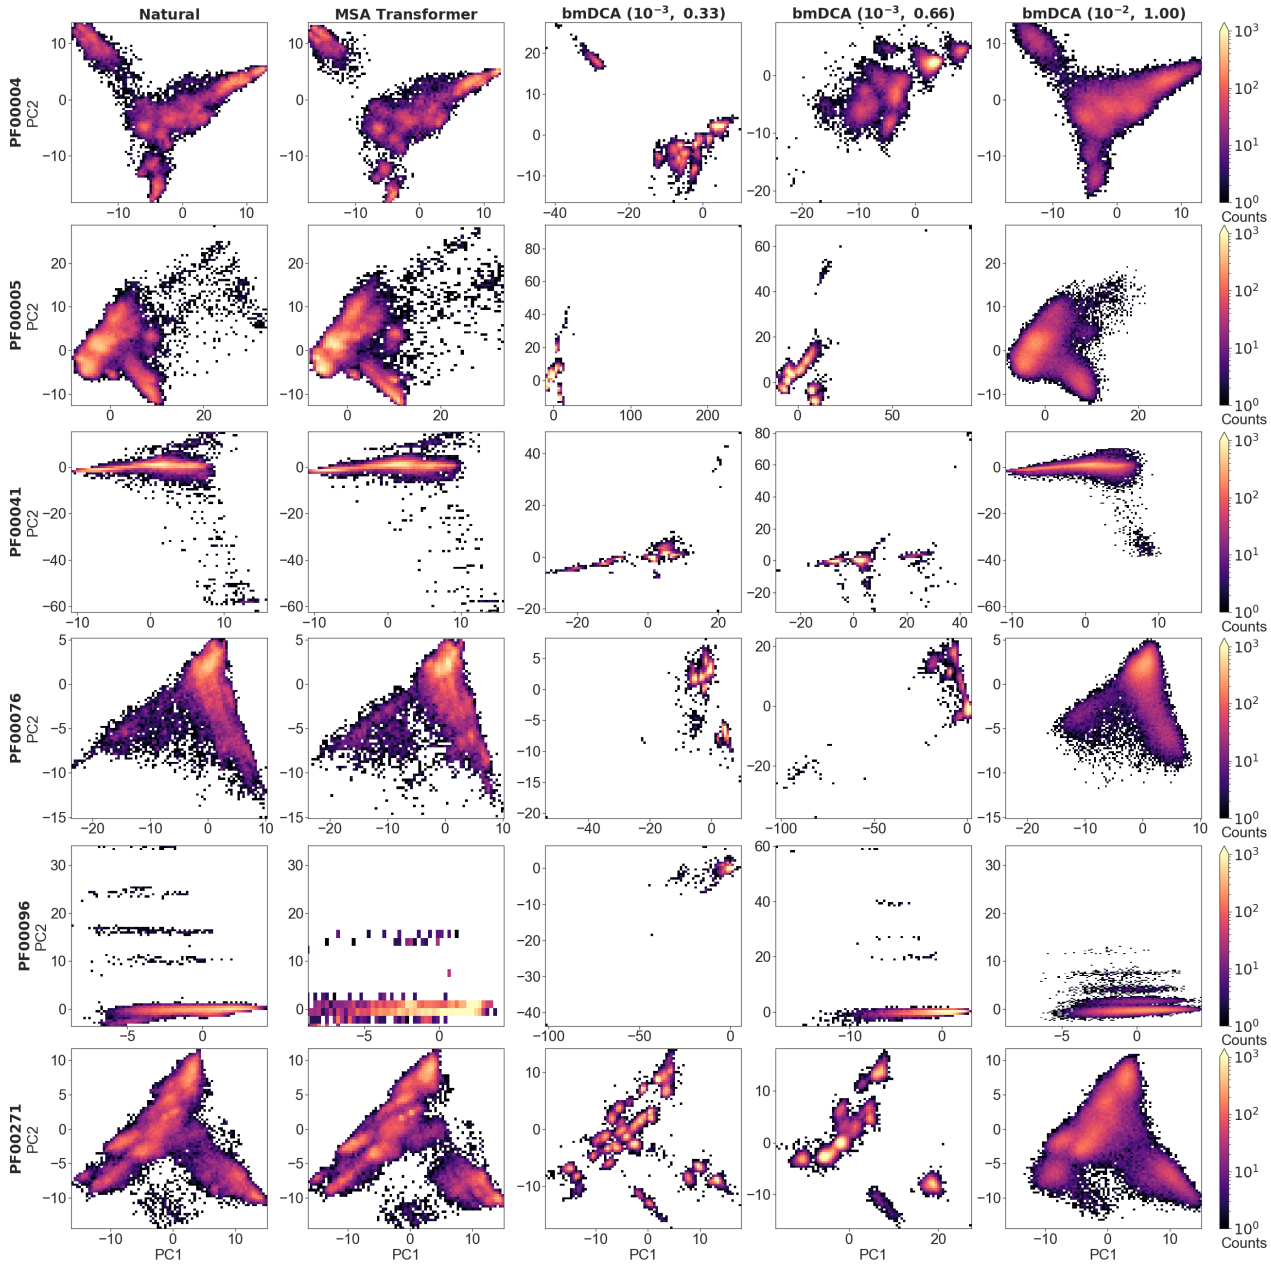

Figure S5: **Distribution of sequences in sequence space for all large protein families in our dataset.** We show the distribution of one-hot encoded natural and synthetic sequences projected in the subspace of the first two principal components of the natural MSA, as in Fig. 6 for families PF00072 and PF00153, but also including bmDCA at $(\lambda, T) = (10^{-3}, 0.66)$.

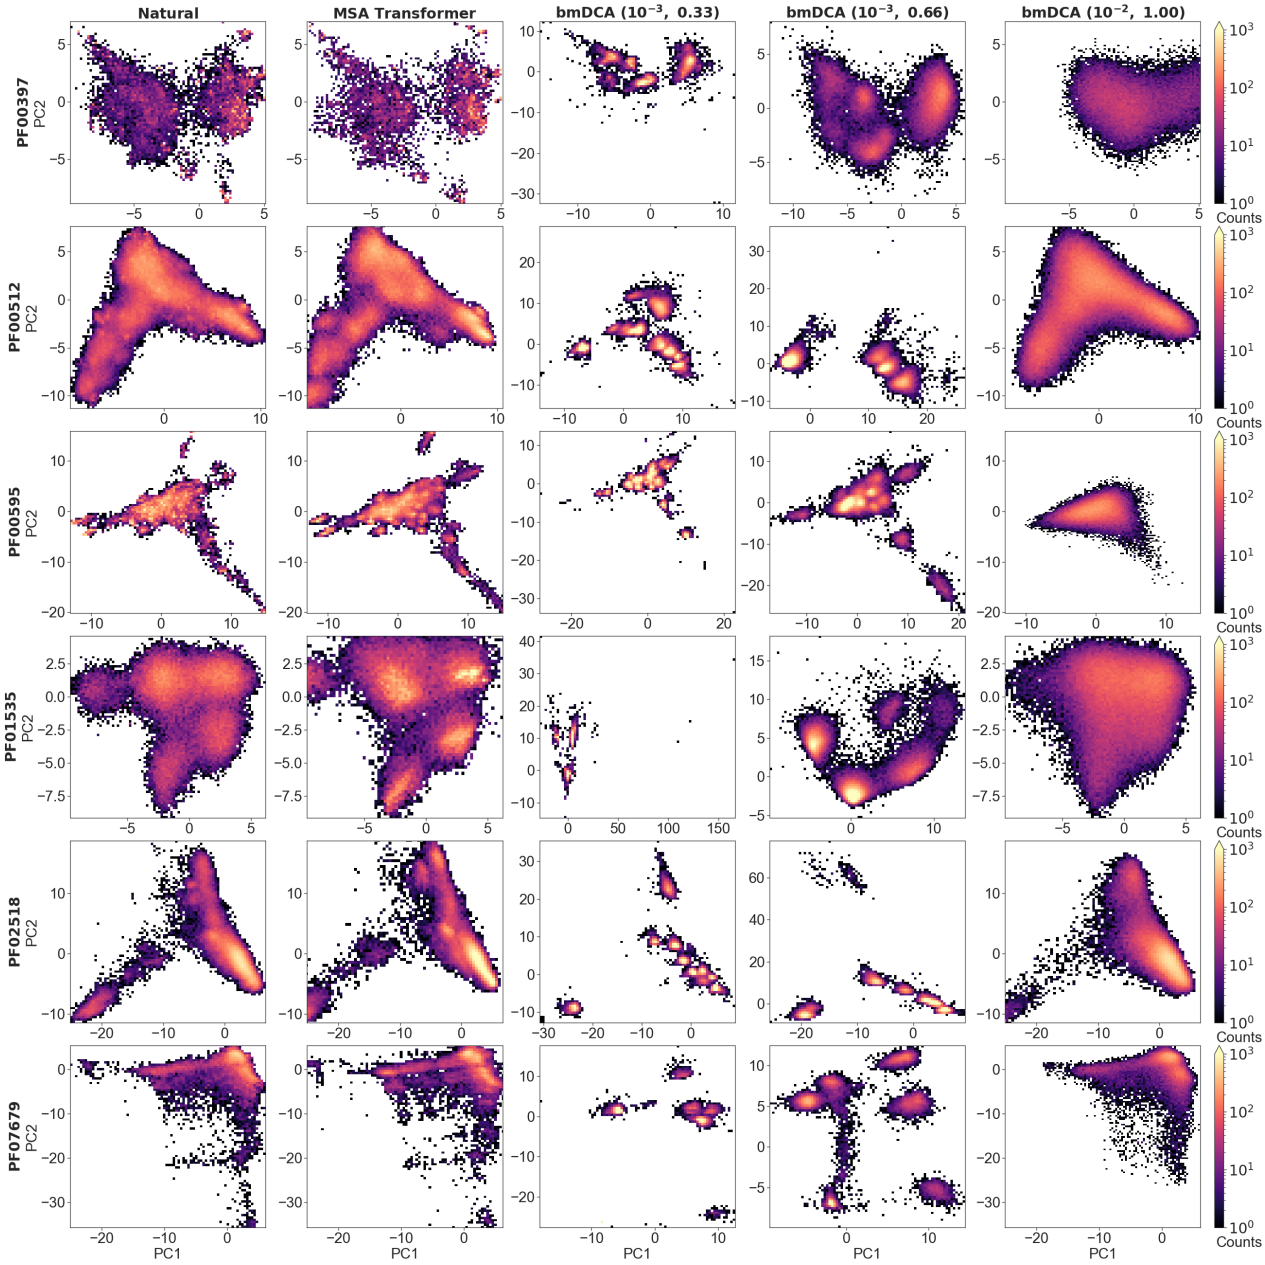

Figure S5: (Continued from previous page.)

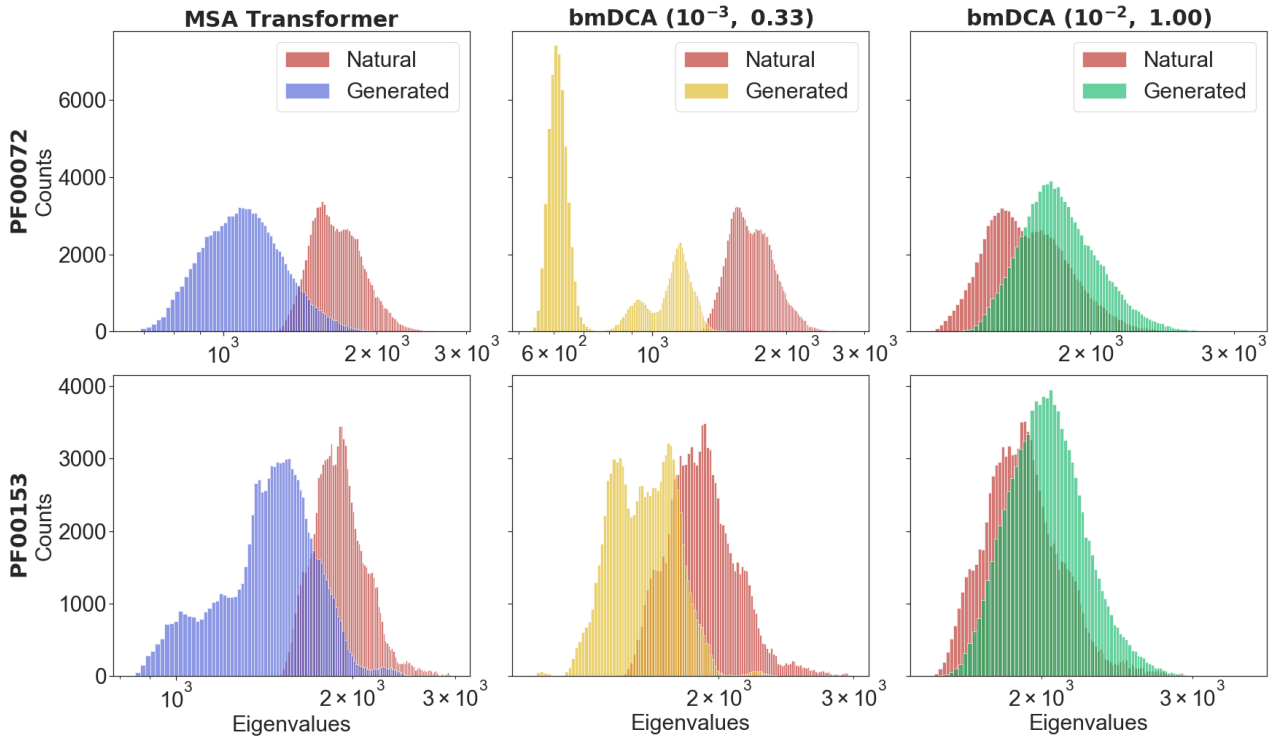

Figure S6: **Comparing phylogenies inferred from natural and generated MSAs for families PF00072 and PF00153.** We show the averaged spectra of modified graph Laplacian (MGL) matrices computed, for each tree inferred from an MSA, by using the leaves of multiple sub-trees made of 500 randomly sampled sequences from the MSA of interest. Specifically, in each case, we perform an average over 200 different sub-trees of the histograms of counts of the eigenvalues (see "Characterizing the distribution of sequences in MSAs"). We compare MSAs generated using either MSA Transformer (left) or bmDCA (center and right) to the natural MSA.

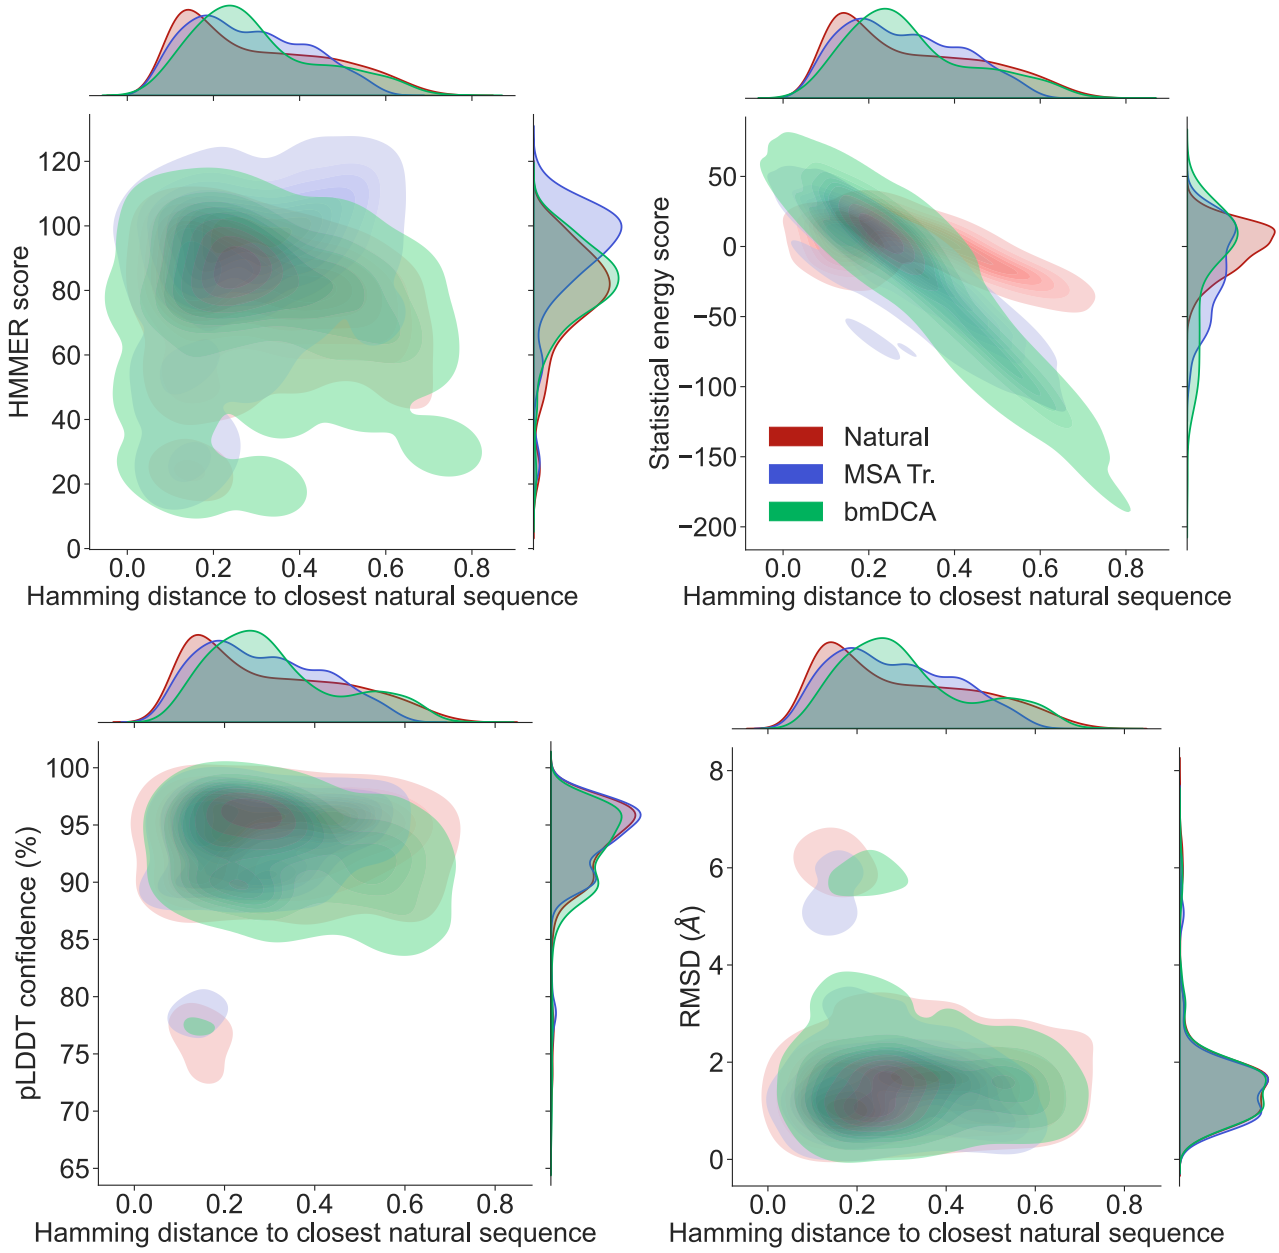

Figure S7: **Homology, coevolution and structural scores vs. distance to the natural MSA, for the chorismate mutase family.** We assess the performance of our generative method in the case of the chorismate mutase family, by comparing our generated sequences ("MSA Tr.") to natural ones, and to the sequences generated in Ref. [13] using bmDCA at various values of the sampling temperature and of the regularization strength. Our synthetic sequences were generated using the iterative masking procedure based on MSA Transformer, starting from the natural alignment used in Ref. [13]. Our sequences score similarly to natural sequences and to the bmDCA-generated sequences from Ref. [13], which were experimentally validated there.

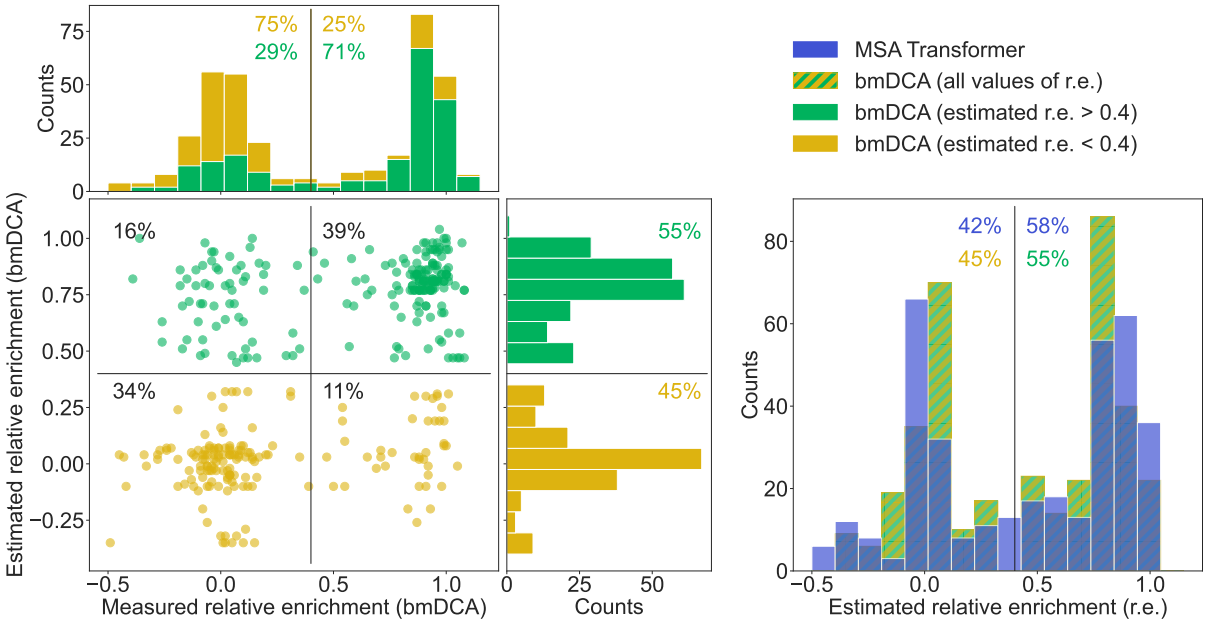

Figure S8: **Retrospective evaluation based on [13] for the chorismate mutase family. Left**: the estimated relative enrichment (r.e.) scores of the bmDCA-generated sequences are plotted versus their experimentally measured counterparts from [13]. We estimate the expected r.e. of these generated sequences as the r.e. of the closest natural sequence measured in [13]. We observe that high estimated r.e. is associated with high measured relative enrichment, as 71% of sequences with estimated r.e. > 0.4 (green) also have measured r.e. > 0.4. Note that in the top marginals (showing the measured r.e. for bmDCA-generated sequences), the two histograms are stacked on top of each other. Thus, the total histogram shows the distribution of all measured r.e. values for bmDCA-generated sequences. **Right**: overlaid histograms of estimated r.e. are shown for our MSA-Transformer–generated sequences and for the bmDCA-generated ones from [13]. Throughout this analysis, we restrict to the top 33% of sequences in terms of pLDDT scores, as it was shown in [62] that good structural scores help to select functional sequences.

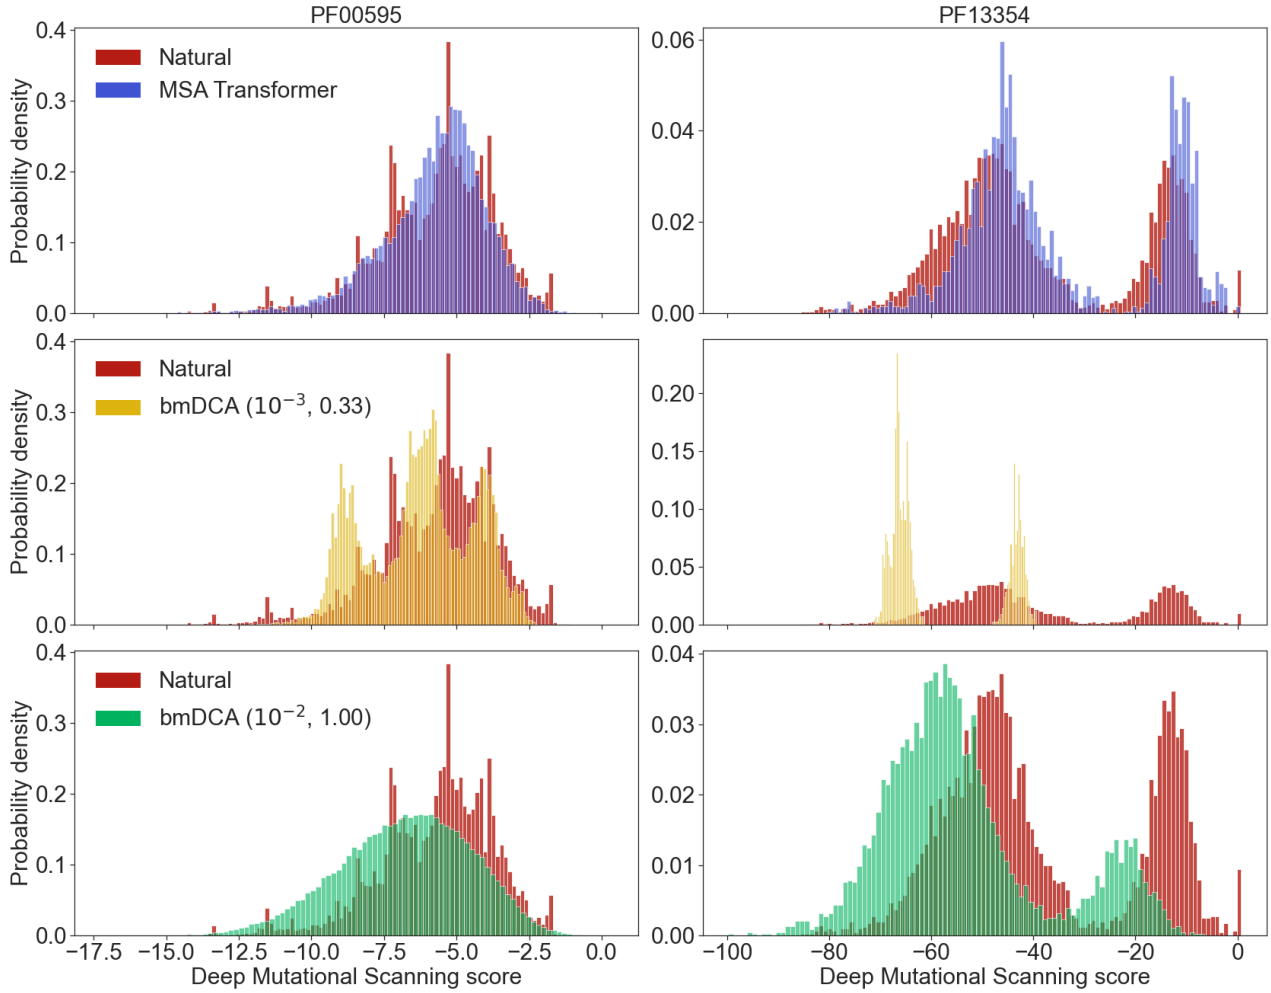

Figure S9: **Deep mutational scanning scores for families PF00595 and PF13354.** In deep mutational scanning (DMS) experiments, the fitness effects of all possible one-point mutations from a reference natural sequence are measured. To assign a DMS score to natural and generated sequences, we align each of them to the reference sequence, and sum the experimentally measured fitness effects of the relevant amino acids at each position. Higher values of the DMS score are better, as they mean higher fitness. We show normalized histograms of the DMS scores for sequences in the natural and generated MSAs of protein families PF00595 and PF13354, based on the DMS experiments in [63] and [64], respectively. For generated sequences, we restrict to those whose Hamming distance to their closest natural neighbor is larger than $\delta = 0.2$.

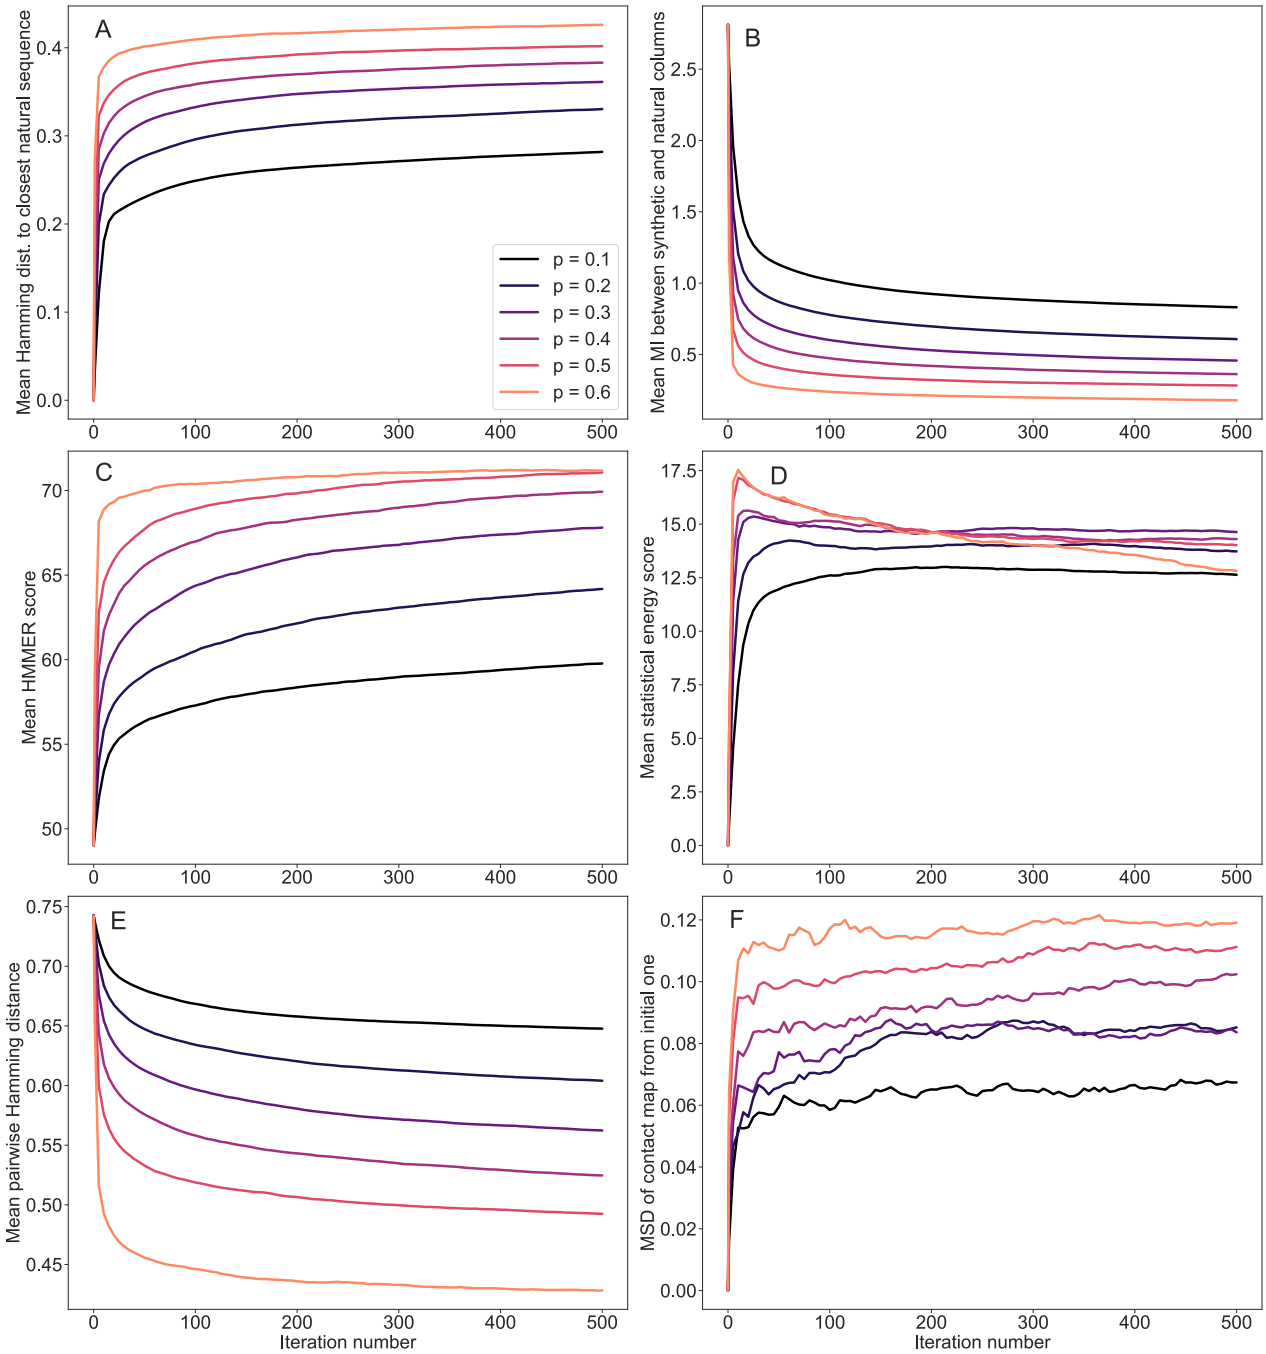

Figure S10: **Evolution of mean scores during the iterative masking procedure, for family PF00153.** Average scores of the generated sequences are reported for different iteration numbers and masking probabilities. The scores employed are: (**A**) Hamming distances to the closest natural sequence, (**B**) Mutual information between synthetic and natural columns of the MSAs, (**C**) HMMER scores – see "Scoring individual sequences", (**D**) Statistical energy scores (negative DCA statistical energies, shifted by their mean value for natural sequences – see "Scoring individual sequences"), (**E**) Pairwise Hamming distances between the generated sequences, (**F**) Mean square deviations (MSD) of the predicted contact maps at each iteration from the initial one (zero iterations). The synthetic MSAs, comprising 5000 sequences, were generated with MSA Transformer starting from natural sequences of family PF00153.

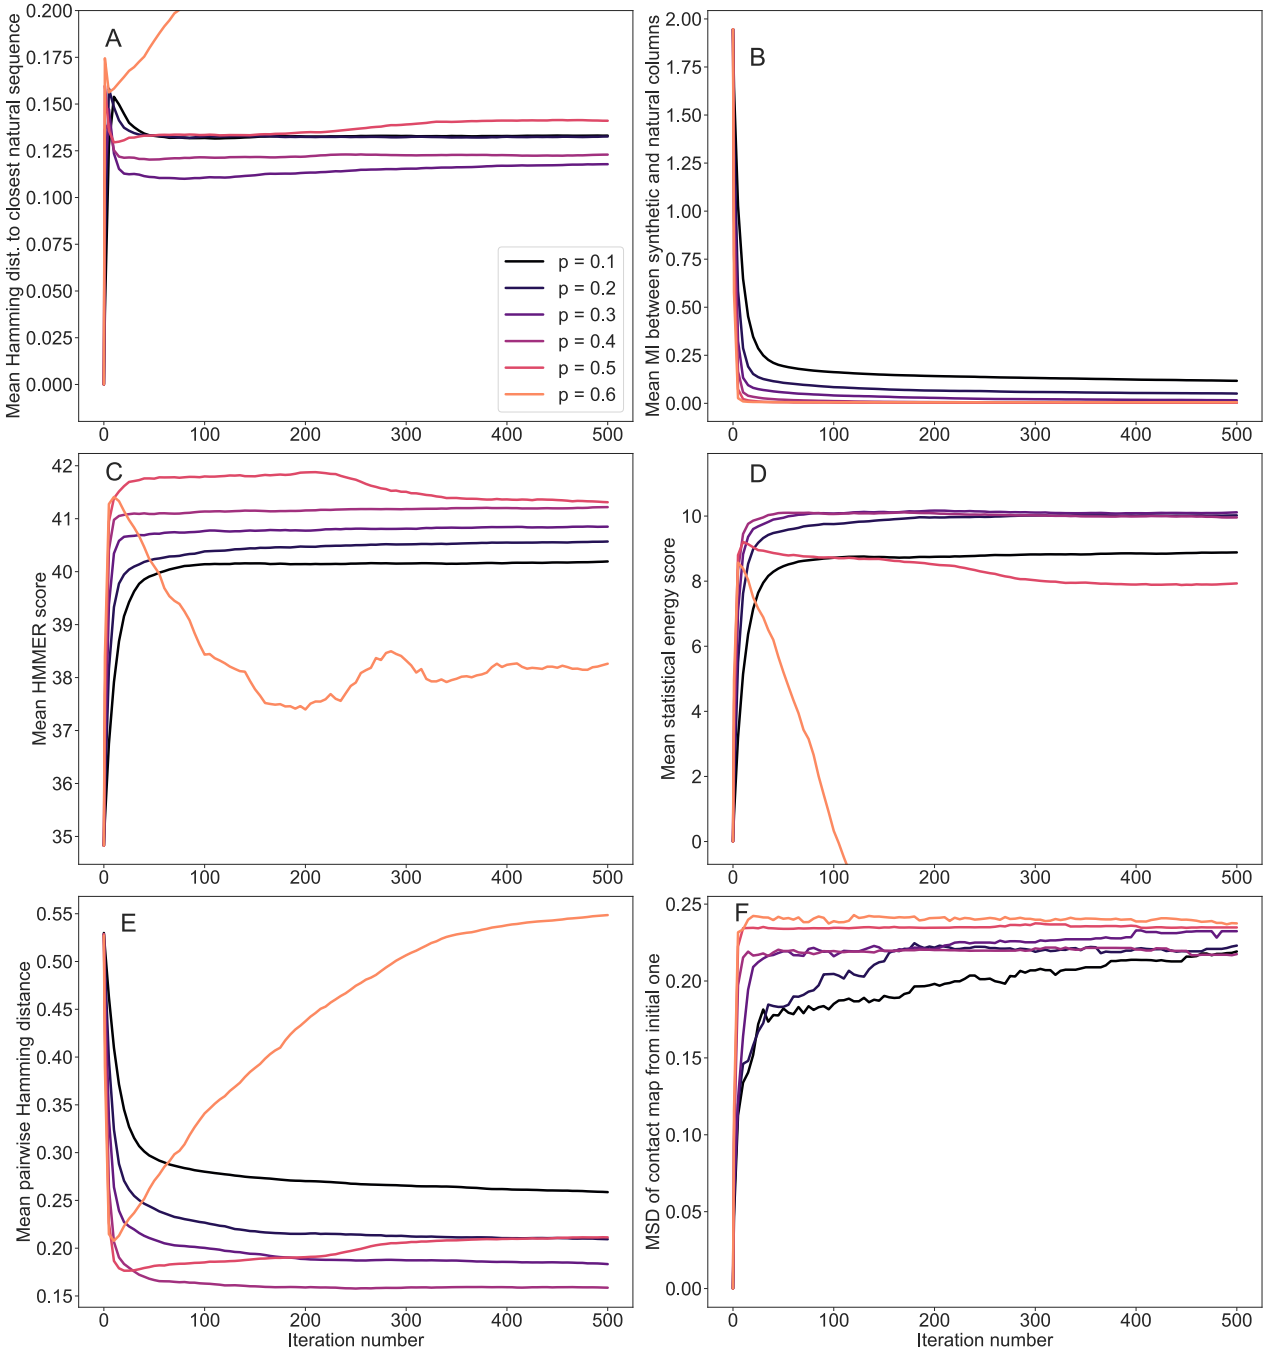

Figure S11: **Evolution of mean scores during the iterative masking procedure, for family PF00096.** Same as in Fig. S10, but for PF00096. This family has the shortest length among those considered here, namely $L = 23$ (see Table 1).

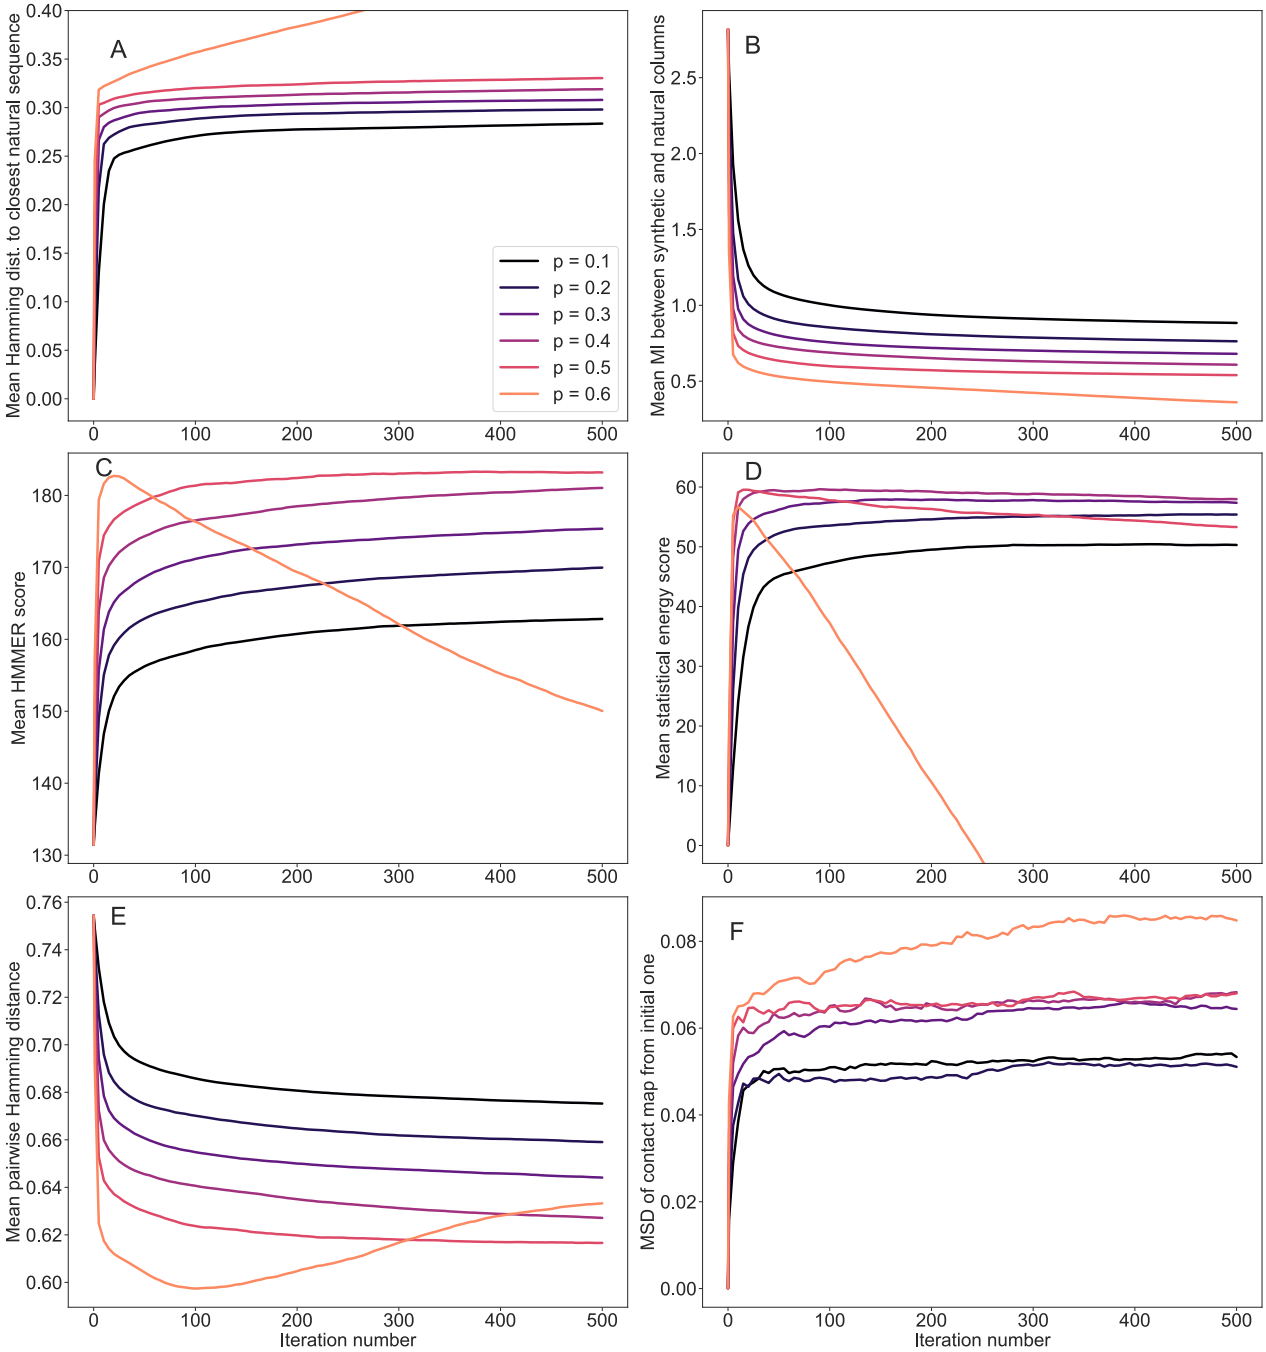

Figure S12: **Evolution of mean scores during the iterative masking procedure, for family PF13354.** Same as in Fig. S10, but for PF13354. This family has the largest length among those considered here, namely $L = 198$ (see Table S1).

Figure S13: **Evolution of inferred contact maps during the iterative masking procedure, for family PF00153.** Contact maps obtained from MSA Transformer for different iteration numbers and masking probabilities are reported. Probabilities of contacts are computed using the logistic regression on the output attention matrices from Ref. [25]. The input of the model consists of 100 different sequences chosen uniformly at random from the synthetic MSA generated at each iteration of our iterative masking procedure. These contact maps were employed to compute the mean square deviations shown in Fig. S10**F**.

Supplementary tables

Pfam ID	Family name	L	M	PDB ID	Resol.
PF01356	A_amylase_inhib	68	51	1OK0	0.93 Å
PF03440	APT	87	14	6RO0	2.13 Å
PF04008	Adenosine_kin	154	342	1WVQ	1.45 Å
PF06351	Allene_ox_cyc	175	378	2BRJ	1.50 Å
PF06355	Aegerolysin	131	322	6MYI	1.15 Å
PF16747	Adhesin_E	125	31	6GUT	1.63 Å
PF18648	ADPRTs_Tse2	155	9	5AKO	2.40 Å
PF13354	Beta-lactamase2	198	4642	6QW8	1.10 Å
-	Chorismate mutase [13]	96	1130	1ECM	2.20 Å

Table S1: **Other Pfam families and natural MSAs used in our analysis.** L denotes the length of an MSA and M its depth. The reference experimental PDB structures used for our RMSD calculations, and their resolutions, are also reported.

	HMMER score		**Statistical energy score**		**pLDDT confidence**		**RMSD**	
Pfam ID	MSA Tr. $\geq$ bmDCA1	MSA Tr. $\geq$ bmDCA2	MSA Tr. $\geq$ bmDCA1	MSA Tr. $\geq$ bmDCA2	MSA Tr. $\geq$ bmDCA1	MSA Tr. $\geq$ bmDCA2	MSA Tr. $\leq$ bmDCA1	MSA Tr. $\leq$ bmDCA2
PF00004	0	1.0	0	1.0	0.20	1.0	0.33	0.96
PF00005	0.93	1.0	0	1.0	$6.5 \cdot 10^{-19}$	1.0	0.78	0.96
PF00041	0.99	1.0	0	1.0	1.0	1.0	$2.9 \cdot 10^{-14}$	$5.9 \cdot 10^{-3}$
PF00072	$4.7 \cdot 10^{-12}$	1.0	0	1.0	$9.1 \cdot 10^{-6}$	1.0	$3.4 \cdot 10^{-6}$	0.96
PF00076	$9.5 \cdot 10^{-134}$	1.0	0	1.0	$6.5 \cdot 10^{-19}$	1.0	$9.2 \cdot 10^{-5}$	1.0
PF00096	0.04	1.0	0	1.0	0.01	0.92	0.92	$2.4 \cdot 10^{-5}$
PF00153	0.91	1.0	0	1.0	0.98	0.84	$9.3 \cdot 10^{-10}$	1.0
PF00271	$5.9 \cdot 10^{-30}$	1.0	0	1.0	0.33	1.0	0.02	0.88
PF00397	0	1.0	0	1.0	$1.5 \cdot 10^{-3}$	1.0	$4.8 \cdot 10^{-10}$	0.38
PF00512	0	1.0	0	1.0	$4.3 \cdot 10^{-3}$	1.0	0.96	0.67
PF00595	0.83	1.0	0	1.0	$1.0 \cdot 10^{-15}$	1.0	$7.2 \cdot 10^{-24}$	1.0
PF01535	0.98	1.0	0	1.0	1.0	1.0	0.78	$4.1 \cdot 10^{-8}$
PF02518	0	1.0	0	1.0	1.0	1.0	1.0	1.0
PF07679	1.0	1.0	0	1.0	1.0	1.0	$4.3 \cdot 10^{-3}$	1.0

Table S2: **p-values of the Kolmogorov–Smirnov test comparing the distributions of homology, coevolution, and structure-based scores across natural and synthetic MSAs.** For each score except the RMSD, we test the Null Hypothesis that the scores of MSA Transformer generated sequences are greater or equal than those of bmDCA generated sequences, in the (stringent) sense that their cumulative distribution function is always above. Here, bmDCA1 stands for bmDCA with $(\lambda, T) = (10^{-3}, 0.33)$ and bmDCA2 for bmDCA with $(\lambda, T) = (10^{-2}, 1)$. For the RMSD, the null hypothesis is that the scores of MSA-Transformer–generated sequences are smaller or equal than those of bmDCA-generated sequences (recall that smaller RMSDs are better). In all cases, a p-value close to one (resp. zero) means that the Null Hypothesis tested should be accepted (resp. rejected). Reported zero p-values are too small to be properly assessed by the algorithm.

	HMMER score			Statistical energy score			pLDDT confidence (%)			RMSD (Å)		
Pfam ID	Natural	MSA Tr.	bmDCA	Natural	MSA Tr.	bmDCA	Natural	MSA Tr.	bmDCA	Natural	MSA Tr.	bmDCA
PF00004	0.5 (0.2)	0.6 (0.2)	**0.8** (0.2)	0 (1)	0.8 (0.9)	**1.6** (0.1)	85.4 (4.1)	**85.8** (4.5)	81.7 (0.7)	3.4 (0.8)	**2.8** (0.7)	3.6 (0.5)
PF00005	0.7 (0.1)	0.8 (0.1)	0.8 (0.1)	0 (1)	1.8 (0.9)	**3.1** (0.2)	83.0 (6.9)	89.0 (4.2)	**91.6** (1.6)	3.8 (1.2)	2.8 (1.0)	2.8 (0.8)
PF00041	0.6 (0.1)	**0.9** (0.2)	0.5 (0.1)	0 (1)	1.5 (1.0)	**4.9** (0.5)	90.0 (4.5)	**92.0** (3.2)	79.2 (2.7)	2.1 (2.1)	**2.9** (0.5)	3.4 (2.2)
PF00072	0.7 (0.1)	**0.9** (0.1)	0.8 (0.1)	0 (1)	2.1 (0.8)	**3.8** (0.3)	94.5 (3.4)	**94.9** (1.9)	94.1 (0.5)	2.4 (0.3)	2.3 (0.1)	**2.1** (0.1)
PF00076	0.6 (0.1)	0.8 (0.2)	0.8 (0.1)	0 (1)	1.5 (0.8)	**3.5** (0.2)	82.2 (4.4)	84.6 (4.8)	**87.6** (1.5)	1.8 (0.5)	1.4 (0.6)	1.4 (0.1)
PF00096	0.8 (0.0)	0.9 (0.0)	0.9 (0.0)	0 (1)	2.2 (0.8)	**2.8** (0.3)	93.0 (2.0)	**94.0** (0.8)	93.7 (0.2)	0.6 (0.1)	0.4 (0.1)	0.4 (0.0)
PF00153	0.5 (0.1)	**0.6** (0.1)	0.5 (0.1)	0 (1)	0.6 (0.8)	**2.6** (0.3)	65.0 (5.0)	**66.6** (6.2)	64.9 (4.3)	5.1 (1.8)	4.4 (1.5)	**4.3** (1.1)
PF00271	0.5 (0.1)	0.5 (0.1)	0.5 (0.2)	0 (1)	1.0 (0.9)	**2.4** (0.3)	78.4 (4.6)	**86.4** (5.4)	83.8 (2.2)	2.0 (0.8)	2.3 (0.6)	**1.8** (0.1)
PF00397	0.7 (0.1)	0.8 (0.1)	0.8 (0.1)	0 (1)	0.5 (0.9)	**2.1** (0.2)	88.1 (2.2)	**88.9** (2.1)	88.2 (1.0)	0.9 (0.3)	0.9 (0.3)	0.9 (0.1)
PF00512	0.5 (0.1)	**0.8** (0.2)	0.7 (0.2)	0 (1)	1.5 (1.0)	**3.2** (0.3)	91.0 (4.0)	**90.2** (4.0)	89.5 (1.5)	2.1 (0.6)	**2.2** (0.5)	3.1 (0.2)
PF00595	0.7 (0.1)	0.7 (0.1)	0.7 (0.1)	0 (1)	0.5 (0.9)	**2.6** (0.3)	93.4 (4.5)	94.0 (4.8)	**95.1** (0.8)	1.8 (0.4)	1.7 (0.5)	**1.4** (0.2)
PF01535	0.5 (0.1)	**0.9** (0.2)	0.6 (0.1)	0 (1)	2.3 (1.1)	**4.1** (0.2)	82.4 (6.2)	**94.3** (5.5)	77.9 (3.6)	1.0 (1.1)	**0.4** (0.7)	0.5 (0.4)
PF02518	0.6 (0.2)	**0.8** (0.2)	0.7 (0.2)	0 (1)	1.9 (0.9)	**3.5** (0.2)	88.0 (6.0)	**91.0** (6.3)	73.6 (2.3)	4.1 (0.9)	**3.9** (0.5)	4.7 (1.1)
PF07679	0.5 (0.1)	**0.7** (0.2)	0.4 (0.1)	0 (1)	1.7 (1.0)	**5.2** (0.6)	93.5 (3.8)	**95.3** (2.9)	89.8 (2.2)	1.3 (1.0)	1.2 (0.5)	1.2 (0.2)

Table S3: **Median homology, coevolution, and structure-based scores in natural and synthetic sequences.** We report the median values of each of the scores shown in Fig. 1, as well as their standard deviations (between parentheses), for natural sequences, for sequences generated by our method based on MSA Transformer, and for sequences generated by bmDCA at low temperature, i.e. with $(\lambda, T) = (10^{-3}, 0.33)$ (denoted by "bmDCA"). Scores are normalized as Fig. 1, except that, for statistical energy, we subtract the median of natural scores instead of the mean for clarity (therefore, all natural MSAs have median 0 and standard deviation 1 for this score). For all scores, the best median among those of the two synthetic MSAs is shown in bold, and for pLDDT, it is shown in red if it is better than that the other synthetic MSA by a margin larger than the largest standard deviation. Recall that higher values are better for all scores, except RMSD, for which the opposite holds.

	Natural sequences	**MSA Transformer**				
		Iterative masking			**Context (greedy)**	
Score		Greedy	$T = 0.5$	$T = 1.0$	Fixed	Variable
Distance	0.155	0.271	0.305	0.514	0.232	0.362
HMMER	48.0	58.2	58.1	48.4	58.7	63.8
−Energy	0	13.0	8.5	−42.0	−15.4	−13.2
$\rho\,[C_{ij}]$	0.94	0.84	0.84	0.62	0.73	0.81
$\rho\,[C_{ijk}]$	0.89	0.80	0.76	0.41	0.66	0.77

Table S4: **Comparing different generation methods of MSA Transformer.** Various scores are shown for the natural MSA of protein family PF00153 and for synthetic MSAs generated in different ways from this family (each synthetic MSA comprises 10,000 sequences). For generation using MSA Transformer (see "Using MSA Transformer to generate sequences via an iterative masking procedure"), our standard iterative masking procedure is shown with its default greedy sampling (corresponding to $T = 0$) and two higher temperatures. Variants of the procedure where only the first sequence is masked ("Context", either fixed or variable, both with greedy sampling) are also shown. We report the mean Hamming distance to the closest natural sequence, which is not itself in the case of natural sequences ("Distance") as well as the mean HMMER and statistical energy scores ("−Energy") described in "Scoring individual sequences". Note that statistical energy scores are shifted by the mean value obtained for the natural MSA (which is −235.8). We also report the Pearson correlations between the two- and three-body statistics of the natural and the generated MSAs, denoted respectively by $\rho\,[C_{ij}]$ and $\rho\,[C_{ijk}]$ (for the natural MSA we report the Pearson correlation between two halves of this MSA), as illustrated in Fig. S3.

Type	λ	T	Distance	$M_{\text{eff}}^{(0.2)}$	HMMER	$-$Energy	$\rho\left[C_{ij}\right]$	$\rho\left[C_{ijk}\right]$	pLDDT (%)	RMSD (Å)
Natural	-	-	0.193	40180	90.3	0	0.99	0.88	93.6	2.5
MSA Tr.	-	-	0.348	9304	119.1	59.1	0.73	0.53	94.7	2.35
					Same as above, with synthetic context:				95.1	2.37
bmDCA	0.01	1	0.55	73062	66.5	-37.0	0.96	0.58	84.3	2.58
					Same as above, with synthetic context:				83.9	2.70
bmDCA	0.01	0.66	0.294	18911	101.7	92.2	0.48	0.11	94.2	2.61
bmDCA	0.01	0.33	0.251	12	103.2	118.3	0.42	0.05	94.2	2.55
bmDCA	0.001	1	0.525	73062	86.9	-18.3	0.97	0.63	89.7	2.44
bmDCA	0.001	0.66	0.296	21294	103.9	89.3	0.48	0.19	94.3	2.6
bmDCA	0.001	0.33	0.274	14	107.7	109.6	0.4	0.13	94.0	2.14
					Same as above, with synthetic context:				94.2	2.24

Table S5: **Impact of regularization strength and sampling temperature on sequence generation by bmDCA, for family PF00072.** We compare MSAs obtained using bmDCA with different regularization strengths λ (for inference) and sampling temperatures T (for generation) with the natural and the MSA-Transformer-generated MSAs. In each case, we report the average of the Hamming distances of each sequence to their closest natural neighbor, which is not itself in the case of natural sequences ("Distance"), as well as the effective MSA depth, the scores defined in "Scoring individual sequences", and the Pearson correlation coefficients of the two- and three-body connected correlations computed from natural and generated MSAs ($\rho\left[C_{ij}\right]$ and $\rho\left[C_{ijk}\right]$). For MSA Tr., bmDCA $(0.01, 1)$ and bmDCA $(0.001, 0.33)$, we also computed structural scores by feeding the entire synthetic MSA to Alphafold as context MSA (instead of using the natural MSA as context, see "Scoring individual sequences"). Structural scores are then very similar to those obtained using natural context.