Independent Training Trajectories Learn the Same Interaction Parameters in Coarse-Grained Protein Models

Marcos Lequerica-Mateos [1], Jonathan Martin [2], José N. Onuchic [3,4,*], Faruck Morcos [2,5,†] and Ivan Coluzza [3,6,7,8,‡]

[1] *CICA-Centro Interdisciplinar de Química e Bioloxía, Rúa as Carballeiras, Universidade da Coruña, A Coruña 15071, Spain*

[2] *Department of Biological Sciences, University of Texas at Dallas, Richardson, Texas 75080, USA*

[3] *Center for Theoretical Biological Physics, Rice University, Houston, Texas 77005, USA*

[4] *Department of Physics and Astronomy, Department of Chemistry, Department of BioSciences, Rice University, Houston, Texas 77005, USA*

[5] *Departments of Bioengineering, Physics, and Center for Systems Biology, University of Texas at Dallas, Richardson, Texas 75080, USA*

[6] *Smalley-Curl Institute, Rice University, Houston, Texas 77005, USA*

[7] *Ken Kennedy Institute, Rice University, Houston, Texas 77005, USA*

[8] *Chemistry Department, Rice University, Houston, Texas 77005, USA*

(Received 15 May 2026; accepted 14 September 2026; published 4 November 2026)

This work introduces a self-consistent training framework for learning residue–residue interaction parameters in matrix-based coarse-grained protein models. By combining direct coupling analysis with a physics-based design model, we construct a feedback loop in which the interaction matrix is iteratively updated from sequences generated by the model itself. We provide a mathematical foundation for this behavior by showing that, for a given iterative map satisfying explicit contact-locality and spectral assumptions, the training process is a contraction, guaranteeing convergence to a unique fixed point independent of the initial interaction matrix. This fixed-point property is general for coarse-grained models in which the residue-residue energy is represented by a distance-dependent term modulated by a type-dependent coupling matrix, while the learned interaction parameters remain specific to the underlying model physics. We then examine the empirical consequence of this framework in Caterpillar, a coarse-grained potential for folding and design. Independent training trajectories, defined by different target structures and stochastic sampling, converge to statistically equivalent interaction matrices. Each trajectory recovers the full interaction solution without requiring ensemble averaging, and the averaged trajectory reaches the same result. The resulting framework provides a link between sequence statistics and physically modeled interaction energies. Applied to protein design, it yields interaction matrices associated with improved structural fidelity and accurate contact prediction, while clarifying how learned sequence patterns relate to the interactions that stabilize protein structure.

DOI: 10.1103/dmcv-sx8q

I. INTRODUCTION

Protein design plays a central role in biotechnology, medicine, and fundamental studies of biomolecular function [1–5]. The ability to engineer sequences with prescribed structural and functional properties enables the development of enzymes, therapeutic agents, and synthetic biomaterials. Existing strategies rely primarily on rational design or directed evolution [6]. Rational approaches depend on detailed structural knowledge, which is not always available, whereas directed evolution requires extensive experimental screening and optimization. A longstanding complementary goal is to formulate reduced physical models that capture how sequence encodes foldability and stability while remaining simple enough to analyze mechanistically [7–10].

A major objective in protein design is to establish a quantitative relationship between amino acid sequence and structural stability. Machine learning methods, including AlphaFold and Rosetta-based approaches [11–13], have significantly improved predictive accuracy. These models capture complex statistical patterns in sequence and structure data, yet their internal representations remain difficult to interpret in terms of physical interactions [14]. This gap motivates complementary frameworks in which learned parameters can be related directly to underlying biophysical principles. In particular, maximum-entropy and inverse-statistical-physics approaches provide an interpretable route for connecting sequence statistics to effective couplings and to quantities that can be related to biophysical energetics [15–18]. Interpretable approaches provide a route for analyzing the statistical structure learned by data-driven models while preserving a direct connection to physical energetics. This perspective is also consistent with the broader energy-landscape view of protein

*Contact author: jonuchic@rice.edu
†Contact author: faruckm@utdallas.edu
‡Contact author: ic33@rice.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

folding, in which effective interactions are expected to encode the funnel-like organization that stabilizes native structure while suppressing competing states [9,19]. Recent advances in sequence design using neural networks [20] and reviews of generative modeling in protein science [21] further underscore the need for such frameworks.

The present work develops a self-consistent training framework based on Direct Coupling Analysis (DCA) integrated with a coarse-grained physical model. Building on earlier correlated-mutation analyses [22] and later global inference approaches [23], DCA uses maximum-entropy or Potts-model inference to disentangle direct from indirect correlations in sequence alignments and thereby infer effective residue–residue couplings [15,16,24–26]. Its predictive performance has improved through advances in statistical inference, including pseudolikelihood maximization [27], and has been validated through extensive benchmarking [28,29]. Because DCA yields an explicit pairwise statistical model, it is especially well matched to coarse-grained frameworks in which residue energetics are also represented through pairwise interaction matrices. Applications to designed sequences indicate that DCA can recover interaction patterns in the absence of evolutionary constraints [30,31]. More broadly, related statistical models have been used to probe energetic couplings and evolutionary sectors, to quantify context-dependent mutational effects, and to enable generative protein design [32–36]. These properties make DCA suitable for a closed-loop procedure in which interaction parameters are learned directly from sequences generated by a physical model.

More generally, coarse-grained protein models provide a well-established framework for representing protein folding, dynamics, and design using reduced physically motivated interactions [37–41]. We implement this framework within the Caterpillar model, a coarse-grained protein design and folding approach defined by physically motivated interaction terms [42]. Throughout this paper, interaction matrix is the primary term for the object that encodes effective residue–residue interactions. In related contexts, the same object may also be described as a coupling matrix, energy matrix, or potential. Here we use a single term to keep the mapping between inference and design explicit.

The workflow begins from a target protein structure and an initial interaction matrix M_{initial}, which provides a generic starting point and is not derived from sequences associated with that target. Given an input matrix M_i, the Caterpillar model generates an ensemble of designed sequences compatible with the target fold through Monte Carlo sampling. The resulting multiple sequence alignment is analyzed with DCA to infer an updated interaction pattern M_{i+1}. A positive global scaling factor then sets the energy scale, equivalently the effective design temperature, for the next physical simulation; it is not propagated as an independent matrix update. In the implementation studied here, the initial condition is $M_{\text{initial}} = M^1_{\text{CAT}}$.

The central hypothesis of the paper is that this self-consistent procedure converges to a unique model-specific interaction matrix. At the mathematical level, the relevant object is a fixed iterative map acting on normalized interaction patterns. For a given map satisfying the stated contraction assumptions, we show convergence to a unique fixed point independent of the initial interaction matrix. This fixed-point property is not restricted to Caterpillar, but applies to the class of coarse-grained models in which a type-dependent matrix $M(S_i, S_j)$ modulates a model-specific distance function $\Gamma(r_{ij})$. The identity and physical interpretation of the converged fixed point remain determined by the underlying model.

The empirical question is distinct from the theorem-level statement. In the specific Caterpillar implementation, the refinement depends on the target structure used during sequence generation and on Monte Carlo sampling fluctuations. Different targets and independent simulations therefore define different realizations of the overall procedure. We test whether these target-conditioned stochastic trajectories converge to statistically equivalent interaction matrices.

The interaction matrix recovered at convergence is model-specific because it reflects the effective energetics of the underlying coarse-grained representation. This convergence can be rationalized by separating directional interactions, which stabilize local secondary-structure motifs, from the longer-range sequence-dependent interactions that mainly tune packing and select among compatible local arrangements[43–47]. Because the refinement targets these packing-dominated effective interactions rather than the hydrogen-bond network itself, an interaction matrix learned from one protein family can remain transferable to another family even when their secondary-structure content differs substantially. The class-general result concerns the existence and uniqueness of a self-consistent solution within this family of models. The converged interaction matrix then defines the parameters that stabilize foldable sequences for a given structural topology while remaining consistent with the statistical structure extracted from the generated sequences.

II. METHODS

A. Overview of the approach

This study builds upon the Caterpillar protein model [48] and the mean-field approximation of direct coupling analysis (mfDCA) [24] to introduce a self-consistent iterative framework for refining the interaction matrix through sequence design and analysis. Detailed descriptions of the Caterpillar model, including its energy terms and coarse-grained representation, are provided in the Supplemental Material (The Caterpillar protein model; Fig. S1), together with the full mathematical formulation of mfDCA (Direct Coupling Analysis) [49]. Here we focus on the elements required to define the iterative framework and the objects that are passed between the design and inference steps.

B. mfDCA: Mean-field direct coupling analysis

mfDCA refers to a family of computational methods developed to improve the accuracy of covariance-based approaches for inferring protein three-dimensional structure from evolutionary sequence data. Its key contribution is the ability to filter out indirect correlations, which are the main source of noise in traditional covariance analyses. Each DCA variant uses its own strategy to isolate direct statistical couplings, but all share a common mathematical foundation built around the

following Hamiltonian:

$$H_{\mathrm{mfDCA}} = -\sum_{1\leqslant i<j<L} e_{ij}(A_i, A_j) - \sum_{i=1}^{L} h_i(A_i). \quad (1)$$

Here the coupling term $e_{ij}(A_i, A_j)$ represents the inferred interaction energy between residue types A_i and A_j at positions i and j of the sequence, respectively. The local-field term $h_i(A_i)$ corresponds to the estimated propensity of observing residue type A_i at position i. In the present framework, these inferred couplings are mapped to an interaction matrix so that the output of mfDCA can be compared directly with the matrix used in the Caterpillar design model. The full mfDCA formulation and matrix-construction details are provided in the Supplemental Material (Direct Coupling Analysis) [49].

C. The Caterpillar protein model

To design the sequences, we employed the Caterpillar model, a protein framework capable of folding and designing sequences in which side chains are coarse-grained as self-avoiding spheres with a diameter of 2 Å. The Caterpillar-based sequence design approach guarantees that solutions refold to the target structure with high precision (2–3 Å). Although these artificial sequences are not intended for experimental validation, they encode the structural logic needed to test the self-consistency of the interaction matrix in a controlled environment. At this stage, they should be regarded as computational constructs for evaluating the framework rather than as biologically meaningful sequences ready for experimental use.

In design simulations, exploration of the vast sequence space is guided by a Monte Carlo scheme that minimizes the following energy function:

$$H_{\mathrm{CAT}} = E_H + \sum_i E_{\mathrm{sol}}(A_i) + \sum_{i<j} E_{\mathrm{res}}(i, j). \quad (2)$$

Here E_H corresponds to the hydrogen-bond energy, which remains constant throughout sequence-design simulations, $E_{\mathrm{sol}}(A_i)$ represents the solvation contribution modeled by penalizing buried hydrophilic residues and exposed hydrophobic ones, and E_{res} denotes the residue–residue interaction term. Further details of the Caterpillar Hamiltonian, including the hydrogen-bond, side-chain, and residue-solvent terms, are provided in the Supplemental Material (The Caterpillar protein model; Fig. S1 [49]).

The residue–residue interaction component of the energy is computed as

$$E_{\mathrm{res}}(i, j) = M^1_{\mathrm{CAT}}(S_i, S_j)\Gamma(r_{ij}), \quad (3)$$

where i and j are the residue indices along the sequence, $M^1_{\mathrm{CAT}}(S_i, S_j)$ are the entries of a 20×20 matrix defining interaction strengths between amino-acid types S_i and S_j, and $\Gamma(r_{ij})$ is the Caterpillar distance function evaluated at the Cα–Cα distance between residues i and j. The matrix M^1_{CAT} was derived in previous work [42], where it was optimized using the Maximum Entropy Principle. Specifically, residue-residue and residue-solvent parameters were tuned so that the hydrophobicity profiles and total interaction energies of the designed sequences matched those of naturally occurring sequences. Caterpillar-designed sequences, which fold to the target structure while remaining free from evolutionary noise and alignment artifacts, provide a controlled environment for benchmarking mfDCA. This setting allows us to test whether mfDCA recovers the interaction patterns embedded during the design process and therefore whether the inference step is suitable for feedback into the iterative framework (see Supplemental Material: Accurate Recovery of Embedded Interactions via mfDCA; Table S1 [49]).

D. Sequence design algorithm

To minimize external biases, we employ the sequence design algorithm recently introduced in Ref. [50]. By using revised mutation probabilities and optimized acceptance rules, this algorithm improves the efficiency of sequence space exploration (see Supplemental Material: Virtual Move Parallel Tempering for Sequence Space Sampling; Sequence Design; Parameters of the Caterpillar simulations [49]). The design strategy aims to replicate the sequence heterogeneity characteristic of natural proteins by relaxing constraints on amino acid usage. Unlike conventional algorithms that fix the global amino acid composition to a specific target, this approach allows the composition to vary dynamically through stochastic reshuffling. By permitting evolution without preimposed distributions, the system spontaneously converges to a "permutation composition"—a metric of selective amino acid utilization—that aligns closely with natural sequences rather than the theoretical maximum. This unconstrained methodology was validated through folding free energy calculations, which confirmed high-precision refolding to native structures, and by mfDCA, which demonstrated contact prediction accuracy exceeding 82%.

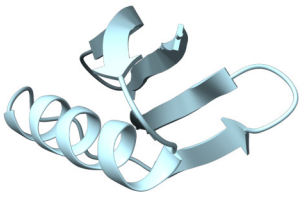
E. Protein dataset selection

To evaluate the effectiveness of the self-consistent iterative refinement framework, we selected six protein structures from the Protein Data Bank (PDB), with lengths ranging from 56 to 220 residues. This dataset was chosen to represent the diversity found in naturally occurring proteins, encompassing a range of three-dimensional configurations and varying degrees of structural disorder (see Fig. 1).

Extensive design simulations were performed for each target structure, producing between 20 000 and 300 000 candidate sequences. To maintain a representative and diverse set of sequences, only those with less than 90% sequence identity to one another were retained. This threshold balances diversity without overly restricting the number of sequences for analysis. Additionally, subsets were limited to a maximum of 20 000 sequences and required a minimum of 10 000 sequences to proceed to further analysis.

F. Folding simulations and validation metrics

Folding simulations were conducted to evaluate the performance of the self-consistent iterative framework at the level of structural behavior. Sequences with the lowest free energy, $F(E, \ln(N_p))$, from each design generation were selected, where N_p represents the number of possible sequence permutations based on length. These simulations mapped folding



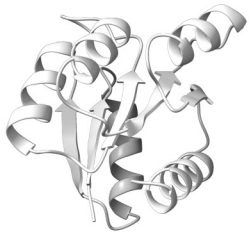
Protein: IgG-binding protein G
PDB Code: 1PGB
Length: 56



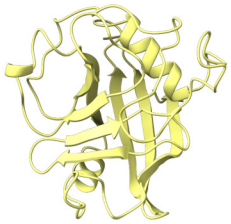
Protein: Barstar
PDB Code: 1A19
Length: 89



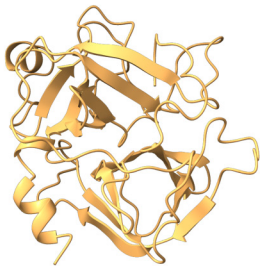
Protein: Peptidyl-prolyl cis-trans isomerase FKBP1A
PDB Code: 2PPN
Length: 107



Protein: Transcriptional regulatory protein WalR
PDB Code: 1NXW
Length: 118



Protein: Peptidyl-prolyl cis-trans isomerase 1
PDB Code: 1A33
Length: 174



Protein: Anionic Tripsin
PDB Code: 1MBQ
Length: 220

FIG. 1. Selected protein structures and corresponding dataset information. Each panel shows the target structure together with its protein name, PDB identifier, and sequence length.

free energy landscapes, $F(DRMSD)$, against the Distance Root-Mean-Square Deviation (DRMSD) from the native PDB structures to quantify how closely the designed sequences refolded to their original configurations. Folding quality was evaluated using two quantities: the DRMSD at which the free energy minimum was located and the DRMSD error within 1–2 K_bT units of energy. These metrics provide a direct readout of folding accuracy and stability. We tested the folding simulations of Caterpillar-designed sequences and verified that, as in previous publications [42,50], the designed sequences reliably adopt native-like conformations. The DRMSD values ranged from 4.1 to 5.5 Å, confirming the structural fidelity of the designed sequences relative to their respective PDB structures. In addition, the free energy landscapes exhibited distinct funnel shapes, indicative of stable folding pathways directed toward the target structures (see Supplemental Material: Folding Simulations; Fig. S5 [49]). These observations justify using DRMSD and landscape shape as downstream indicators of the quality of the refined interaction matrix.

G. mfDCA application and PPV metrics

The accuracy of mfDCA in predicting residue-residue interactions was evaluated by calculating the positive predictive values (PPVs) for the designed sequences, using the top direct information (DI) values as contact predictors. PPVs

TABLE I. PPV for contact predictions using mfDCA on designed sequences. The top $L/5$ DI scores exhibit PPVs nearing perfect prediction, while extending to the top $5L$ maintains high accuracy.

Protein	L	PPV $L/5$	PPV $5L$
1PGB	56	0.89	0.82
1A19	89	0.86	0.95
2PPN	107	0.96	0.96
1NXW	118	1.00	0.98
1A33	174	0.97	0.96
1MBQ	220	0.90	0.96

quantify the fraction of predicted residue-residue contacts that are true contacts and therefore provide a direct measure of DCA reliability. Specifically, we calculated PPVs for two sets of predicted contacts: the top $L/5$ residue pairs, where L is the sequence length, and the top $5L$ residue pairs. The $L/5$ subset focuses on the highest-ranked contacts, whereas the $5L$ subset provides a broader assessment of accuracy across a larger number of predicted contacts. These complementary metrics allow us to evaluate both the precision of DCA in identifying the strongest interactions and its overall reliability across a wider range of predicted pairs.

The Caterpillar model is particularly useful in this context because it generates sequences without biological functional constraints. The absence of evolutionary pressure reduces confounding noise, allowing the designed sequences to reflect primarily the physical interaction rules embedded within the model. For Caterpillar-designed sequences, we generally observe high PPVs for the top DI pairs, indicating that mfDCA captures the structural logic encoded by the design model (see Table I).

H. Interaction matrix extraction

The extraction of the interaction matrix is a central component of this work and builds on a previously introduced method [30]. mfDCA first returns position-dependent couplings $e_{ij}(\alpha, \beta)$, whereas the Caterpillar Hamiltonian uses a position-independent 20×20 matrix whose indices are amino-acid types. Positional and structural dependence enters the physical model through $\Gamma(r_{ij})$. We define the selected pair set

$$\mathcal{P} = \{(i, j) : |i - j| > 4,\ 6\,\text{Å} < r_{ij} < 11\,\text{Å}\} \tag{4}$$

and project the position-dependent couplings onto the amino-acid-type matrix according to

$$M^{\text{DCA}}_{\alpha\beta} = \frac{1}{|\mathcal{P}|} \sum_{(i,j)\in\mathcal{P}} \frac{e_{ij}(\alpha, \beta)}{\Gamma(r_{ij})}. \tag{5}$$

The sequence-separation cutoff removes local backbone correlations, while the spatial window focuses the average on nonlocal residue-residue-mediated interactions and suppresses hydrogen-bond contributions. The normalization is by the number of selected pairs. This projection produces the position-independent interaction matrix compared with the Caterpillar matrix used in design (see Supplemental Material: Direct Coupling Analysis; Accurate Recovery of Embedded Interactions via mfDCA; Table S1 [49]).

After extraction, the inferred matrix is rescaled to set the global energy scale used in the next Caterpillar design simulation. Specifically, M^N_{DCA} is divided by the slope of the regression line obtained by comparing the design matrix M^N_{CAT} with the recovered matrix M^N_{DCA}. Multiplying the interaction matrix by a positive global factor is equivalent to changing the effective Monte Carlo design temperature. At every generation, however, the next interaction matrix is inferred anew from the correlations of the newly generated sequence ensemble. The regression factor is therefore not propagated as an independent matrix update and does not define an additional nonlinear operation in the fixed-point Jacobian. It remains physically necessary to balance the residue-residue term against the other contributions to the Caterpillar Hamiltonian.

I. Self-consistent iterative refinement framework

The primary innovation of this work lies in the self-consistent iterative refinement process, termed the *Design Generation* framework (see Fig. 2), which systematically improves the interaction matrix by integrating sequence design with mfDCA analysis. The formal map analyzed for convergence is $\mathcal{F} = \mathcal{I} \circ \mathcal{G}$ on normalized interaction patterns, equivalently on projective classes $[M] = \{\lambda M : \lambda > 0\}$. The generation step $\mathcal{G}$ produces a designed sequence ensemble at the effective design temperature set by the input matrix amplitude, and the inference step $\mathcal{I}$ returns a newly extracted normalized interaction pattern from that ensemble. The subsequent regression scaling selects the amplitude used in the next physical simulation but is not a third operation in the analyzed map.

Starting from an initial interaction matrix (M^1_{CAT}), sequences are designed using the Caterpillar model. mfDCA then infers position-dependent couplings e_{ij}, which are projected onto an updated amino-acid-type interaction pattern (M^N_{DCA}). The extracted matrix is rescaled as described above and used in the next generation as M^{N+1}_{CAT}. The matrix pattern is inferred anew at each generation; only its physical amplitude is calibrated before feedback into Caterpillar.

J. Independent and averaged refinement paths

To assess the flexibility of this self-consistent framework, we explored two refinement paths. In the Independent Paths approach, each protein's interaction matrix was iteratively refined separately, allowing the design process to evolve along target-specific trajectories. In the Averaged Paths approach, interaction matrices derived from multiple proteins were averaged at each iteration to produce a generalized matrix ($M^{\text{avg}}_{\text{DCA}}$). These two settings provide different empirical realizations of the same overall workflow, while the theorem-level argument in the Results section applies to a fixed map defined by a specific choice of generation and inference steps.

K. Statistical analysis and iterative improvement

Correlations between the updated interaction matrices, M_{DCA}, and the design matrices, M_{CAT}, were evaluated using

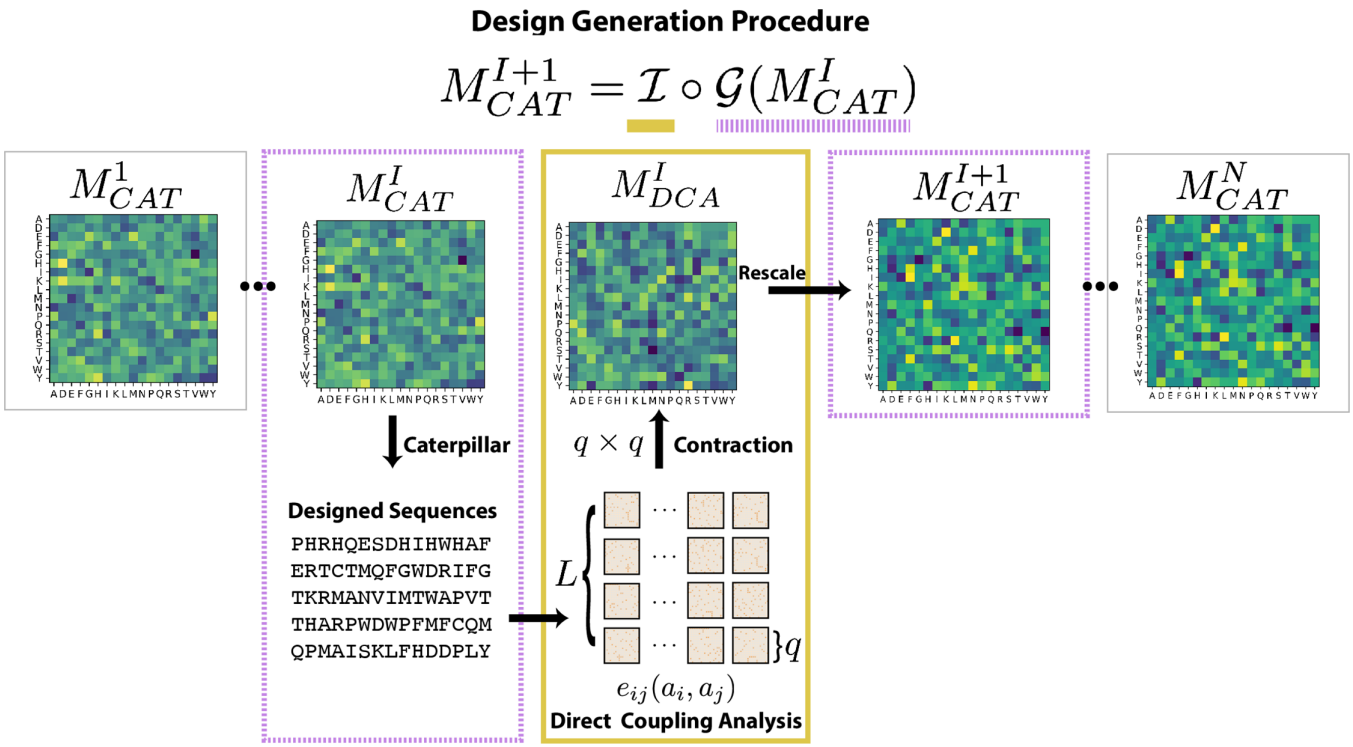


FIG. 2. Schematic of the Design Generation process. Starting with M_{CAT}, sequences are designed and subjected to DCA to extract M_{DCA}. After scaling, M_{DCA} is used as M_{CAT} for the next generation. This iterative process allows continuous refinement and optimization of the interaction matrix, maintaining equilibrium between energy components.

R^2 metrics. This analysis quantifies the iterative improvement of matrix alignment across generations. Additional comparisons based on rotational alignment and Z-score statistics are introduced in the Results section to evaluate whether independently refined matrices are statistically equivalent beyond direct element-wise correlation.

III. RESULTS AND DISCUSSION

A. Convergence to a fixed-point solution

At the theorem level, we ask whether the self-consistent framework converges for a fixed iterative map. One normalized update is

$$M^{(n+1)} = \mathcal{F}(M^{(n)}) = \mathcal{I} \circ \mathcal{G}(M^{(n)}), \tag{6}$$

where $\mathcal{G}$ generates a sequence ensemble at the effective design temperature associated with the input matrix and $\mathcal{I}$ infers a new normalized interaction pattern through mfDCA. The map acts on projective classes $[M] = \{\lambda M : \lambda > 0\}$: the positive regression factor used in the physical implementation sets the amplitude for the next design simulation but is not propagated as an independent matrix update. In the theorem-level discussion $\mathcal{F}$ is fixed; target-conditioned and Monte Carlo-sampled Caterpillar pathways considered below are separate empirical realizations.

The model class consists of coarse-grained Hamiltonians whose type-dependent matrix modulates a model-specific distance function,

$$H = -\frac{1}{2} \sum_{k \neq l} \sum_{g,h} M_{gh} \Gamma(r_{kl}) \delta(X_k, g) \delta(X_l, h). \tag{7}$$

Thus the fixed-point property is not restricted to Caterpillar, but the map, fixed-point location, and physical interpretation remain model-specific.

The Banach fixed-point theorem guarantees convergence to a unique fixed point independent of the initial matrix if, for a fixed map, there is a $k < 1$ such that

$$\|\mathcal{F}(M) - \mathcal{F}(M')\| \leqslant k \|M - M'\|. \tag{8}$$

For a differentiable map, a sufficient local condition is $\|J\| < 1$, where $J = \partial \mathcal{F} / \partial M$. The chain-rule and linear-response derivation in the Supplemental Material [49] yields the estimate

$$\|J\| \leqslant \epsilon, \qquad \epsilon = \frac{1}{(L(q-1))^2} \sum_{(ia)(jb)} |(C^{-1})_{(ia)(jb)}|. \tag{9}$$

The localization assumption entering this bound is formulated on the physical contact graph, not as literal bandedness in sequence order. Excluded volume gives a minimum representative-point separation $a > 0$, and the finite contact cutoff of $\Gamma(r_{ij})$ gives a contact radius R. Every neighbor of a residue therefore lies in a fixed-radius ball while neighboring representative points remain separated by at least a. A sphere-packing argument gives

$$\deg(i) \leqslant z_c(a, R) < \infty, \tag{10}$$

where z_c is independent of protein length. Closest-packing residue coordination supports this physical picture [51], while protein contact-network studies support the use of a local contact graph [52].

We decompose the regularized covariance into a contact-local part and a residual,

$$C = C_G + R, \tag{11}$$

where C_G retains covariance blocks on the bounded-degree contact graph and R contains weak indirect and finite-sampling correlations. We assume length-independent spectral bounds for C_G and $\|C_G^{-1}R\| < 1$, so that

$$C^{-1} = \left(I + C_G^{-1}R\right)^{-1} C_G^{-1} = \sum_{n=0}^{\infty} \left(-C_G^{-1}R\right)^n C_G^{-1}. \tag{12}$$

This identity gives perturbative stability in operator norm. Entrywise localization of the full inverse additionally requires R to belong to a compatible off-diagonally localized matrix class, for which inverse-closedness results apply [53].

Ordinary band-matrix inverse decay provides the historical starting point [54,55], but the contact graph is not generally one-dimensional. For any polynomial p,

$$(p(C_G))_{uv} = 0 \quad \text{when} \quad \deg p < d_G(u, v), \tag{13}$$

because a power C_G^k connects only vertices joined by a graph path of length at most k. Combining this polynomial locality with polynomial approximation of $f(x) = 1/x$ on the spectrally bounded interval containing $\sigma(C_G)$ gives the sparse matrix-function estimate [56]

$$\left|\left(C_G^{-1}\right)_{uv}\right| \leqslant K \exp[-\gamma d_G(u, v)], \tag{14}$$

with constants controlled by the spectral bounds rather than matrix size.

To express this localization in sequence order, we use compact-globule scaling. A compact segment containing $s = |i - j|$ residues occupies volume proportional to s, so its characteristic linear extent scales as $s^{1/3}$ [57–59]. Identifying contact-graph distance with this compact spatial scale is a physical assumption of the present model,

$$d_G(i, j) \propto |i - j|^{1/3}, \tag{15}$$

and yields the stretched-exponential sequence-localization bound

$$|(C^{-1})_{(ia)(jb)}| \leqslant K' \exp[-\alpha |i - j|^{1/3}]. \tag{16}$$

This is the precise sense in which the inverse covariance is effectively or pseudo-banded along sequence order. Because

$$\sum_{s \geqslant 0} \exp(-\alpha s^{1/3}) < \infty, \tag{17}$$

the absolute row sum remains bounded, the total absolute weight grows at most linearly with L, and $\epsilon = O(1/L)$. Hence $\epsilon < 1$ for sufficiently large systems within this explicitly defined contact-local, spectrally bounded class; this is not a claim for arbitrary covariance matrices.

For the Caterpillar implementation, the numerical value $\epsilon = 0.3908$ was computed from the covariance matrix at convergence of the averaged pathway, which combines the sequence ensembles generated for all proteins in the dataset. It is a numerical check of the contraction condition for that implementation, rather than the sole basis of the analytical scaling argument. Under the stated assumptions, a fixed map converges to a unique model-specific fixed point; agreement among different target-conditioned paths is tested separately below.

We now turn to the empirical results for the specific Caterpillar implementation. The following results address trajectory-level behavior under specific model realizations. The discussion proceeds in three steps. We first examine convergence of the interaction matrices across generations. We then evaluate whether different target-conditioned stochastic trajectories are statistically equivalent. Finally, we assess whether this convergence is reflected in contact prediction and folding behavior.

B. Self-consistent refinement of residue-residue interaction parameters

This subsection examines how the iterative refinement process improves the alignment between the inferred interaction matrices M_{DCA} and the design interaction matrices M_{CAT} over successive generations. At each step I, the interaction matrices M_{DCA}^I, obtained from mfDCA, were compared with the corresponding design matrices M_{CAT}^I, which encode the interaction patterns used in the Caterpillar model at that generation. These matrices were also evaluated against the original design matrix M_{CAT} used as the initial input for the algorithm (see Fig. 3). In the first generation, the correlation coefficient R^2 values ranged from 0.76 to 0.84 across the six proteins (see Table I). This strong initial correlation is consistent with earlier findings [30] and indicates that mfDCA already captures much of the interaction structure embedded in the design process even in the absence of evolutionary constraints.

As the design iterations progressed, the correlation between M_{DCA}^N and M_{CAT}^N consistently improved, indicating that the feedback loop between design and mfDCA progressively aligns the inferred and applied interaction matrices (see Table I). A consistent scaling effect was observed across iterations, as expected from the design temperature and the statistical invariance of coupling values under constant shifts. The normalization described in the Methods section removes this scale ambiguity while preserving the relative interaction structure, which allows successive generations to be compared on a common basis.

Figure 3 shows this progressive increase in alignment across generations for both refinement protocols. The steady improvement in R^2 values indicates that the interaction matrices derived through the iterative design process become increasingly consistent with the matrices used to generate the sequences. This pattern is consistent with the expected behavior of a self-consistent refinement procedure and provides the empirical starting point for the stronger comparison across independent trajectories.

By implementing the design generation process over six generations, we observed clear improvements in matrix alignment across three testing scenarios. Both the Independent Paths method (see Supplemental Material: Figs. S2 and S3 for the alternative initial matrix and its convergence behavior [49]) and the Averaged Paths method converged to near-perfect correlation (R^2 approaching 1) by the sixth generation. This result shows that the iterative protocol stabilizes the interaction matrix within the Caterpillar implementation. We

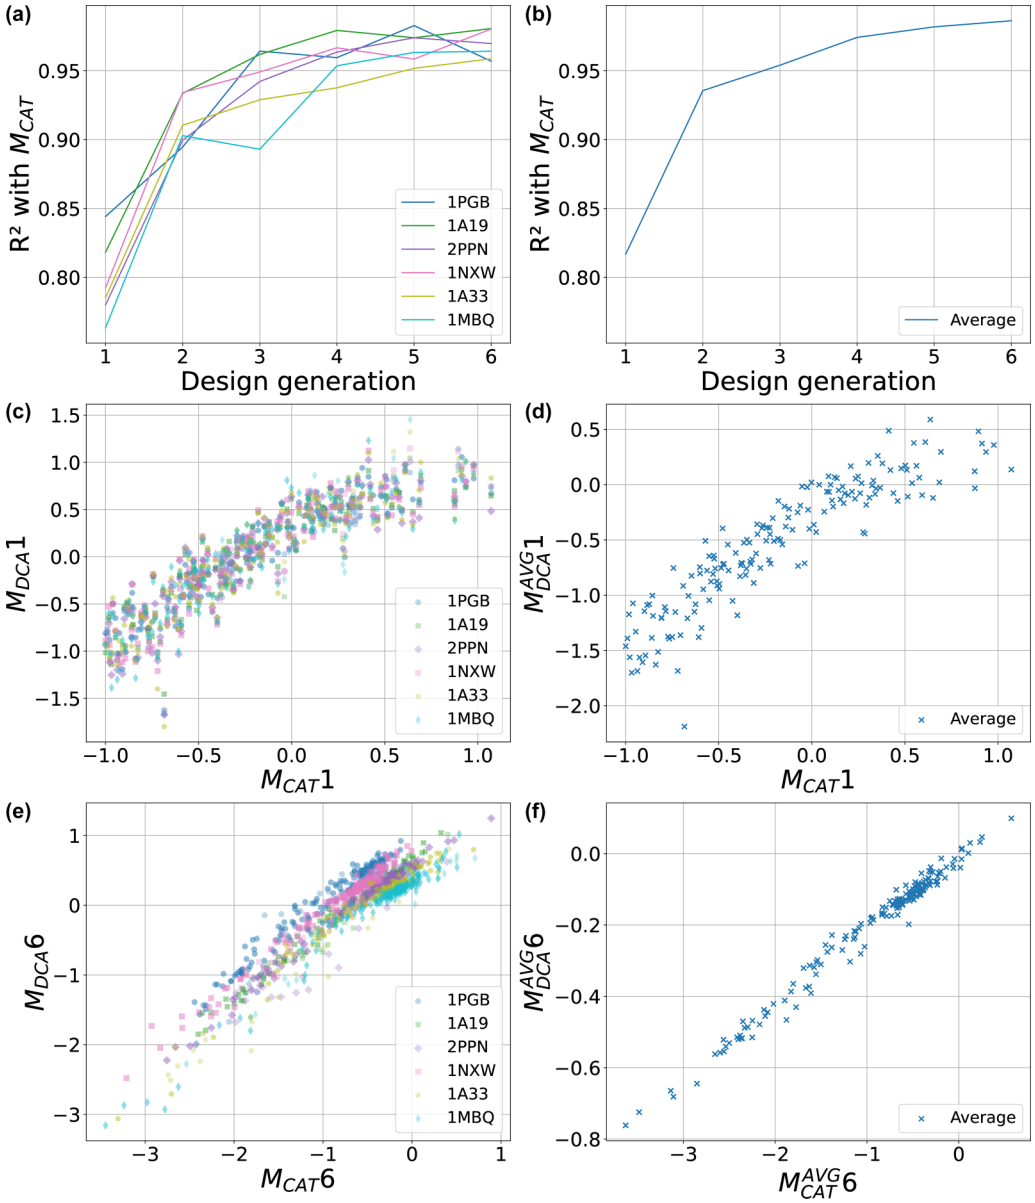


FIG. 3. Evolution of interaction matrix correlations across Design Generations for Independent and Averaged Paths. (a) R^2 values between recovered and previous generation matrices for each protein in the Independent Paths method, showing rapid convergence to high correlation. (b) R^2 values for the Averaged Paths method, similarly reaching near-perfect correlation. (c) Correlation between the recovered M_{DCA} and original M_{CAT} matrices in the first generation (Independent Paths). (d) Same as (c) for the Averaged Paths, showing strong initial correlation with the Caterpillar matrix. (e) Correlation between the recovered and used matrices in the sixth generation for the Independent Paths method, demonstrating nearly perfect alignment. (f) Same as (e) for the Averaged Paths, further confirming the high accuracy of interaction matrix recovery by the final generation.

now move from matrix-level convergence to trajectory-level comparisons.

Convergence to a common interaction matrix across independent paths

We now address a distinct empirical question. Different target structures and independent Monte Carlo simulations define different refinement trajectories, so agreement among their converged matrices is not implied by initialization-independence alone. To compare these trajectories, we first define what statistical equivalence means in this context. Direct element-wise comparison is insufficient because interaction matrices can represent the same effective physical structure in different orthogonal bases. A meaningful comparison therefore requires alignment in a common basis and a statistical test against an appropriate random-matrix baseline.

Two interaction matrices are considered statistically equivalent when, after an orthogonal rotation into a common interaction basis, their similarity exceeds the expectation from a null ensemble of random matrices. The rotation is required because residue–residue interaction matrices define effective Hamiltonians up to an orthogonal transformation of the interaction basis, so physically equivalent interaction structures

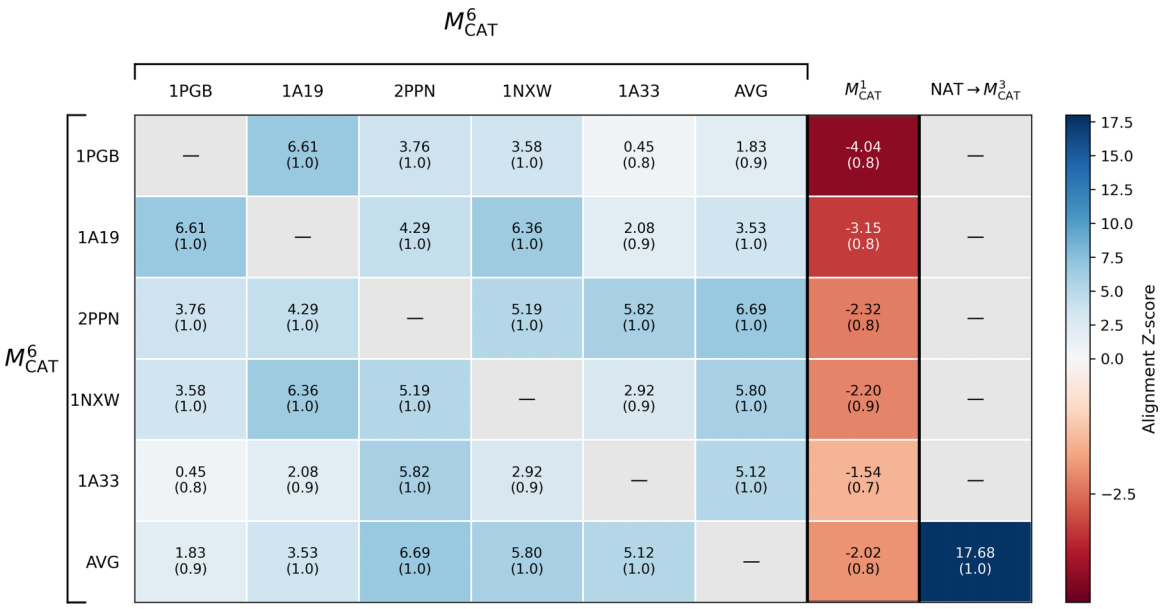


FIG. 4. Z-scores for initial alignment via Q-reconstruction between interaction matrices of selected proteins at Gen 6 (M^6_{CAT}), the synthetic averaged matrix (AVG), the initial matrix (M^1_{CAT}), and the matrix derived from natural sequences (NAT → M^3_{CAT}). Pearson correlation coefficients (r) for each comparison, approximated to a single significant digit, are provided in parentheses. Higher Z-scores indicate statistical significance relative to a null distribution of random symmetric matrices. Negative values signify structural divergence exceeding the null expectation, indicating that the matrices cannot be aligned more effectively than random symmetric matrices. Hence the M^1_{CAT} matrix is spectrally different from the AVG → M^3_{CAT} matrix. The heatmap makes the relative statistical similarity among independently converged, averaged, initial, and natural-sequence-derived interaction matrices directly visible while the table retains the exact Z-scores and Pearson correlations.

need not appear identical in their original coordinate representations [30]. The null-model comparison is required because standard correlation metrics can yield high values for complex symmetric matrices even when the apparent agreement is induced by generic spectral structure rather than by a shared physical interaction pattern.

We quantify this equivalence with a two-tier Z-score analysis. The relevant background is the Wigner semicircle distribution [60], which implies that random matrices of this type possess inherent spectral features that can generate nonphysical correlations. We therefore construct independent null ensembles for each matrix comparison, rotate the matrices into the eigenbasis of a reference matrix, and evaluate $Z_{\text{Score}} = [r_{\text{Pearson}}(A^{\text{rot}}, B) - \mu_{\text{null}}]/\sigma_{\text{null}}$. This quantity measures how many standard deviations (σ_{null}) the observed similarity exceeds the random expectation, where $r_{\text{Pearson}}(A^{\text{rot}}, B)$ is the Pearson correlation between the appropriately rotated target matrices A and B, and μ_{null} is the average correlation obtained from the corresponding null ensemble of random matrices. Meaningful similarity is established only when the Z-score exceeds the null expectation ($Z > 1$). The derivation of the null-model construction, eigenbasis rotation, Pearson-based similarity metric, and conservative Z-score procedure is provided in the Supplemental Material [Statistical Significance of Matrix Similarity (Z-Score)] [49].

Figure 4 reports the resulting Z-scores and provides a graphical summary of the same Q-reconstruction comparisons. With the appropriate orthogonal rotation, all Gen 6 (M^6_{CAT}) matrices from the independent paths become statistically equivalent. Most interprotein comparisons yield Z-scores well above 3.0, indicating that these separate target-conditioned paths converge to a shared effective interaction topology rather than merely reflecting stochastic noise. In the context of this statistical test, the high Z-scores are notable because the target proteins differ in fold and secondary-structure composition: after the orthogonal-basis alignment, the common signal is therefore unlikely to be only a target-specific hydrogen-bond pattern. Instead, it is consistent with the Caterpillar separation between directional terms that stabilize local secondary-structure motifs and sequence-dependent matrix terms that encode packing preferences among compatible local arrangements [43–47]. This interpretation explains why independently learned matrices can align across protein families while still preserving the local structural information required by each target. In the specific Caterpillar implementation, this is the key empirical result: each independent refinement trajectory recovers the full interaction solution without requiring ensemble averaging.

The same table also shows that the comparisons between each independent M^6_{CAT} matrix and the averaged matrix are statistically significant (compare rows 1PGB, 1A19, 2PPN, 1NXW, and 1A33 vs the column AVG in Fig. 4). The averaged path therefore reaches the same common converged matrix rather than generating that matrix only through averaging. The negative Z-score between the M^1_{CAT} and M^6_{CAT} matrices (see the AVG → M^1_{CAT} entry in Fig. 4) indicates no significant similarity, which in turn shows that the converged solution is not a trivial recovery of the input matrix.

The identification of this shared effective interaction solution is further supported by initializing the refinement protocol

INDEPENDENT PATHS METHOD						
	1PGB	1A19	2PPN	1NXW	1A33	1MBQ
M^1_{CAT}	0.89	0.86	0.96	1.0	0.97	0.9
M^2_{CAT}	0.96	0.98	0.96	1.0	0.97	0.95
M^3_{CAT}	1.0	1.0	0.98	1.0	1.0	0.93
M^4_{CAT}	1.0	1.0	1.0	1.0	1.0	0.92
M^5_{CAT}	1.0	1.0	1.0	1.0	1.0	0.94
M^6_{CAT}	1.0	1.0	1.0	1.0	1.0	0.9
AVERAGED PATHS METHOD						
	1PGB	1A19	2PPN	1NXW	1A33	1MBQ
M^1_{CAT}	0.82	0.95	0.96	0.98	0.96	0.96
M^2_{CAT}	0.86	0.94	0.93	0.98	0.97	0.94
M^3_{CAT}	0.88	0.98	0.95	0.98	0.97	0.95
M^4_{CAT}	0.92	0.98	0.96	0.99	0.97	0.96
M^5_{CAT}	0.86	0.98	0.98	1.0	0.99	0.97
M^6_{CAT}	0.9	0.98	0.97	1.0	0.99	0.98

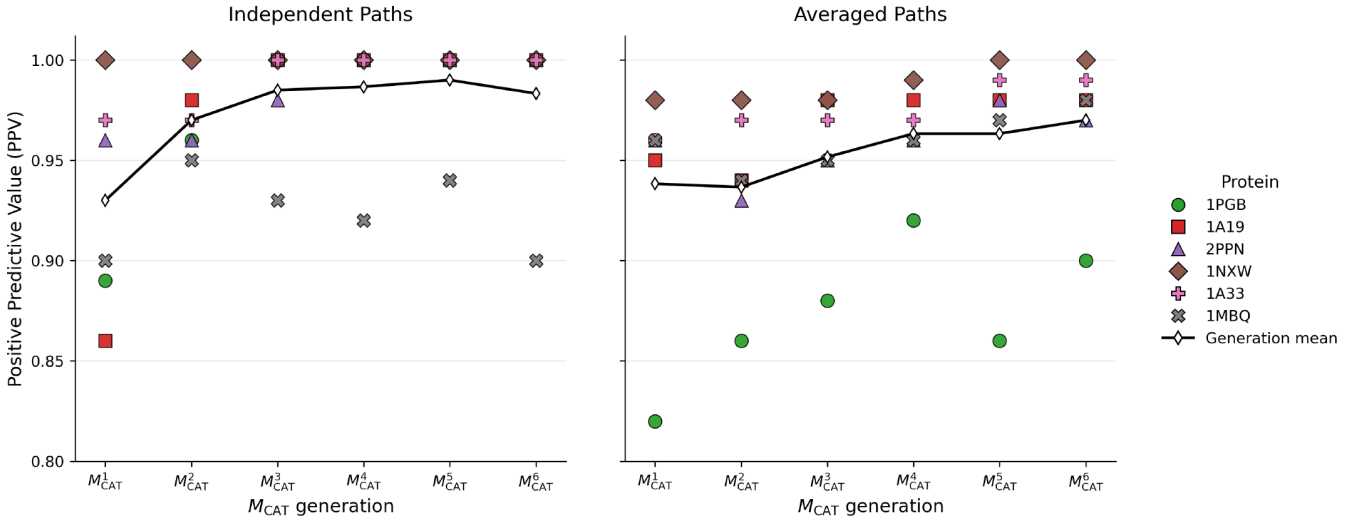


FIG. 5. Positive Predictive Values (PPVs) for DI contact predictions across generations for the top $L/5$ contacts. The upper panel reports the exact PPV values for six proteins across six generations using the Independent Paths and Averaged Paths protocols. The lower panel displays the same generation-resolved values graphically, highlighting the progression toward high contact-prediction accuracy.

with matrices derived from natural sequence data. An initial coupling matrix, NAT, was inferred from natural multiple sequence alignments for the protein families in the dataset [24]; the protein IDs, MSA sizes, and convergence of the natural averaged path are provided in the Supplemental Material (see Matrix obtained from natural sequences [49]). Iterative application of the map $\mathcal{F}$ to this naturally derived matrix produced a convergence trajectory toward a fixed point NAT $\rightarrow M^3_{\mathrm{CAT}}$ (only three iterations needed for convergence; see Fig. S4 [49]), which is statistically equivalent to the model-specific matrix M^6_{CAT} (see Fig. 4).

This comparison has important limitations. No folding validation is performed for the naturally seeded matrices, and no energy-landscape comparison is provided for those trajectories. The present analysis therefore establishes only statistical alignment at the level of the interaction matrix.

Under that limitation, the result suggests that the effective interaction physics encoded by natural sequence statistics is statistically aligned, up to an orthogonal transformation of the interaction basis, with the fixed-point interaction matrix obtained within the Caterpillar framework. In this sense, the natural and Caterpillar-derived interaction matrices are consistent with different coordinate representations of the same underlying interaction structure. This interpretation remains conditional until the naturally seeded potentials are tested by direct folding and landscape analysis.

		Averaged Path		Independent Path	
	M^1_{CAT}	M^3_{CAT}	M^6_{CAT}	M^3_{CAT}	M^6_{CAT}
1PGB	1.7	1.7	1.5	1.7	1.7
1A19	3.3	2.7	1.9	2.7	2.4
2PPN	2.9	1.6	1.8	1.6	1.5
1NXW	3.7	2.6	2.2	2.7	2.3
1A33	1.8	1.8	1.7	2.3	2.2
1MBQ	3.0	3.3	2.8	2.6	2.0

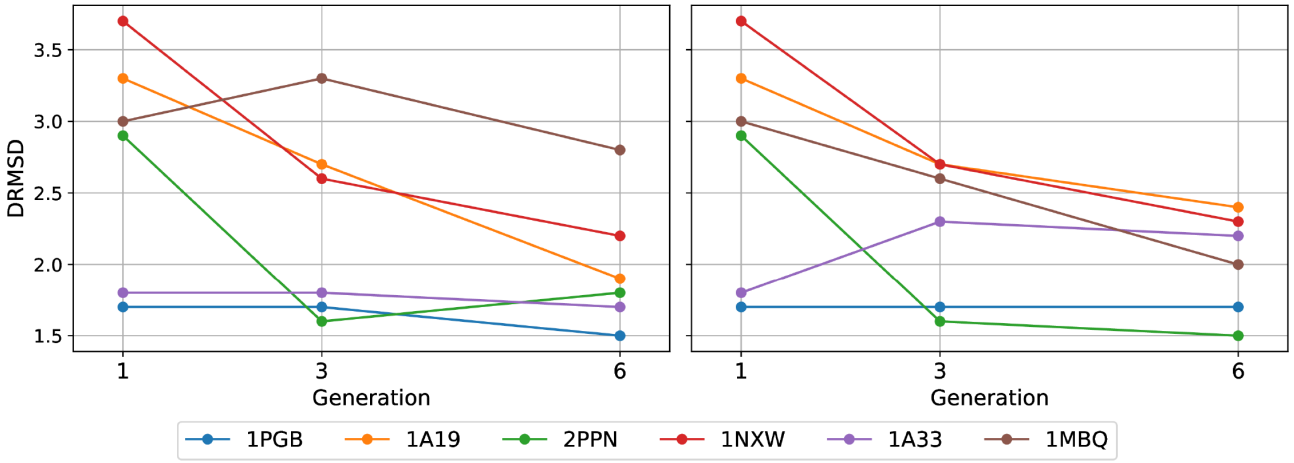


FIG. 6. Distance Root-Mean-Square Deviation (DRMSD) of the free-energy minima for the lowest-free-energy structure sampled at the folding temperature for each simulation and target structure. The upper panel reports the exact DRMSD values. The lower panel illustrates the changes in the energy minima across generations for the averaged path (left) and independent path (right).

C. Superior predictive power and structural fidelity of the converged matrix

1. Improvements on mfDCA predictive power

We next examine the functional consequence of matrix convergence for contact prediction. The improved alignment between M_{DCA} and M_{CAT} matrices is reflected in the predictive power of mfDCA, with a representative example of DI-based contact-map recovery shown in the Supplemental Material (Fig. S6 [49]). Positive predictive values (PPVs) for DI contact predictions were calculated across six generations for six proteins using both Independent Paths and Averaged Paths methods (see Fig. 5). These values provide a direct measure of mfDCA accuracy because they quantify the proportion of predicted residue-residue contacts that are true contacts.

In the first generation, PPV values were already high, ranging from 0.82 to 0.96, which indicates that the initial matrices captured much of the relevant interaction structure. As the iterative process progressed, the PPV values steadily increased, with most proteins achieving near-perfect PPVs (PPV $= 1.0$) by Generation 6 in the Independent Paths method. This pattern shows that greater self-consistency of the interaction matrix is accompanied by more accurate recovery of residue-residue contact information.

The Averaged Paths method showed the same trend, although the improvement was more gradual. By the sixth generation, the averaged method reached PPVs close to 1.0 for most proteins, indicating that the common converged interaction structure remains predictive at the contact level even when learned through the averaged trajectory.

Figure 5 therefore links matrix convergence to contact prediction by combining the exact PPV values with a graphical representation of the generation-resolved trends for the independent and averaged pathways. Improved matrix consistency is reflected in improved PPV, providing the first step in the empirical chain connecting convergence of the interaction matrix to improved design performance.

2. Improved refolding of designed sequences

We next test whether the convergence of the interaction matrix is reflected at the level of folding behavior. To do so, we conducted folding simulations on sequences from Generation 3 and Generation 6 and compared their stability and structural fidelity with those of Generation 1. The relevant observables are the DRMSD values and the shapes of the corresponding free-energy landscapes (see Fig. 6).

Generation 3 folding simulations showed enhanced stability relative to Generation 1, with a more clearly defined funnel-shaped energy landscape (see Supplemental Material: Fig. S5 for representative free-energy landscapes and Figs. S9 and S11 for representative folded structures from Generation 3 independent and averaged paths, respectively [49]). This funnel-like shape indicates that the folding pathways for Generation 3 sequences are more consistently directed toward the target structures. In both the Averaged Paths and

Independent Paths methods, the improved matrix consistency observed during refinement is therefore already reflected in more reliable native-like folding by Generation 3.

By Generation 6, the folding simulations showed a further sharpening of this behavior, with more pronounced funnel-shaped energy landscapes that guide the folding process toward the target structures (see Supplemental Material: Fig. S5 for representative free-energy landscapes and Figs. S10 and S12 for representative folded structures from Generation 6 independent and averaged paths, respectively [49]). This refined landscape corresponds to a high degree of structural fidelity, with Generation 6 sequences folding into stable, native-like conformations that closely match the intended designs. The reduction in DRMSD values at Generation 6 completes the empirical chain established in this section: convergence of the interaction matrix is reflected first in improved PPV and then in improved folding accuracy.

IV. CONCLUSION

In this work we introduced a self-consistent iterative refinement scheme that progressively optimizes residue-residue interaction parameters. At the mathematical level, we established a class-general basis for convergence through the Banach fixed-point theorem. For a given iterative map satisfying the contact-locality and spectral assumptions, the DCA-based update is a contraction for coarse-grained models whose residue-residue energy is a distance-dependent term modulated by a type-dependent coupling matrix. This guarantees convergence to a unique fixed point independent of the initial interaction matrix. The fixed-point property is not specific to Caterpillar, but the location and physical meaning of the fixed point remain specific to the underlying model physics.

In the specific Caterpillar implementation, the empirical question is whether different target-conditioned stochastic refinement trajectories approach the same effective solution. Despite differences in starting conditions, target structures, and Monte Carlo fluctuations, independent paths converge to statistically equivalent interaction matrices, and the averaged path reaches the same solution. This indicates that the shared solution is recovered at the level of individual trajectories rather than produced only by averaging. The result also clarifies what is transferable in the model: the converged matrix should be interpreted primarily as a self-consistent description of sequence-dependent packing preferences, whereas the directional terms of the coarse-grained model supply the local secondary-structure stabilization. This division of roles is consistent with previous work on directional interactions and heteropolymer design [43–47], and explains why a common interaction matrix can support families with different secondary-structure content without implying that the same hydrogen-bond pattern is being transferred. The same convergence is reflected in systematic improvements of PPV and DRMSD metrics across generations, linking convergence of the interaction matrix to enhanced contact prediction and folding accuracy.

The broader interpretation of these results should remain tied to the demonstrated scope of the study. The self-consistent iterative map offers a class-general strategy for refining amino acid energetics across diverse physical models within the same matrix-based coarse-grained framework, and a mathematically grounded route for tuning design models toward foldable sequences. The statistical alignment between the learned interaction matrices and matrices extracted from natural sequence data suggests that a more realistic representation of natural amino acid physics may be related to the present results through a coordinate transformation or matrix rotation. This suggestion remains conditional because no folding validation or energy-landscape comparison is provided for the naturally seeded matrices in the present work. Within those limits, the framework offers a controlled basis for linking sequence statistics to model-specific physical interactions and for clarifying the interaction rules that govern protein design in coarse-grained models.

ACKNOWLEDGMENTS

The work at the Center for Theoretical Biological Physics was sponsored by the NSF (Grants No. PHY-2019745 and No. PHY-2607619), by NSF Grant No. PHY-2210291, and by the Welch Foundation (Grant No. C-1792). J.N.O. is a CPRIT Scholar in Cancer Research sponsored by the Cancer Prevention and Research Institute of Texas. J.M. and F.M. acknowledge support from the National Institutes of Health (Grants No. R35GM133631 and No. R01AI178692). F.M. acknowledges support from the NSF CAREER award (MCB-1943442). The authors acknowledge the High Performance Computing at The University of Texas at Dallas (HPC@UTD) for providing computing resources and support. This work was supported by the FLUFET project, which has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101130125.

Generative artificial intelligence tools were used to assist literature discovery and the organization of the mathematical derivation concerning inverse-covariance decay for sparse contact graphs; all references, assumptions, mathematical statements, and conclusions were independently examined and verified by the authors.

DATA AVAILABILITY

There are no publicly available research data or software supporting this paper. Requests for further information or data should be sent to the authors.

[1] B. I. Dahiyat and S. L. Mayo, De novo protein design: Fully automated sequence selection, Science **278**, 82 (1997).

[2] 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**, 1364 (2003).

[3] G. L. Butterfoss and B. Kuhlman, Computer-based design of novel protein structures, Annu. Rev. Biophys. Biomol. Struct. **35**, 49 (2006).

[4] P.-S. Huang, S. E. Boyken, and D. Baker, The coming of age of *de novo* protein design, Nature (London) **537**, 320 (2016).

[5] E. Marcos and D. Silva, Essentials of *de novo* protein design: Methods and applications, WIREs Comput. Mol. Sci. **8**, e1374 (2018).

[6] L. Sellés Vidal, M. Isalan, J. T. Heap, and R. Ledesma-Amaro, A primer to directed evolution: Current methodologies and future directions, RSC Chem. Biol. **4**, 271 (2023).

[7] K. Yue and K. A. Dill, Inverse protein folding problem: Designing polymer sequences, Proc. Natl. Acad. Sci. USA **89**, 4163 (1992).

[8] E. I. Shakhnovich and A. M. Gutin, Formation of unique structure in polypeptide chains: Theoretical investigation with the aid of a replica approach, Biophys. Chem. **34**, 187 (1989).

[9] J. N. Onuchic and P. G. Wolynes, Theory of protein folding, Curr. Opin. Struct. Biol. **14**, 70 (2004).

[10] W. G. Noid, Perspective: Coarse-grained models for biomolecular systems, J. Chem. Phys. **139**, 90901 (2013).

[11] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, *et al.*, Highly accurate protein structure prediction with AlphaFold, Nature (London) **596**, 583 (2021).

[12] N. Anand, R. R. Eguchi, I. I. Mathews, C. P. Perez, A. Derry, R. B. Altman, and P.-S. Huang, Protein sequence design with a learned potential, Nat. Commun. **13**, 746 (2022).

[13] S. Chaudhury, S. Lyskov, and J. Gray, PyRosetta: A script-based interface for implementing molecular modeling algorithms using Rosetta, Bioinformatics **26**, 689 (2010).

[14] T. Zheng, A. Rondina, G. Micklem, and P. Liò, Challenges and guidelines in deep generative protein design: Four case studies, in *Proceedings of the ICML 2025 Workshop on Multi-modal Foundation Models and Large Language Models for Life Sciences*, Vancouver, Canada (2025), arXiv:2411.18568.

[15] R. R. Stein, D. S. Marks, and C. Sander, Inferring pairwise interactions from biological data using maximum-entropy probability models, PLoS Comput. Biol. **11**, e1004182 (2015).

[16] 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**, 032601 (2018).

[17] R. M. Levy, A. Haldane, and W. F. Flynn, Potts Hamiltonian models of protein co-variation, free energy landscapes, and evolutionary fitness, Curr. Opin. Struct. Biol. **43**, 55 (2017).

[18] C. Malbranke, D. Bikard, S. Cocco, R. Monasson, and J. Tubiana, Machine learning for evolutionary-based and physics-inspired protein design: Current and future synergies, Curr. Opin. Struct. Biol. **80**, 102571 (2023).

[19] P. G. Wolynes, As simple as can be? Nat. Struct. Biol. **4**, 871 (1997).

[20] J. Dauparas, I. Anishchenko, N. Bennett, H. Bai, R. J. Ragotte, L. F. Milles, B. I. M. Wicky, A. Courbet, R. J. de Haas, N. Bethel, P. J. Y. Leung, T. F. Huddy, S. Pellock, D. Tischer, F. Chan, B. Koepnick, H. Nguyen, A. Kang, B. Sankaran, A. K. Bera, *et al.*, Others robust deep protein sequence design using ProteinMPNN, Science **378**, 49 (2022).

[21] H. Khakzad, I. Igashov, A. Schneuing, C. Goverde, M. Bronstein, and B. Correia, A new age in protein design empowered by deep learning, Cell Syst. **14**, 925 (2023).

[22] U. Göbel, C. Sander, R. Schneider, and A. Valencia, Correlated mutations and residue contacts in proteins, Proteins: Struct. Funct. Bioinf. **18**, 309 (1994).

[23] 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. USA **106**, 67 (2009).

[24] F. Morcos, A. Pagnani, B. Lunt, A. Bertolino, D. S. Marks, C. Sander, R. Zecchina, J. N. Onuchic, T. Hwa, and M. Weigt, Direct-coupling analysis of residue coevolution captures native contacts across many protein families, Proc. Natl. Acad. Sci. USA **108**, E1293 (2011).

[25] 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**, e28766 (2011).

[26] A. S. Lapedes, B. G. Giraud, and C. Jarzynski, Using sequence alignments to predict protein structure and stability with high accuracy, Los Alamos Report No. LA-UR-02-4481 (2002), arXiv:1207.2484.

[27] M. Ekeberg, C. Lövkvist, Y. Y. Lan, M. Weigt, and E. Aurell, Improved contact prediction in proteins: Using pseudolikelihoods to infer Potts models, Phys. Rev. E **87**, 012707 (2013).

[28] H. Kamisetty, S. Ovchinnikov, and D. Baker, Assessing the utility of coevolution-based residue–residue contact predictions in a sequence- and structure-rich era, Proc. Natl. Acad. Sci. USA **110**, 15674 (2013).

[29] H. Jacquin, A. Gilson, E. I. Shakhnovich, S. Cocco, and R. Monasson, Test and benchmark of statistical inverse methods for protein structure prediction, PLoS Comput. Biol. **12**, e1004889 (2016).

[30] J. Martin, M. Lequerica Mateos, J. N. Onuchic, I. Coluzza, and F. Morcos, Machine learning in biological physics: From biomolecular prediction to design, Proc. Natl. Acad. Sci. USA **121**, e2311807121 (2024).

[31] F. M. Ortega, F. Hossain, V. V. Volobouev, G. Meloni, H. Torabifard, and F. Morcos, Generative landscapes and dynamics to design functional multidomain artificial transmembrane transporters, ACS Cent. Sci. **11**, 1452 (2025).

[32] S. W. Lockless and R. Ranganathan, Evolutionarily conserved pathways of energetic connectivity in protein families, Science **286**, 295 (1999).

[33] N. Halabi, O. Rivoire, S. Leibler, and R. Ranganathan, Protein sectors: Evolutionary units of three-dimensional structure, Cell **138**, 774 (2009).

[34] 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**, 268 (2016).

[35] T. A. Hopf, J. B. Ingraham, F. J. Poelwijk, C. P. I. Schärfe, M. Springer, C. Sander, and D. S. Marks, Mutation effects predicted from sequence co-variation, Nat. Biotechnol. **35**, 128 (2017).

[36] W. P. Russ, M. Figliuzzi, C. Stocker, P. Barrat-Charlaix, M. Socolich, P. Kast, D. Hilvert, R. Monasson, S. Cocco, M. Weigt, and R. Ranganathan, An evolution-based model for designing chorismate mutase enzymes, Science **369**, 440 (2020).

[37] V. Tozzini, Coarse-grained models for proteins, Curr. Opin. Struct. Biol. **15**, 144 (2005).

[38] C. Clementi, Coarse-grained models of protein folding: Toy models or predictive tools? Curr. Opin. Struct. Biol. **18**, 10 (2008).

[39] S. Kmiecik, D. Gront, M. Kolinski, L. Wieteska, A. E. Dawid, and A. Kolinski, Coarse-grained protein models and their applications, Chem. Rev. **116**, 7898 (2016).

[40] A. Davtyan, N. P. Schafer, W. Zheng, C. Clementi, P. G. Wolynes, and G. A. Papoian, AWSEM-MD: Protein structure prediction using coarse-grained physical potentials and bioinformatically based local structure biasing, J. Phys. Chem. B **116**, 8494 (2012).

[41] J. K. Noel, F. Morcos, and J. N. Onuchic, Sequence co-evolutionary information is a natural partner to minimally-frustrated models of biomolecular dynamics, F1000Research **5**, 106 (2016).

[42] I. Coluzza, Transferable coarse-grained potential for *de novo* protein folding and design, PLoS ONE **9**, e112852 (2014).

[43] C. Cardelli, V. Bianco, L. Rovigatti, F. Nerattini, L. Tubiana, C. Dellago, and I. Coluzza, The role of directional interactions in the designability of generalized heteropolymers, Sci. Rep. **7**, 4986 (2017).

[44] C. Cardelli, L. Tubiana, V. Bianco, F. Nerattini, C. Dellago, and I. Coluzza, Heteropolymer design and folding of arbitrary topologies reveals an unexpected role of alphabet size on the knot population, Macromolecules **51**, 8346 (2018).

[45] C. Cardelli, F. Nerattini, L. Tubiana, V. Bianco, C. Dellago, F. Sciortino, and I. Coluzza, General methodology to identify the minimum alphabet size for heteropolymer design, Adv. Theory Simul. **2**, 1900031 (2019).

[46] I. Coluzza, P. D. J. van Oostrum, B. Capone, E. Reimhult, and C. Dellago, Design and folding of colloidal patchy polymers, Soft Matter **9**, 938 (2012)

[47] I. Coluzza, P. D. J. van Oostrum, B. Capone, E. Reimhult, and C. Dellago, Sequence controlled self-knotting colloidal patchy polymers, Phys. Rev. Lett. **110**, 075501 (2013).

[48] I. Coluzza, A coarse-grained approach to protein design: Learning from design to understand folding, PLoS ONE **6**, e20853 (2011).

[49] See Supplemental Material at http://link.aps.org/supplemental/10.1103/dmcv-sx8q for model and sampling details, mathematical derivations, and supplementary analyses of matrix convergence and protein folding. It also provides the detailed mathematical and computational framework underpinning the study, including the Caterpillar Hamiltonian, sequence-sampling protocol, projection of position-dependent mfDCA couplings onto the amino-acid interaction matrix, and the full Jacobian derivation. It states the contact-locality, spectral, and residual-localization assumptions used to obtain sparse-graph inverse-covariance decay, compact-globule sequence localization, and the resulting contraction bound. It also contains the matrix-alignment analysis, natural-sequence construction, folding free-energy landscapes, contact-map recovery, and DRMSD comparisons across refinement pathways, which includes Refs. [61–69].

[50] M. Lequerica-Mateos, J. Martin, J. N. Onuchic, F. Morcos, and I. Coluzza, Designing proteins: Mimicking natural protein sequence heterogeneity, J. Chem. Phys. **161**, 194102 (2024).

[51] Z. Bagci, R. L. Jernigan, and I. Bahar, Residue coordination in proteins conforms to the closest packing of spheres, Polymer **43**, 451 (2002).

[52] A. R. Atilgan, P. Akan, and C. Baysal, Small-world communication of residues and significance for protein dynamics, Biophys. J. **86**, 85 (2004).

[53] S. Jaffard, Propriétés des matrices 'bien localisées' près de leur diagonale et quelques applications, Ann. Inst. Henri Poincaré C Anal. Non Linéaire **7**, 461 (1990).

[54] S. Demko, Inverses of band matrices and local convergence of spline projections, SIAM J. Numer. Anal. **14**, 616 (1977).

[55] S. Demko, W. F. Moss, and P. W. Smith, Decay rates for inverses of band matrices, Math. Comput. **43**, 491 (1984).

[56] M. Benzi and N. Razouk, Decay bounds and $O(n)$ algorithms for approximating functions of sparse matrices, Electron. Trans. Numer. Anal. **28**, 16 (2007).

[57] A. Y. Grosberg, S. K. Nechaev, and E. I. Shakhnovich, The role of topological constraints in the kinetics of collapse of macromolecules, J. Physique **49**, 2095 (1988).

[58] T. G. Dewey, Protein structure and polymer collapse, J. Chem. Phys. **98**, 2250 (1993).

[59] M. V. Imakaev, K. M. Tchourine, S. K. Nechaev, and L. A. Mirny, Effects of topological constraints on globular polymers, Soft Matter **11**, 665 (2015).

[60] E. P. Wigner, On the distribution of the roots of certain symmetric matrices, Ann. Math. **67**, 325 (1958).

[61] T. E. Creighton, *Proteins: Structures and Molecular Properties*, 2nd ed. (W. H. Freeman, New York, 1993).

[62] A. Irbäck, F. Sjunnesson, and S. Wallin, Three-helix-bundle protein in a Ramachandran model, Proc. Natl. Acad. Sci. USA **97**, 13614 (2000).

[63] S. Miyazawa and R. L. Jernigan, Estimation of effective interresidue contact energies from protein crystal structures: Quasi-chemical approximation, Macromolecules **18**, 534 (1985).

[64] J. Kyte and R. F. Doolittle, A simple method for displaying the hydropathic character of a protein, J. Mol. Biol. **157**, 105 (1982).

[65] I. Coluzza and D. Frenkel, Virtual-move parallel tempering, ChemPhysChem **6**, 1779 (2005).

[66] G. C. Boulougouris and D. Frenkel, Novel Monte Carlo scheme for systems with short-ranged interactions, J. Chem. Phys. **122**, 244106 (2005).

[67] G. C. Boulougouris and D. Frenkel, Monte Carlo sampling of a Markov web, J. Chem. Theory Comput. **1**, 389 (2005).

[68] R. A. Fisher, Frequency distribution of the values of the correlation coefficient in samples of an indefinitely large population, Biometrika **10**, 507 (1915).

[69] D. T. Jones, S. M. Kandathil, and A. Valencia, High precision in protein contact prediction using fully convolutional neural networks and minimal sequence features, Bioinformatics **34**, 3308 (2018).