Revealing evolutionary constraints on proteins through sequence analysis

Shou-Wen Wang[1,2,3◕¤], Anne-Florence Bitbol[4◕*], Ned S. Wingreen[3,5*]

1 Department of Engineering Physics, Tsinghua University, Beijing, 100086, China
2 Beijing Computational Science Research Center, Beijing, 100094, China
3 Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, NJ 08544, USA
4 Sorbonne Université, CNRS, Laboratoire Jean Perrin (UMR 8237), F-75005 Paris, France
5 Department of Molecular Biology, Princeton University, Princeton, NJ 08544, USA

◕These authors contributed equally to this work.
¤Current Address: Department of Systems Biology, Harvard Medical School, Boston, MA 02115, USA.
* anne-florence.bitbol@sorbonne-universite.fr (A.-F. B.), wingreen@princeton.edu (N. S. W.)

Abstract

Statistical analysis of alignments of large numbers of protein sequences has revealed "sectors" of collectively coevolving amino acids in several protein families. Here, we show that selection acting on any functional property of a protein, represented by an additive trait, can give rise to such a sector. As an illustration of a selected trait, we consider the elastic energy of an important conformational change within an elastic network model, and we show that selection acting on this energy leads to correlations among residues. For this concrete example and more generally, we demonstrate that the main signature of functional sectors lies in the small-eigenvalue modes of the covariance matrix of the selected sequences. However, secondary signatures of these functional sectors also exist in the extensively-studied large-eigenvalue modes. Our simple, general model leads us to propose a principled method to identify functional sectors, along with the magnitudes of mutational effects, from sequence data. We further demonstrate the robustness of these functional sectors to various forms of selection, and the robustness of our approach to the identification of multiple selected traits.

Author summary

Proteins play crucial parts in all cellular processes, and their functions are encoded in their amino-acid sequences. Recently, statistical analyses of protein sequence alignments have demonstrated the existence of "sectors" of collectively correlated amino acids. What is the origin of these sectors? Here, we propose a simple underlying origin of protein sectors: they can arise from selection acting on any collective protein property. We find that the main signature of these functional sectors lies in the low-eigenvalue modes of the covariance matrix of the selected sequences. A better understanding of protein sectors will make it possible to discern collective protein properties directly from sequences, as well as to design new functional sequences, with far-reaching applications in synthetic biology.

Introduction

Proteins play crucial roles in all cellular processes, acting as enzymes, motors, receptors, regulators, and more. The function of a protein is encoded in its amino-acid sequence. In evolution, random mutations affect the sequence, while natural selection acts at the level of function, however our ability to predict a protein's function directly from its sequence has been very limited. Recently, the explosion of available sequences has inspired new data-driven approaches to uncover the principles of protein operation. At the root of these new approaches is the observation that amino-acid residues which possess related functional roles often evolve in a correlated way. In particular, analyses of large alignments of protein sequences have identified "sectors" of collectively correlated amino acids [1–6], which has enabled successful design of new functional sequences [3]. Sectors are spatially contiguous in the protein structure, and in the case of multiple sectors, each one may be associated with a distinct role [4, 7]. What is the origin of these sectors, and can we identify them from sequence data in a principled way?

To address these questions, we developed a general physical model that naturally gives rise to sectors. Specifically, motivated by the observation that many protein properties reflect additive contributions from individual amino acids [8–10], we consider any additive trait subject to natural selection. As a concrete example, we study a simple elastic-network model that quantifies the energetic cost of protein deformations [11], which we show to be an additive trait. We then demonstrate that selection acting on any such additive trait automatically yields collective correlation modes in sequence data. We show that the main signature of the selection process lies in the small-eigenvalue modes of the covariance matrix of the selected sequences, but we find that some signatures also exist in the widely-studied large-eigenvalue modes. Finally, we demonstrate a principled method to identify sectors and to quantify mutational effects from sequence data alone.

Results

Selection on an additive trait

We focus on selection on an additive scalar trait

$$T(\vec{\alpha}) = \sum_{l=1}^{L} \Delta_l(\alpha_l)\,, \tag{1}$$

where $\vec{\alpha} = (\alpha_1, \ldots, \alpha_L)$ is the amino-acid sequence considered, L is its length, and $\Delta_l(\alpha_l)$ is the mutational effect on the trait T of a mutation to amino acid α_l at site l. Mutational effects can be measured with respect to a reference sequence $\vec{\alpha}^0$, satisfying $\Delta_l(\alpha_l^0) = 0$ for all l.

Eq. 1 is very general as it amounts to saying that, to lowest order, mutations have an additive effect on the trait T, which can be any relevant physical property of the protein, say its binding affinity, catalytic activity, or thermal stability [12]. System-specific details are encoded by the single-site mutational effects $\Delta_l(\alpha_l)$, which can be measured experimentally. The assumption of additivity is experimentally validated in many cases. For instance, protein thermal stability, measured through folding free energy, is approximately additive [8, 13]. Importantly, we allow selection to act on a phenotype that is a nonlinear function of T. Permitting a phenotypic nonlinearity on top of our additive trait model is motivated by the fact that actual phenotype data from recent high-throughput mutagenesis experiments were accurately modeled via a nonlinear mapping of an underlying additive trait [10].

Protein sectors are usually defined operationally as collective modes of correlations in amino-acid sequences. However, the general sequence-function relation in Eq. 1 suggests an operational definition of a *functional* protein sector, namely as the set of sites with dominant mutational effects on a trait under selection. Selection can take multiple forms. To be concrete, we first consider a simple model of selection, assuming a favored value T^* of the trait T, and using a Gaussian selection window. We subsequently show that the conclusions obtained within this simple model are robust to different forms of selection. Our Gaussian selection model amounts to selecting sequences according to the following Boltzmann distribution:

$$P(\vec{\alpha}) = \frac{\exp(w(\vec{\alpha}))}{\sum_{\vec{\alpha}} \exp(w(\vec{\alpha}))}\,, \tag{2}$$

where the fitness $w(\vec{\alpha})$ of a sequence is given by

$$w(\vec{\alpha}) = -\frac{\kappa}{2}\left(T(\vec{\alpha}) - T^*\right)^2 = -\frac{\kappa}{2}\left(\sum_{l=1}^{L} \Delta_l(\alpha_l) - T^*\right)^2 . \tag{3}$$

The selection strength κ sets the width of the selection window.

Such selection for intermediate values of a trait can be realistic, e.g. for protein stability [8]. However, the form of selection can vary, for example selection can be for a nonlinear transform of a trait to be above a certain threshold [10], and several relevant selection variants are investigated below. Crucially, while the trait is additive (Eq. 1), the fact that fitness (Eq. 3) and selection (Eq. 2) are nonlinear functions of the trait leads to coupling between mutations. This phenomenon is known as global [10, 14] or nonspecific [9] epistasis, and its relevance has been shown in evolution experiments [14], over and above contributions from specific epistasis [9, 15]. The focus of this paper is on global epistasis, and we do not include specific epistasis. Studying the interplay of these two types of epistasis will be an interesting future direction.

A toy model yielding a concrete example of an additive trait

Elastic-network model

To illustrate how additive traits naturally arise, we consider the elastic energy associated with a functionally important protein deformation. We explicitly derive the additivity of this trait in the regime of small deformations and weak mutational effects. This concrete example is relevant since functional deformation modes are under selection in proteins [16–18], and dynamical domains possess a signature in sequence data [19]. Moreover, elastic-network models have elucidated a variety of protein properties [11,20–22], including the emergence of allostery [23–29]. Thus motivated, we begin by building an elastic-network model [11,20] for a well-studied PDZ protein domain (Fig. 1(a,b)) [30,31] and computing the relationship between its "sequence" and the energetic cost of a functionally-relevant conformational change.

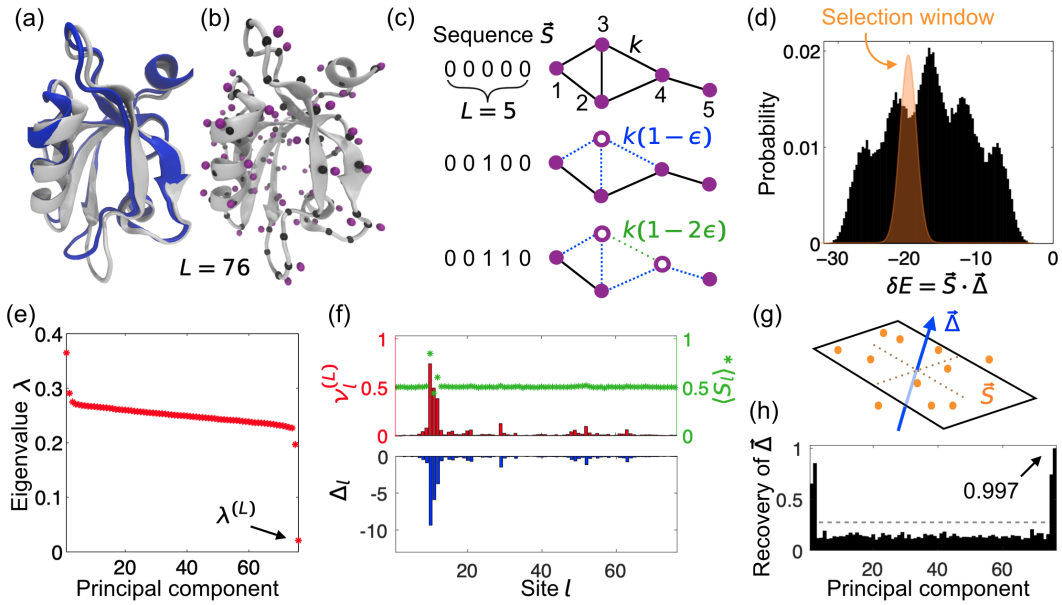

Fig 1. Selection applied to an elastic protein model leads to a statistical signature among sequences. (a) Cartoon representation of the third PDZ domain of the rat postsynaptic density protein 95 from the RSCB PDB [32], (gray: ligand free, 1BFE; blue: ligand bound, 1BE9 (ligand not shown)). (b) Elastic network model for 1BFE, where each amino-acid residue is represented by its alpha carbon (Cα, black node) and beta carbon (Cβ, purple node). Nearby nodes interact through a harmonic spring [20] (S1 Appendix). (c) Relation between protein sequence $\vec{S}$ and elastic network: 0 denotes the reference state, while 1 denotes a mutated residue, which weakens interactions of the corresponding Cβ with all its neighbors by ϵ. (d) Histogram of the energy δE required to deform the domain from its ligand-free to its ligand-bound conformation, for randomly sampled sequences where 0 and 1 are equally likely at each site. Sequences are selectively weighted using a narrow Gaussian window (orange) around δE^*. (e) Eigenvalues of the covariance matrix C for the selectively weighted protein sequences. (f) Upper panel: last principal component $\nu_l^{(L)}$ of C (red) and average mutant fraction $\langle S_l \rangle_*$ (green) at site l after selection; lower panel: effect Δ_l of a single mutation at site l on δE. (g) Schematic representation of the selected ensemble in sequence space, where each dot is a highly-weighted sequence; thus dots are restricted to a narrow region around a plane perpendicular to $\vec{\Delta}$. (h) Recovery of $\vec{\Delta}$ for all principal components $\vec{\nu}^{(j)}$, with maximum Recovery=1 (Eq. 8). Gray dashed line: random expectation of Recovery (S1 Appendix).

To build the elastic-network model of the PDZ domain, we replace each of the $L = 76$ amino-acid residues by its corresponding alpha carbon Cα and beta carbon Cβ, as shown in Fig. 1(b). Every pair of carbons within a cutoff distance d_c is then connected with a harmonic spring [11]. Following a previous analysis of the same PDZ domain [20], we set $d_c = 7.5$ Å and assign spring constants as follows: a) 2 for Cα-Cα pairs if adjacent along the backbone, 1 otherwise; b) 1 for Cα-Cβ pairs; c) 0.5 for Cβ-Cβ pairs.

Within our elastic model, the energetic cost of a small deformation from the equilibrium structure is

$$E = \frac{1}{2} \sum_{i,j} \left(\boldsymbol{r}_i - \boldsymbol{r}_i^0 \right) M_{ij} \left(\boldsymbol{r}_j - \boldsymbol{r}_j^0 \right) = \frac{1}{2} \, \delta \boldsymbol{r}^T M \delta \boldsymbol{r}, \tag{4}$$

where $\boldsymbol{r}_i$ is the position of the ith carbon atom, $\boldsymbol{r}_i^0$ is its equilibrium position, and the Hessian matrix M contains the second derivatives of the elastic energy with respect to atomic coordinates. Here, we take $\delta \boldsymbol{r}$ to be the conformational change from a ligand-free state (1BFE) to a ligand-bound state (1BE9) of the same PDZ domain (Fig. 1(a)). This conformational change is central to PDZ function, so its energetic cost has presumably been under selection during evolution. Any other coherent conformational change would also be suitable for our analysis. Note that our aim is not to analyze conformational changes in all their richness, but to provide a minimal concrete example of a relevant additive trait, and to analyze the impact of selection acting on this trait on the associated family of sequences.

To mimic simply the effect of mutation and selection within our toy model, we introduce "mutations" of residues that weaken the spring constants involving their beta carbons by a small fraction ϵ. In practice, we take $\epsilon = 0.2$. We represent mutations using a sequence $\vec{S}$ with $S_l \in \{0, 1\}$, where l is the residue index: $S_l = 0$ denotes the reference state, while $S_l = 1$ implies a mutation (Fig. 1(c)). The sequence $\vec{S}$ and the spring network fully determine the Hessian matrix M, and thus the energy cost E of a conformational change (Eq. 4). Note that here $\vec{S}$ is a binary sequence, which represents a simplification compared to real protein sequences $\vec{\alpha}$ where each site can feature 21 states (20 amino acids, plus the alignment gap). We start with the binary model for simplicity, and we then extend our results to a more realistic 21-state model. Note that binary representations of actual protein sequences, with a consensus residue state and a "mutant" state, have proved useful in sector analysis [4], although more recent approaches for diverse protein families have employed full 21-state models [7]. Binary representations are also appropriate to analyze sets of sufficiently close sequences, notably HIV proteins, allowing identification of their sectors [5] and predictions of their fitness landscapes [33].

Deformation energy as an additive trait

Focusing on mutations that weakly perturb the elastic properties of a protein, we perform first-order perturbation analysis: $M = M^{(0)} + \epsilon M^{(1)} + o(\epsilon)$. Using Eq. 4 yields $E = E^{(0)} + \epsilon E^{(1)} + o(\epsilon)$, with $E^{(1)} = \delta \boldsymbol{r}^T M^{(1)} \delta \boldsymbol{r}/2$. Both $M^{(1)}$ and $E^{(1)}$ can be expressed as sums of contributions from individual mutations. We define Δ_l as the first-order energy cost $\epsilon E^{(1)}$ of a single mutation at site l of the sequence. To leading order, the effect of mutations on the energy cost of a deformation reads

$$\delta E = E - E^{(0)} = \sum_{l=1}^{L} S_l \Delta_l. \tag{5}$$

This equation corresponds to the binary-sequence case of the general additive trait defined in Eq. 1. Hence, the deformation energy in our toy model of a protein as a sequence-dependent elastic network constitutes a practical example of an additive trait.

Within our functional definition, a protein sector is the set of sites with dominant mutational effects on the trait under selection. The vector $\vec{\Delta}$ of mutational effects for our elastic-network model of the PDZ domain is shown in Fig. 1(f). The magnitudes of mutational effects are strongly heterogeneous (S1 Appendix, Fig. S1). Here, the amino acids with largest effects, which constitute the sector, correspond to those that move most upon ligand binding. (Note that the ligand-binding deformation of PDZ is well-described by one low-frequency normal mode of the elastic network [20]: hence, our sector significantly overlaps with the sites that are most involved in this mode.)

How is such a functionally-defined sector reflected in the statistical properties of the sequences that survive evolution? To answer this question, we next analyze sequences obtained by selecting on the trait δE. While for concreteness, we use the mutational effects obtained from our elastic model, the analysis is general and applies to any additive trait. Indeed, we later present some examples using synthetically-generated random mutational effect vectors, both binary and more realistic 21-state ones (see Figs. 3, 4, and S1 Appendix).

Signature of selection in sequences

For our elastic model of the PDZ domain, the distribution of the additive trait δE for random sequences is shown in Fig. 1(d). We use the selection process introduced in Eqs. 2-3 to limit sequences to a narrower distribution of δEs, corresponding, e.g., to a preferred ligand-binding affinity. The fitness of a binary sequence $\vec{S}$, a particular case of Eq. 3, reads:

$$w(\vec{S}) = -\frac{\kappa}{2}\left(\sum_{l=1}^{L}\Delta_l S_l - \delta E^*\right)^2 . \tag{6}$$

Here, the selection strength κ sets the width of the selection window, and δE^* is its center. For all selections, we take $\kappa = 10/(\sum_l \Delta_l^2)$, so that the width of the selection window scales with that of the unselected distribution. We have confirmed that our conclusions are robust to varying selection strength, provided $\kappa \sum_l \Delta_l^2 \gg 1$ (see Fig. S3).

Although mutations have additive effects on the trait δE, the nonlinearities involved in fitness and selection give rise to correlations among sites. For instance, if $\delta E^* = 0$ and if $\Delta_l < 0$ for all l, as in Fig. 1, a mutation at a site with large $|\Delta_l|$ will decrease the likelihood of additional mutations at all other sites with large $|\Delta_l|$.

Previous approaches to identifying sectors from real protein sequences have relied on modified forms of Principal Component Analysis (PCA). So we begin by asking: can PCA identify sectors in our physical model? PCA corresponds to diagonalizing the covariance matrix C of sequences: it identifies the principal components (eigenvectors) $\vec{\nu}^{(j)}$ associated with progressively smaller variances (eigenvalues) $\lambda^{(j)}$. We introduce $\langle\cdot\rangle_*$ to denote ensemble averages over the selectively weighted sequences, reserving $\langle\cdot\rangle$ for averages over the unselected ensemble. The mutant fraction at site l in the selected ensemble is $\langle S_l\rangle_* = \sum_{\vec{S}} S_l P(\vec{S})$, and the covariance matrix C reads

$$C_{ll'} = \Big\langle (S_l - \langle S_l\rangle_*)\cdot(S_{l'} - \langle S_{l'}\rangle_*)\Big\rangle_* . \tag{7}$$

To test the ability of PCA to identify a functional sector, we employed the selection window shown in orange in Fig. 1(d). The resulting eigenvalues are shown in Fig. 1(e). One sees outliers. In particular, why is the last eigenvalue so low? Due to the narrow selection window, according to Eq. 6 the highly-weighted sequences satisfy $\sum_l S_l\Delta_l = \vec{S}\cdot\vec{\Delta} \approx \delta E^*$. This means that in the L-dimensional sequence space, the data points for the highly-weighted sequences lie in a narrow region around a plane perpendicular to $\vec{\Delta}$ (Fig. 1(g)). Hence, the data has exceptionally small variance in this direction, leading to a particularly small eigenvalue of C. Moreover, the corresponding last principal component $\vec{\nu}^{(L)}$ points in the direction with the smallest variance and is consequently parallel to $\vec{\Delta}$ (Fig. 1(f)). Formally, in Eq. 6, $\vec{\Delta}$ appears in a quadratic coupling term where it plays the part of a repulsive pattern in a generalized Hopfield model [34,35]: alone, such a term would penalize sequences aligned with $\vec{\Delta}$. But here, $\vec{\Delta}$ also appears in a term linear in $\vec{S}$ and as a result Eq. 6 penalizes sequences that fail to have the selected projection onto $\vec{\Delta}$.

In this example, the last principal component accurately recovers the functional sector corresponding to the largest elements of the mutational-effect vector $\vec{\Delta}$. More generally, to quantify the recovery of $\vec{\Delta}$ by a given vector $\vec{\nu}$, we introduce

$$\text{Recovery} = \frac{\sum_l |\nu_l\Delta_l|}{\sqrt{\sum_l \nu_l^2}\sqrt{\sum_l \Delta_l^2}}, \tag{8}$$

which is nonnegative, has a random expectation of $(\sqrt{2/\pi L})\sum_l |\Delta_l|/\sqrt{\sum_l \Delta_l^2}$ for $L \gg 1$ (S1 Appendix), and saturates at 1 (including the case of parallel vectors). For our test case, Fig. 1(h) shows Recovery for all principal components. The last one features the highest Recovery, almost 1, confirming that it carries substantial information about $\vec{\Delta}$. The second-to-last principal component and the first two also provide a value of Recovery substantially above random expectation. Outlier eigenvalues arise from the sector, and accordingly, we find that the number of modes with high Recovery often corresponds to the number of sites with strong mutational effects. A more formal analysis of this effect will be an interesting topic for further study.

In our model, $\vec{\Delta}$ is fundamentally a direction of *small variance*. So why do the first principal components also carry information about $\vec{\Delta}$? Qualitatively, when variance is decreased in one direction due to a repulsive pattern $\vec{\Delta}$, variance tends to increase in orthogonal directions involving the same sites. To illustrate this effect, let $L = 3$ and $\vec{\Delta} = (-1, 1, 0)$, and consider the sequences $\vec{S}$ satisfying $\vec{\Delta}\cdot\vec{S} = 0$ (namely $(0,0,0)$; $(1,1,0)$; $(0,0,1)$; $(1,1,1)$). The last principal component is $\vec{\Delta}$, with zero variance, and the first principal component is $(1,1,0)$: Recovery is 1 for both of them. This selection conserves the trace of the covariance matrix (i.e. the total variance), so that decreasing the variance along $\vec{\Delta} = (-1,1,0)$ necessarily increases it along $(1,1,0)$. This simple example provides an intuitive understanding of why the large-eigenvalue modes of the covariance matrix also carry information about $\vec{\Delta}$.

It is worth remarking that Eq. 6 is a particular case of a general fitness function with one- and two-body terms (known as fields and couplings in Ising or Potts models in physics). Here, the values of these one- and two-body terms are constrained by their expressions in terms of $\vec{\Delta}$. In practice, several traits might be selected simultaneously (see below), yielding more independent terms among the fields and couplings. More generally, such one- and two-body descriptions have been very successfully employed via Direct Coupling Analysis (DCA) to identify strongly coupled residues that are in contact within a folded protein [36–38], to investigate folding [39], and to predict fitness [33, 40–45] and conformational changes [46, 47], as well as protein-protein interactions [48, 49]. A complete model of protein covariation in nature should necessarily incorporate both the collective modes described here and the strongly coupled residue pairs which are the focus of DCA.

ICOD method

An important concern is whether the last principal component is robust to small and/or noisy datasets. Indeed, other directions of small variance can appear in the data. As a second example, we applied a different selection window, centered in the tail of the distribution of δEs from our elastic model of the PDZ domain (Fig. 2(a), inset). This biased selection generates strong conservation, $\langle S_l \rangle_* \approx 1$, for some sites with significant mutational effects. Extreme conservation at one site now dictates the last principal component, and disrupts PCA-based recovery of $\vec{\Delta}$ (Fig. 2(a,b)).

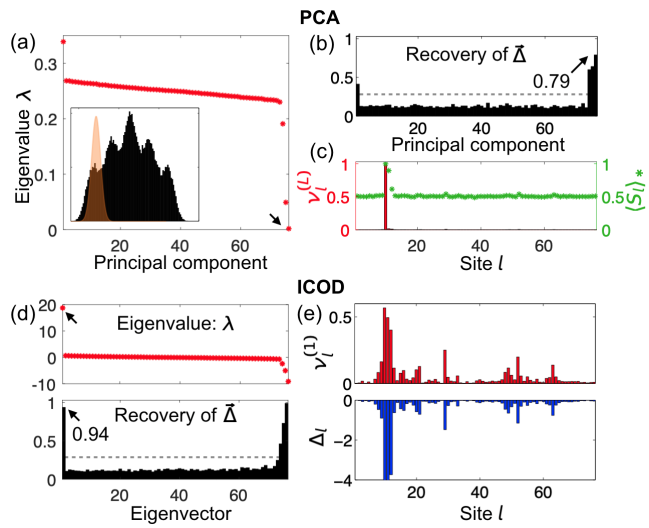

Fig 2. Recovery of mutational-effect vector $\vec{\Delta}$ from sequence analysis in the case of strongly biased selection. (a-c) Principal Component Analysis (PCA) performs poorly due to strong conservation at some sites of large mutational effect. (a) Eigenvalues of covariance matrix obtained for strongly biased selection around $\delta E^*_{\text{biased}}$ (inset, orange window) for same model proteins as in Fig. 1. (b) Recovery of $\vec{\Delta}$ for all principal components. (c) Last principal component $\nu_l^{(L)}$ (red) and average mutant fraction $\langle S_l \rangle_*$ (green) at site l. (d-e) The ICOD method performs robustly. (d) Eigenvalues of $\tilde{C}^{-1}_{ll'}$ (Eq. 10) (upper) and Recovery of $\vec{\Delta}$ for all eigenvectors (lower). (e) Leading eigenvector $\nu_l^{(1)}$ (upper) and mutational effect Δ_l at site l (lower, same as in Fig. 1(f)). Gray dashed lines in (b,d): random expectation of Recovery (S1 Appendix).

To overcome this difficulty, we developed a more robust approach that relies on inverting the covariance matrix. Previously, the inverse covariance matrix was successfully employed in Direct Coupling Analysis (DCA) to identify strongly coupled residues that are in contact within a folded protein [36–38]. The fitness in our model (Eq. 6) involves one and two-body interaction terms, and constitutes a particular case of the DCA Hamiltonian (S1 Appendix). A small-coupling approximation [37, 38, 50, 51] (S1 Appendix) gives

$$C^{-1}_{ll'} \approx (1 - \delta_{ll'})\, \kappa \Delta_l \Delta_{l'} + \delta_{ll'} \left(\frac{1}{P_l} + \frac{1}{1 - P_l} \right), \tag{9}$$

where P_l denotes the probability that site l is mutated. Since we are interested in extracting $\vec{\Delta}$, we can simply set to zero the diagonal elements of C^{-1}, which are dominated by conservation effects, to obtain a new matrix

$$\tilde{C}^{-1}_{ll'} \approx (1 - \delta_{ll'}) \kappa \Delta_l \Delta_{l'}. \tag{10}$$

The first eigenvector of $\tilde{C}^{-1}$ (associated with its largest eigenvalue) should accurately report $\vec{\Delta}$ since, except for its zero diagonal, $\tilde{C}^{-1}$ is proportional to the outer product $\vec{\Delta} \otimes \vec{\Delta}$. We call this approach the *Inverse Covariance Off-Diagonal* (ICOD) method. As shown in Fig. 2(d-e), ICOD overcomes the difficulty experienced by PCA for biased selection, while performing equally well as PCA for unbiased selection (Fig. S2, S1 Appendix). Removing the diagonal elements of C^{-1} before diagonalizing is crucial: otherwise, the first eigenvector of C^{-1} is the same as the last eigenvector of C and suffers from the same shortcomings for strong conservation. Here too, eigenvectors associated to both small and large eigenvalues contain information about $\vec{\Delta}$ (Figs. 2(b,d)).

Selection on multiple traits

An important challenge in sector analysis is distinguishing multiple, independently evolving sectors [4, 7, 52]. We can readily generalize our fitness function (Eqs. 3, 6) to allow for selection on multiple additive traits:

$$w(\vec{S}) = -\sum_{i=1}^{N} \frac{\kappa_i}{2} \left(\sum_{l=1}^{L} \Delta_{i,l} S_l - T_i^* \right)^2 , \tag{11}$$

where N is the number of distinct additive traits $T_i(\vec{S}) = \sum_l \Delta_{i,l} S_l$ under selection, $\vec{\Delta}_i$ is the vector of mutational effects on trait T_i, κ_i is the strength of selection on this trait, and T_i^* is the associated selection bias. For example, $\vec{\Delta}_1$ might measure how mutations change a protein's binding affinity, while $\vec{\Delta}_2$ might be related to its thermal stability, etc.

In S1 Appendix Fig. S5, we consider selection on two distinct additive traits, using synthetically-generated random mutational-effect vectors $\vec{\Delta}_1$ and $\vec{\Delta}_2$ (S1 Appendix). Note that these mutational effects are thus unrelated to our toy model of protein elastic deformations: as stated above, our approach holds for any additive trait under selection. ICOD then yields *two* large outlier eigenvalues of the modified inverse covariance matrix $\tilde{C}^{-1}$. The associated eigenvectors accurately recover both $\vec{\Delta}_1$ and $\vec{\Delta}_2$, after a final step of Independent Component Analysis (ICA) [7, 53, 54] that successfully disentangles the contributions coming from the two constraints (see S1 Appendix).

Performance in sector recovery

We further tested the performance of ICOD by systematically varying the selection bias, both for our toy model of PDZ elastic deformations and for more general synthetically-generated random mutational-effect vectors (Fig. 3). ICOD achieves high Recovery of these various mutational-effect vectors for both single and double selection over a broad range of selection biases T^*, albeit performance falls off in the limit of extreme bias.

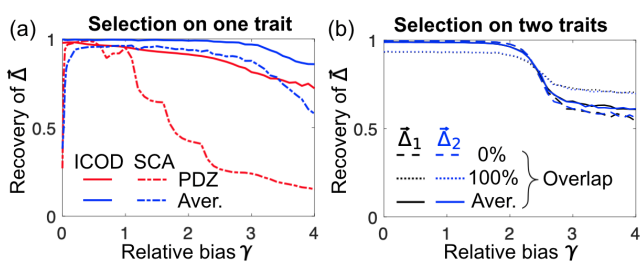

Fig 3. Average recovery of mutational-effect vectors $\vec{\Delta}$ as a function of relative selection bias $\gamma \equiv (T^* - \langle T \rangle)/\sqrt{\langle (T - \langle T \rangle)^2 \rangle}$ on the selected additive trait T. (a) Selection on a single trait. Different $\vec{\Delta}$s are used to generate sequence ensembles: the elastic-network $\vec{\Delta}$ from Fig. 1 (red); synthetic $\vec{\Delta}$s (S1 Appendix) with number of sites of large mutational effect (sector sites) ranging from 1 to 100, for sequences of length $L = 100$ (blue). Recovery is shown for ICOD (solid curves) and for SCA [4, 7] (dashed curves). (b) Selection on two distinct traits. Different pairs of synthetic $\vec{\Delta}$s (S1 Appendix) are used to generate sequence ensembles (with $L = 100$): "0%" indicates two non-overlapping sectors, each with 20 sites; "100%" indicates two fully overlapping sectors, each with 100 sites; "Aver." indicates average Recovery over 100 cases of double selection, where the single-sector size increases from 1 to 100, and the overlap correspondingly increases from 0 to 100. ICA was applied to improve Recovery (S1 Appendix).

How does ICOD compare with other approaches to identifying sectors? We compared the performance of ICOD with Statistical Coupling Analysis (SCA), the original PCA-based method [4, 7]. In SCA, the covariance matrix C is

reweighted by a site-specific conservation factor ϕ_l, the absolute value is taken, $\tilde{C}^{(\mathrm{SCA})}_{ll'} = |\phi_l C_{ll'} \phi_{l'}|$, and sectors are identified from the leading eigenvectors of $\tilde{C}^{(\mathrm{SCA})}$. We therefore tested the ability of the first eigenvector of $\tilde{C}^{(\mathrm{SCA})}$ to recover $\vec{\Delta}$ for a single selection. We found that the square root of the elements of the first SCA eigenvector can provide high Recovery of $\vec{\Delta}$ (Figs. 3, S13, S14) (S1 Appendix). However, the performance of SCA relies on conservation through ϕ_l, and it has been shown that residue conservation actually dominates sector identification by SCA in certain proteins [52]. Consequently, for unbiased selection, SCA breaks down (Fig. 3(a), dashed curves) and cannot identify sector sites (S1 Appendix Fig. S17). ICOD does not suffer from such shortcomings, and performs well over a large range of selection biases. Note that in SCA, only the top eigenvectors of $\tilde{C}^{(\mathrm{SCA})}$ convey information about sectors (Figs. S13, S15).

We also compared ICOD with another PCA-based approach [34], which employs an inference method specific to the generalized Hopfield model, and should thus be well adapted to identifying sectors within our physical model (Eq. 6). Overall, this specialized approach performs similarly to ICOD, being slightly better for very localized sectors, but less robust than ICOD for strong selective biases and small datasets (S1 Appendix). Exactly as for PCA and ICOD, within this method, the top Recovery is obtained for the bottom eigenvector of the (modified) covariance matrix, consistent with $\vec{\Delta}$ in our model being a repulsive pattern [34], but large Recoveries are also obtained for the top eigenvectors (Fig. S18).

Robustness to different forms of selection

To assess the robustness of functional sectors to selections different from the simple Gaussian selection window of Eqs. 2-3, we selected sequences with an additive trait T above a threshold T_t, and varied this threshold. For instance, a fluorescent protein might be selected to be fluorescent enough, which could be modeled by requiring that (a nonlinear transform of) an additive trait be sufficiently large [10]. As shown in Fig. 4, the corresponding sectors are identified by ICOD as well as those resulting from our initial Gaussian selection window. In Fig. 4(d), we show the performance of both ICOD and SCA at recovering sectors arising from selection with a threshold. Consistent with previous results (see Fig. 3), we find that ICOD is more robust than SCA to extreme selections. We also successfully applied ICOD to other forms of selection: Fig. S8 shows the case of a quartic fitness function replacing the initial quadratic one (Eq. 3) in the Boltzmann distribution (Eq. 2) and Fig. S9 shows the case of a rectangular selection window (S1 Appendix). These results demonstrate the robustness of functional sectors, and of ICOD, to different plausible forms of selection.

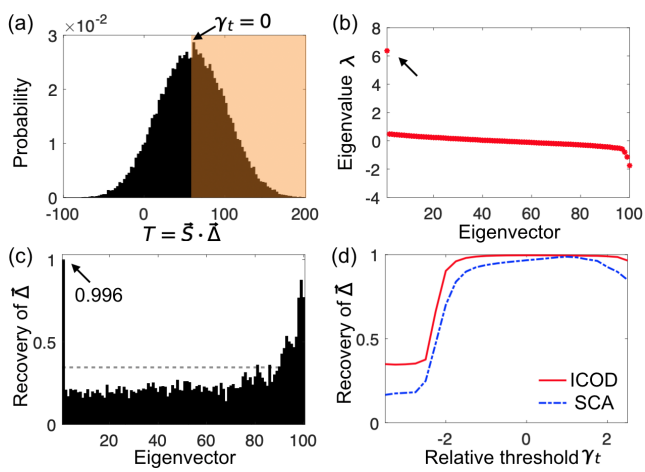

Fig 4. Identification of sectors that result from threshold-based selection. (a) Histogram of the additive trait $T(\vec{S}) = \vec{S} \cdot \vec{\Delta}$ for randomly sampled sequences where 0 and 1 are equally likely at each site. Sequence length is $L = 100$, mutational effects are synthetically generated with 20 sector sites (see S1 Appendix). Sequences are selected if they have a trait value $T(\vec{S}) > T_t$ (orange shaded region). Selection is shown for $T_t = \langle T \rangle$, or equivalently $\gamma_t = 0$, in terms of the relative threshold $\gamma_t \equiv (T_t - \langle T \rangle)/\sqrt{\langle (T - \langle T \rangle)^2 \rangle}$. (b) Eigenvalues of the ICOD-modified inverse covariance matrix $\tilde{C}^{-1}$ (Eq. 10) of the selected sequences for $\gamma_t = 0$. (c) Recovery of $\vec{\Delta}$ for all eigenvectors of $\tilde{C}^{-1}$ for $\gamma_t = 0$. Gray dashed line: random expectation of Recovery. (d) Recovery of $\vec{\Delta}$ for ICOD and for SCA as functions of the relative selection threshold γ_t. The data in (d) is averaged over 100 realizations of $\vec{\Delta}$.

Extension to 21-state sequences and to natural sequences

So far, we have considered binary sequences, with only one type of mutation with respect to the reference state. In the S1 Appendix, we demonstrate that our formalism, including the ICOD method, extends to mutations among q different states. The case $q = 21$, which includes the 20 different amino-acid types plus the alignment gap is the relevant one for real proteins. The single-site mutational effects Δ_l are then replaced by state-specific mutational effects $\Delta_l(\alpha_l)$ with $\alpha_l \in \{1, \ldots, 21\}$ (see Eq. 1). Fig. S10 (S1 Appendix) shows that the generalized version of ICOD performs very well on synthetic data generated for the case $q = 21$. We further demonstrate that sector identification is robust to gauge changes (reference changes) and to the use of pseudocounts (S1 Appendix).

While the main purpose of this article is to propose an operational definition of functional protein sectors and to understand how they can arise, an interesting next question will be to investigate what ICOD can teach us about real data. As a first step in this direction, we applied ICOD to a multiple sequence alignment of PDZ domains. In this analysis, we employed a complete description with $q = 21$, but we compressed the ICOD-modified inverse matrix using the Frobenius norm to focus on overall (and not residue-specific) mutational effects (see S1 Appendix for details). As shown in Figs. 5(a) and (b), both ICOD and SCA identify one strong outlying large eigenvalue, thus confirming that PDZ has only one sector [6]. Recall that due to the inversion step, the largest eigenvalue in ICOD is related to the mode with smallest variance, whose importance was demonstrated above. Furthermore, as seen in Figs. 5(c) and (d), both methods correctly predict the majority of residues found experimentally to have important mutational effects on ligand binding to the PDZ domain shown in Fig. 1(a) [6]. For instance, over the 20 top sites identified by ICOD (resp. SCA), we find that 85% (resp. 75%) of them are also among the 20 experimentally most important sites. Note that for SCA, we recover the result from Ref. [6]. The performance of ICOD is robust to varying the cutoff for removal of sites with a large proportion of gaps (see Fig. S21), but notably less robust than SCA to pseudocount variation (see Fig. S22).

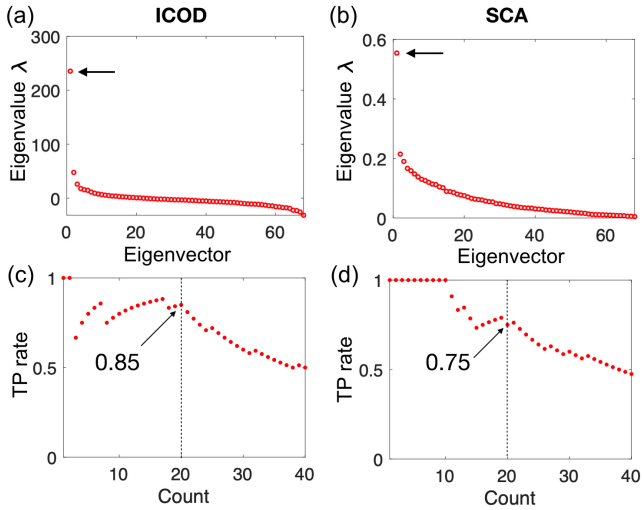

Fig 5. Performance of ICOD and SCA at predicting the 20 sites with largest experimentally-determined mutational effects in a PDZ domain. (a) Eigenvalues of the compressed ICOD-modified inverse covariance matrix $\tilde{C}^{-1}$ (S1 Appendix). (b) Eigenvalues of the SCA matrix. (c) True Positive (TP) rates obtained by taking the first eigenvector $\vec{\nu}^{(1)}$ from the compressed ICOD-modified inverse covariance matrix, generating a ranked list of sites of descending magnitudes of the components $||\nu_l^{(1)}||$ of this eigenvector at each site l (S1 Appendix), and computing the fraction of the top sites in this predicted ordering that are also among the 20 experimentally most important sites [6]. Results are shown versus the number of top predicted sites ("count"). (d) TP rates from SCA, computed as in panel (c). In panels (c) and (d), the TP rate values obtained for the top 20 predicted sites are indicated by arrows. In all panels, a pseudocount ratio $\Lambda = 0.02$ was used, and sites with more than 15% gap state were discarded (see S1 Appendix for details).

Importantly, both ICOD and SCA perform much better than random expectation, which is 29%. Hence, both of these methods can be useful to identify functionally important sites. The slightly greater robustness of SCA to pseudocounts on this particular dataset (see Fig. S22) might come from the fact that many of the experimentally-identified functionally important sites in the PDZ domain are strongly conserved [52], which makes the conservation reweighting step in SCA advantageous. Since residue conservation alone is able to predict most of the experimentally important PDZ sites [52], we also compared conservation to SCA and ICOD: ranking sites by conservation (employing

the conservation score of Ref. [7], see S1 Appendix) indeed identifies 70% of the top 20 experimentally-determined sites with important mutational effects. Interestingly, ICOD scores are slightly more strongly correlated with conservation than SCA scores are correlated with conservation (see Fig. S23), despite the fact that conservation is explicitly used in SCA and not in ICOD.

Overall, this preliminary application to real data highlights the ability of ICOD to identify functionally related amino acids in a principled way that only relies on covariance. We emphasize that the main goal of this paper is to provide insight into the possible physical origins of sectors, and into the statistical signatures of these physical sectors in sequence data. A more extensive application of ICOD and related methods to real sequence data will be the subject of future work.

Discussion

We have demonstrated how sectors of collectively correlated amino acids can arise from evolutionary constraints on functional properties of proteins. Our model is very general, as it only relies on the functional property under any of various forms of selection being described by an underlying additive trait, which has proven to be valid in many relevant situations [8–10, 13].

We showed that the primary signature of functional selection acting on sequences lies in the small-eigenvalue modes of the covariance matrix. In contrast, sectors are usually identified from the large-eigenvalue modes of the SCA matrix [4, 7]. This is not in contradiction with our results because, as we showed, signatures of our functional sectors are often also found in large-eigenvalue modes of the covariance matrix. Besides, the construction of the SCA matrix from the covariance matrix involves reweighting by conservation and taking an absolute value or a norm [4, 7], which can substantially modify its eigenvectors, eigenvalues, and their order. Conservation is certainly important in real proteins, especially in the presence of phylogeny; indeed, the SCA matrix, which includes both conservation and covariance, was recently found to capture well experimentally-measured epistasis with respect to the free energy of PDZ ligand binding [55]. However, the fundamental link we propose between functional sectors and small-eigenvalue modes of the covariance matrix is important, since large-eigenvalue modes of the covariance matrix also contain confounding information about subfamily-specific residues [56] and phylogeny [57], and consistently, some sectors identified by SCA have been found to reflect evolutionary history rather than function [4]. Interestingly, the small-eigenvalue modes are also the ones that contain most information about structural contacts in real proteins [35]. Hence, our results help explain previously observed correlations between sectors and contacts, e.g. the fact that contacts are overrepresented within a sector but not across sectors [58].

We introduced a principled method to detect functional sectors from sequence data, based on the primary signature of these sectors in the small-eigenvalue modes of the covariance matrix. We further demonstrated the robustness of our approach to the existence of multiple traits simultaneously under selection, to various forms of selection, and to data-specific questions such as reference choices and pseudocounts.

Importantly, our modeling approach allowed us to focus on functional selection alone, in the absence of historical contingency and of specific structural constraints, thus yielding insights complementary to purely data-driven methods. The collective modes investigated here are just one source of residue-residue correlations. Next, it will be interesting to study the intriguing interplay between functional sectors, phylogeny, and contacts, and to apply our methods to multiple protein families. Our results shed light on an aspect of the protein sequence-function relationship and open new directions in protein sequence analysis, with implications in synthetic biology, building toward function-driven protein design.

Supporting information

S1 Appendix. Methodological details and further results. In S1 Appendix, we present additional details about our model and methods, as well as additional results.

Acknowledgments

S.-W. W. and N. S. W. acknowledge the Center for the Physics of Biological Function under NSF Grant PHY-1734030. A.-F. B. and N. S. W. acknowledge the Aspen Center for Physics, which is supported by NSF Grant PHY-1607611. S.-W. W. was supported by the NSFC under Grants No. U1430237 and 11635002. S.-W. W. also acknowledges Tsinghua University for supporting a half-year visit in Princeton University. N. S. W. was supported by NSF Grant MCB-1344191 and by NIH Grant R01 GM082938.

S1 Appendix

Contents

1 Supplemental results for elastic network model of PDZ domain **12**

2 Recovery by a random vector **12**

3 Inverse covariance matrix of our sequence ensembles **13**
3.1 Binary sequences . . . 14
3.1.1 From a sector model for binary sequences to an Ising model . . . 14
3.1.2 First-order small-coupling expansion . . . 14
3.1.3 Application to our sector model . . . 15
3.2 Sequences with q possible states at each site . . . 16
3.2.1 From a sector model to a Potts model for sequences . . . 16
3.2.2 First-order small-coupling expansion . . . 16
3.2.3 Application to our sector model . . . 17
3.2.4 Selection on multiple traits . . . 17

4 Robustness of functional sectors and of ICOD **18**
4.1 Robustness of ICOD to selection bias, selection strength, and multiple selections . . . 18
4.2 Robustness of functional sectors to different forms of selection . . . 20
4.3 Multiple states per site and alternative gauge choice . . . 22
4.4 Pseudocounts . . . 24

5 Performance of SCA **25**
5.1 Analytical estimates for $\langle S_l \rangle_*$ and $C_{ll'}$ for a single selection with binary sequences . . . 25
5.2 Analysis of the SCA method . . . 26
5.3 Comparison between ICOD and SCA . . . 28

6 Performance of a method based on the generalized Hopfield model **29**

7 Application of ICOD to a multiple sequence alignment of PDZ domains **32**

1 Supplemental results for elastic network model of PDZ domain

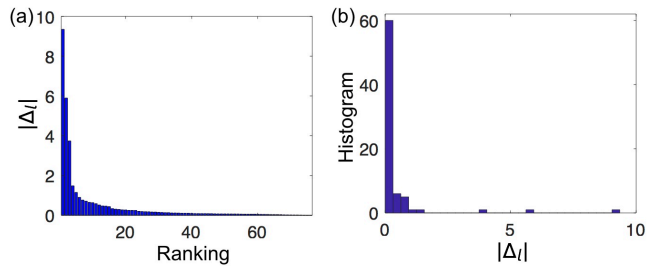

Fig S1. Magnitude of single-site mutational effects Δ_l for the PDZ domain conformational change from Fig. 1. (a) Magnitudes by rank. (b) Histogram of magnitudes. According to our definition, the sites of large magnitude constitute "sector" sites with respect to selection on the energy cost of this conformational change, while all others are "non-sector" sites.

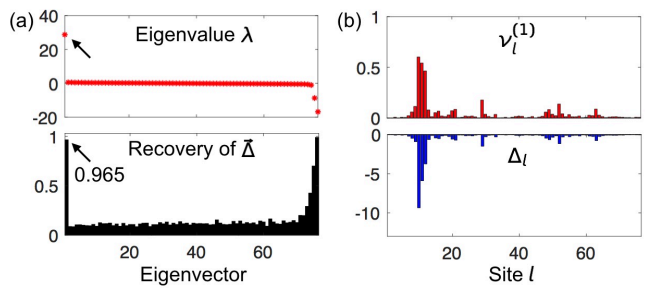

Fig S2. Performance of ICOD for the selected sequence ensemble from Fig. 1. (a) Eigenvalues for ICOD method (upper) and Recovery of $\vec{\Delta}$ for all eigenvectors (lower). (b) Leading eigenvector $\nu_l^{(1)}$ (upper) and mutational effect Δ_l at site l (lower, same as in Fig. 1(f)). The excellent performance of ICOD on this unbiased ensemble of sequences supports the general applicability of the ICOD method to both biased and unbiased sequence ensembles.

2 Recovery by a random vector

Here, we calculate the random expectation of the Recovery of the mutational-effect vector $\vec{\Delta}$ by a generic other vector $\vec{\nu}$, in order to establish a null model to which to compare. For a binary sequence, Recovery, as defined in Eq. 8, can be expressed as

$$\text{Recovery} = \vec{\Delta}' \cdot \vec{\nu'} = \sum_{l=1}^{L} \Delta'_l \, \nu'_l, \tag{S1}$$

with $\Delta'_l = |\Delta_l|/\sqrt{\sum_l \Delta_l^2}$ and $\nu'_l = |\nu_l|/\sqrt{\sum_l \nu_l^2}$. As before, L denotes the length of the sequence and hence the number of components of $\vec{\Delta}$ and $\vec{\nu}$. As $\vec{\nu'}$ is a normalized L-dimensional vector, its components can be expressed in L-dimensional spherical coordinates using $L-1$ angles θ_i:

$$\nu'_l = \left(\prod_{i=1}^{l-1} \sin\theta_i\right) \cos\theta_l \qquad \forall l \in \{1, \ldots, L-1\}\,, \tag{S2}$$

$$\nu'_L = \prod_{i=1}^{L-1} \sin\theta_i\,, \tag{S3}$$

where $\theta_i \in [0, \pi/2]$ for all $i \in \{1, \cdots, L\}$, because all components of $\vec{\nu'}$ are nonnegative. Note that we employ the usual convention that empty products are equal to one: Eq. S2 yields $\nu'_1 = \cos\theta_1$.

The average Recovery for a random vector $\vec{\nu'}$ with an orientation uniformly distributed in the L-dimensional sphere reads:

$$\langle \text{Recovery} \rangle = \frac{\int_\Omega d\Omega \, \sum_l \Delta'_l \nu'_l}{\int_\Omega d\Omega} = \frac{\sum_l \Delta'_l \, I_l}{\int_\Omega d\Omega} \,, \tag{S4}$$

where the angular element is $d\Omega = \prod_{i=1}^{L-1} d\theta_i \, \sin^{L-i-1}(\theta_i)$, the integration domain is $\Omega = [0, \pi/2]^{L-1}$, and we have introduced $I_l = \int_\Omega d\Omega \, \nu'_l$. Using Eq. S2, we obtain for $1 \leq l \leq L-1$

$$I_l = \int_\Omega d\Omega \, \nu'_l = \left(\prod_{i=1}^{l-1} \int_0^{\pi/2} d\theta_i \, \sin^{L-i}(\theta_i) \right) \left(\int_0^{\pi/2} d\theta_l \, \sin^{L-l-1}(\theta_l) \cos(\theta_l) \right) \left(\prod_{i=l+1}^{L-1} \int_0^{\pi/2} d\theta_i \, \sin^{L-i-1}(\theta_i) \right) \,, \tag{S5}$$

and similarly, Eq. S3 yields

$$I_L = \int_\Omega d\Omega \, \nu'_L = \prod_{i=1}^{L-1} \int_0^{\pi/2} d\theta_i \, \sin^{L-i}(\theta_i) \,. \tag{S6}$$

Using the following results valid for $n > -1$:

$$\int_0^{\pi/2} d\theta \, \sin^n(\theta) = \frac{\sqrt{\pi}}{2} \frac{\Gamma\left(\frac{1+n}{2}\right)}{\Gamma\left(\frac{n+2}{2}\right)} \,; \quad \int_0^{\pi/2} d\theta \, \sin^n(\theta) \cos(\theta) = \frac{1}{n+1} \,, \tag{S7}$$

where Γ denotes the Euler Gamma function, which satisfies $\Gamma(x+1) = x\,\Gamma(x)$ for all x, we obtain for $1 \leq l \leq L$:

$$I_l = \frac{\pi^{(L-1)/2}}{2^{L-1}\,\Gamma\left(\frac{L+1}{2}\right)}, \tag{S8}$$

which is independent of l. Besides,

$$\int_\Omega d\Omega = \frac{\pi^{L/2}}{2^{L-1}\,\Gamma(L/2)}. \tag{S9}$$

Combining Eq. S4 with Eqs. S8 and S9 finally yields

$$\langle \text{Recovery} \rangle = \frac{\sum_l \Delta'_l \, I_l}{\int_\Omega d\Omega} = \frac{\Gamma\left(L/2\right)}{\sqrt{\pi}\,\Gamma\left(\frac{L+1}{2}\right)} \sum_l \Delta'_l = \frac{\Gamma\left(L/2\right)}{\sqrt{\pi}\,\Gamma\left(\frac{L+1}{2}\right)} \frac{\sum_l |\Delta_l|}{\sqrt{\sum_l \Delta_l^2}}. \tag{S10}$$

In particular, in the relevant regime $L \gg 1$, an asymptotic expansion of Γ yields:

$$\langle \text{Recovery} \rangle \approx \sqrt{\frac{2}{\pi L}} \frac{\sum_l |\Delta_l|}{\sqrt{\sum_l \Delta_l^2}}. \tag{S11}$$

The maximum expectation of Recovery is obtained when all components of $\vec{\Delta}$, i.e. all mutational effects, are identical:

$$\langle \text{Recovery} \rangle_{\max} = \sqrt{\frac{2}{\pi}} \approx 0.798. \tag{S12}$$

Conversely, the average Recovery becomes minimal when only one component of $\vec{\Delta}$ is nonzero, which constitutes the limit of the case where the mutational effect at one site is dominant:

$$\langle \text{Recovery} \rangle_{\min} = \sqrt{\frac{2}{\pi L}}, \tag{S13}$$

which approaches zero in the limit $L \to \infty$.

3 Inverse covariance matrix of our sequence ensembles

Here, we present a derivation of the small-coupling approximation of the inverse covariance matrix for our artificially-generated sequence ensembles. In this small-coupling limit, the inverse covariance matrix provides an estimate of the energetic couplings used to generate the data. More generally, deducing energetic parameters from observed statistics is a well-known inference problem, also known as an inverse problem. Two-body energetic couplings can be

inferred from the one and two-body frequencies observed in the data, using a standard maximum entropy approach. However, the exact calculation of the energetic terms is difficult, and various approximations have been developed. Following Refs [37, 38], we use the mean-field or small-coupling approximation, which was introduced in Ref. [50] for the Ising spin-glass model. For the sake of completeness, we now review the main steps of the calculation, which follow Ref. [37]. Note that we do not use inference methods specific to low-rank coupling matrices [34, 35] because we wish to retain generality, with the application to real sequence data in mind.

We begin with the case of binary sequences, which is discussed in the main text. Following that, we generalize to cases where more than two states are allowed at each site, such as the 21 possible states for real protein sequence (20 amino acids plus gap).

3.1 Binary sequences

We begin by deriving Eq. 9 from the main text, which provides an approximation for the inverse covariance matrix of the ensembles of our binary artificial sequences. Each sequence $\vec{S}$ is such that $S_l \in \{0, 1\}$ for each site l with $1 \leq l \leq L$, where L is the length of the sequence.

3.1.1 From a sector model for binary sequences to an Ising model

Recall the fitness w of a binary sequence $\vec{S}$ under selection for trait T to be close to T^* (Eq. 6):

$$w(\vec{S}) = -\frac{\kappa}{2}\left(T(\vec{S}) - T^*\right)^2 = -\frac{\kappa}{2}\left(\sum_l \Delta_l S_l - T^*\right)^2 . \tag{S14}$$

We introduce $s_l = 2S_l - 1$: it is an "Ising spin" variable ($S_l = 0 \Leftrightarrow s_l = -1$ and $S_l = 1 \Leftrightarrow s_l = 1$). The fitness in Eq. S14 can be rewritten as

$$w(\vec{s}) = -\frac{\kappa}{2}\left(\sum_l D_l s_l - \alpha\right)^2 , \tag{S15}$$

with $D_l = \Delta_l/2$ and $\alpha = T^* - \sum_l D_l$. Expanding yields

$$w(\vec{s}) = -\frac{\kappa}{2}\left(\sum_{l\neq p} D_l D_p s_l s_p + \sum_l D_l^2 - 2\alpha \sum_l D_l s_l + \alpha^2\right) , \tag{S16}$$

where we have used the fact that $s_l^2 = 1$. The second term and the last term in Eq. S16 are both constants, and therefore our fitness is equivalent to

$$w(\vec{s}) = -\frac{\kappa}{2}\left(\sum_{l\neq p} D_l D_p s_l s_p - 2\alpha \sum_l D_l s_l\right) . \tag{S17}$$

This fitness has the form of a standard Ising Hamiltonian with inter-spin couplings and local fields, albeit with the convention difference in overall sign between fitness and energy.

3.1.2 First-order small-coupling expansion

We next consider the general Ising Hamiltonian with inter-spin couplings and local fields

$$H(\vec{s}) = -\frac{1}{2}\epsilon \sum_{i\neq j} J_{ij} s_i s_j - \sum_i h_i s_i , \tag{S18}$$

where ϵ is a constant to be employed in a small-coupling expansion. With this Hamiltonian, taking thermal energy $k_{\mathrm{B}}T = 1$, the equilibrium probability of finding a particular sequence $\vec{s}$ is

$$P(\vec{s}) = \frac{1}{Z} e^{-H(\vec{s})}, \tag{S19}$$

where $Z = \sum_{\vec{s}} e^{-H(\vec{s})}$.

Introducing $F = -\log Z$, we have

$$\frac{\partial F}{\partial h_i} = -\langle s_i \rangle = -m_i \,,$$
$$\frac{\partial^2 F}{\partial h_i \partial h_j} = -\frac{\partial m_i}{\partial h_j} = \langle s_i \rangle \langle s_j \rangle - \langle s_i s_j \rangle = -C'_{ij} \,, \tag{S20}$$

where, following the Ising terminology, m_i denotes the average magnetization at site i, while C' denotes the covariance matrix in the Ising convention. Note that, using the identity $m_i = 2P_i - 1$, where P_i denotes the probability that $s_i = 1$, we obtain

$$C'_{ij} = \langle s_i s_j \rangle - \langle s_i \rangle \langle s_j \rangle = 4\,(P_{ij} - P_i P_j) = 4\,C_{ij} \,, \tag{S21}$$

where P_{ij} is the probability that $s_i = s_j = 1$, and C denotes the covariance matrix in the Potts convention, which is used in the main text because it allows straightforward generalization to the case where more than two states are possible at each site.

Performing a Legendre transform, we introduce $G = F + \sum_i m_i h_i$, yielding

$$\frac{\partial G}{\partial m_i} = h_i \,, \tag{S22}$$
$$\frac{\partial^2 G}{\partial m_i \partial m_j} = \frac{\partial h_i}{\partial m_j} = {C'_{ij}}^{-1} \,. \tag{S23}$$

We now perform a small-coupling expansion and express G to first order in ϵ (see Eq. S18): $G(\epsilon) \approx G(0) + \epsilon G'(0)$. Since sites are independent for $\epsilon = 0$, it is straightforward to express $G(0)$ and $G'(0)$ as a function of the one-body expectations, represented by m_i, and of the couplings. We obtain

$$G(0) = \sum_i \frac{m_i + 1}{2} \log\left(\frac{m_i + 1}{2}\right) + \frac{1 - m_i}{2} \log\left(\frac{1 - m_i}{2}\right) \,, \tag{S24}$$

and

$$G'(0) = \frac{\partial G}{\partial \epsilon}(0) = -\frac{1}{2} \sum_{i \neq j} J_{ij} m_i m_j \,. \tag{S25}$$

Using these expressions, and taking $\epsilon = 1$ in the expansion, we obtain the following approximation for G:

$$G \approx \sum_i \frac{m_i + 1}{2} \log\left(\frac{m_i + 1}{2}\right) + \frac{1 - m_i}{2} \log\left(\frac{1 - m_i}{2}\right) - \frac{1}{2} \sum_{i \neq j} J_{ij} m_i m_j \,. \tag{S26}$$

Using Eqs. S22 and S23, we obtain the elements of the inverse covariance matrix from Eq. S26:

$${C'_{kl}}^{-1} = -J_{kl} \,,\ \forall l \neq k \,,$$
$${C'_{ll}}^{-1} = \frac{1}{2}\left(\frac{1}{1 + m_l} + \frac{1}{1 - m_l}\right) = \frac{1}{4}\left(\frac{1}{P_l} + \frac{1}{1 - P_l}\right) \,, \tag{S27}$$

where P_l denotes the probability that $s_l = 1$.

Note that Eq. S26 is a first-order small-coupling (or mean-field) approximation. The expansion can be extended to higher order, and the second-order expansion is known as the Thouless, Anderson, and Palmer (TAP) free energy [50, 59].

3.1.3 Application to our sector model

Comparing Eqs. S17 and S18 (with $\epsilon = 1$) allows us to identify the couplings in our sector model as

$$J_{kl} = -\kappa\, D_k D_l = -\kappa\, \Delta_k \Delta_l / 4 \,,\ \forall k \neq l \,. \tag{S28}$$

Note that this expression is in the Ising gauge (also known as the zero-sum gauge). Recall also that the link to the Potts convention is made through $C' = 4\,C$ (Eq. S21), which implies ${C'}^{-1} = C^{-1}/4$. Finally, recall that fitness and energy have opposite signs.

Hence, in the Potts convention, Eq. S27 yields for our sector model:

$$C^{-1}_{kl} = \kappa \Delta_k \Delta_l \,,\ \forall l \neq k \,,$$
$$C^{-1}_{ll} = \frac{1}{P_l} + \frac{1}{1 - P_l} \,. \tag{S29}$$

This corresponds to Eq. 9 in the main text.

3.2 Sequences with q possible states at each site

3.2.1 From a sector model to a Potts model for sequences

Motivated by the fact that a real protein sequence has 21 possible states at each site (20 different amino acids plus gap), we now generalize the above result to the case where q states are possible at each site. We denote these states by α with $\alpha \in \{1, .., q\}$. Our sector model can then be mapped to a q-state Potts model. The length-L vector $\vec{\Delta}$ of single-site mutational effects introduced in the two-state case in the main text is replaced by a $(q-1) \times L$ matrix of mutational effects, each being denoted by $\Delta_l(\alpha_l)$. These mutational effects can be measured with respect to a reference sequence $\vec{\alpha}^0$ satisfying $\Delta_l(\alpha_l^0) = 0$, $\forall l \in \{1, \ldots, L\}$: at each site l, the state present in the reference sequence $\vec{\alpha}^0$ serves as the reference with respect to which the mutational effects at that site are measured. For the sake of simplicity, we will take state q as reference state at all sites. This does not lead to any loss of generality, since it is possible to reorder the states for each l.

The generalization of the fitness function Eq. 6 (Eq. S14) to our q-state model can be written as

$$w(\vec{\alpha}) = -\frac{\kappa}{2}\left(T(\vec{\alpha}) - T^*\right)^2 = -\frac{\kappa}{2}\left(\sum_{l=1}^{L} \Delta_l(\alpha_l) - T^*\right)^2 , \tag{S30}$$

(see Eq. 3 in the main text). Expanding this expression, discarding a constant term, and using the fact that there can only be one state at each site, we find that the fitness of sequences can be expressed as

$$w(\vec{\alpha}) = -\frac{\kappa}{2}\sum_{l \neq k} \Delta_l(\alpha_l)\Delta_k(\alpha_k) - \frac{\kappa}{2}\sum_{l=1}^{L} \Delta_l(\alpha_l)\left(\Delta_l(\alpha_l) - 2\,T^*\right) . \tag{S31}$$

This is a particular case of the more general Potts Hamiltonian

$$H(\vec{\alpha}) = -\frac{1}{2}\sum_{l \neq k} e_{lk}(\alpha_l, \alpha_k) - \sum_{l=1}^{L} h_l(\alpha_l) , \tag{S32}$$

which is the one usually considered in Direct Coupling Analysis (DCA) [37, 38].

In order to identify Eq. S31 and Eq. S32, one must deal with the degeneracies present in Eq. S32, where the number of independent parameters is $L(q-1) + L(L-1)(q-1)^2/2$ [60]. To lift this degeneracy, we choose the gauge usually taken in mean-field DCA [37]: $e_{lk}(\alpha_l, q) = e_{lk}(q, \alpha_k) = h_l(q) = 0$ for all l, k, α_l, α_k. This choice is consistent with taking state q as the reference state for mutational effects (see above), and we will refer to it as the reference-sequence gauge. This gauge choice enables us to identify the couplings between Eq. S31 and Eq. S32:

$$e_{lk}(\alpha_l, \alpha_k) = -\kappa\Delta_l(\alpha_l)\Delta_k(\alpha_k) , \tag{S33}$$

for all $l \neq k$, and all α_l, α_k, with $\Delta_l(q) = 0$ for all l (recalling that fitness and energy have opposite signs).

3.2.2 First-order small-coupling expansion

The derivation of the first-order mean-field or small-coupling approximation for q-state models is very similar to the Ising case presented above. Hence, we will simply review the main results (see Ref. [37]).

We start with the Hamiltonian

$$H(\vec{\alpha}) = -\frac{\epsilon}{2}\sum_{l \neq k} e_{lk}(\alpha_l, \alpha_k) - \sum_{l=1}^{L} h_l(\alpha_l) , \tag{S34}$$

where ϵ has been introduced to perform the small-coupling expansion. Eq. S34 coincides with Eq. S32 for $\epsilon = 1$. Considering $F = -\log(Z)$ with $Z = \sum_{\vec{\alpha}} e^{-H(\vec{\alpha})}$, where $H(\vec{\alpha})$ is the Potts Hamiltonian in Eq. S34, we have for all k and all $\alpha_k < q$:

$$\frac{\partial F}{\partial h_k(\alpha_k)} = -P_k(\alpha_k) , \tag{S35}$$

where $P_k(\alpha_k)$ is the one-body probability. Similarly, we have for all k, l and all $\alpha_k < q$ and $\alpha_l < q$:

$$\frac{\partial^2 F}{\partial h_l(\alpha_l)\partial h_k(\alpha_k)} = -\frac{\partial P_k(\alpha_k)}{\partial h_l(\alpha_l)} = -C_{kl}(\alpha_k, \alpha_l) , \tag{S36}$$

where we have introduced the covariance $C_{kl}(\alpha_k, \alpha_l) = P_{kl}(\alpha_k, \alpha_l) - P_k(\alpha_k)P_l(\alpha_l)$.

We perform a Legendre transform and introduce $G = F - \sum_i \sum_{\alpha_i} h_i(\alpha_i)P_i(\alpha_i)$, yielding

$$\frac{\partial G}{\partial P_k(\alpha_k)} = h_k(\alpha_k)\,, \tag{S37}$$

$$\frac{\partial^2 G}{\partial P_l(\alpha_l)\partial P_k(\alpha_k)} = \frac{\partial h_l(\alpha_l)}{\partial P_k(\alpha_k)} = C^{-1}_{kl}(\alpha_k, \alpha_l)\,, \tag{S38}$$

for all k, l and all $\alpha_k < q$ and $\alpha_l < q$. Note that, in the latter equation, $C^{-1}_{kl}(\alpha, \beta)$ is shorthand for A^{-1}_{ij}, where A is the $(q-1)L \times (q-1)L$ covariance matrix where terms involving the reference state q have been excluded: $A_{ij} = C_{kl}(\alpha, \beta)$, where $i = (q-1)(k-1) + \alpha$ and $j = (q-1)(l-1) + \beta$, with $\alpha \in \{1, \ldots, q-1\}$ and $\beta \in \{1, \ldots, q-1\}$ [61].

We next perform a first-order expansion of G in ϵ, and take $\epsilon = 1$, yielding:

$$G \approx \sum_l \sum_{\alpha_l} P_l(\alpha_l) \log\left(P_l(\alpha_l)\right) - \frac{1}{2} \sum_{l \neq k} \sum_{\alpha_l, \alpha_k} e_{lk}(\alpha_l, \alpha_k) P_l(\alpha_l) P_k(\alpha_k)\,. \tag{S39}$$

Applying Eqs. S37, S38 to Eq. S39, and using $P_l(q) = 1 - \sum_{\alpha_l < q} P_l(\alpha_l)$ gives

$$\begin{aligned} C^{-1}_{kl}(\alpha_k, \alpha_l) &= -e_{kl}(\alpha_k, \alpha_l), \;\; \forall\, l \neq k\,, \\ C^{-1}_{ll}(\alpha_k, \alpha_l) &= \frac{1}{P_k(q)} + \frac{\delta_{\alpha_k \alpha_l}}{P_k(\alpha_k)}\,. \end{aligned} \tag{S40}$$

This result is the standard one found in DCA [37].

3.2.3 Application to our sector model

Combining Eqs. S33 and S40, we obtain for our sector model:

$$\begin{aligned} C^{-1}_{kl}(\alpha_k, \alpha_l) &= \kappa \Delta_k(\alpha_k)\Delta_l(\alpha_l), \;\; \forall\, l \neq k\,, \\ C^{-1}_{ll}(\alpha_k, \alpha_l) &= \frac{1}{P_k(q)} + \frac{\delta_{\alpha_k \alpha_l}}{P_k(\alpha_k)}\,. \end{aligned} \tag{S41}$$

For $q = 2$, Eq. S41 reduces to Eq. S27 (Eq. 9 in the main text), using $1 - P_l = P_l(q)$.

3.2.4 Selection on multiple traits

So far, we have mainly discussed the case where there selection on only one trait (yielding one sector). However, real proteins face various selection pressures. The generalization of the fitness in Eq. S30 to N simultaneous selection on different traits reads

$$w(\vec{S}) = -\sum_{i=1}^{N} \frac{\kappa_i}{2} \left(T_i - T_i^*\right)^2 = -\sum_{i=1}^{N} \frac{\kappa_i}{2} \left(\sum_{l=1}^{L} \Delta_{i,l}(\alpha_l) - T_i^* \right)^2\,, \tag{S42}$$

which corresponds to Eq. 11 in the main text. We choose the reference-state gauge, assuming again for simplicity that the reference state is q at each site. The identification to the general Potts Hamiltonian Eq. S32 (recalling that fitnesses and energies have opposite signs) then yields

$$e_{lk}(\alpha_l, \alpha_k) = -\sum_{i=1}^{N} \kappa_i \Delta_{i,l}(\alpha_l)\Delta_{i,k}(\alpha_k)\,, \tag{S43}$$

which generalizes Eq. S33 to the multiple selection case. Using the small-coupling expansion result in Eq. S40, we obtain the following approximation for the inverse covariance matrix:

$$\begin{aligned} C^{-1}_{kl}(\alpha_k, \alpha_l) &= \sum_{i=1}^{N} \kappa_i \Delta_{i,k}(\alpha_k)\Delta_{i,l}(\alpha_l), \;\; \forall\, l \neq k\,, \\ C^{-1}_{ll}(\alpha_k, \alpha_l) &= \frac{1}{P_k(q)} + \frac{\delta_{\alpha_k \alpha_l}}{P_k(\alpha_k)}\,. \end{aligned} \tag{S44}$$

This generalizes Eq. S41 to the case of simultaneous selection on multiple traits.

4 Robustness of functional sectors and of ICOD

In the main text, we introduced the Inverse Covariance Off-Diagonal (ICOD) method to identify protein sectors from sequence data. The ICOD method exploits the approximate expression derived above for the inverse covariance matrix (Eq. S41); in particular, ICOD makes use of the fact that the off-diagonal elements of C^{-1} are simply related to the elements of the mutational effect vector $\vec{\Delta}$. In this section, we first describe our comparison of ICOD to SCA for single selection, and detail our test of ICOD for double selection, using synthetic binary sequences. Next, we confirm the robustness of the ICOD method to different forms of selection and then show how ICOD can be extended to sequences with more than two states per site, and finally demonstrate its robustness to gauge choice and pseudocounts.

4.1 Robustness of ICOD to selection bias, selection strength, and multiple selections

To quantify the performance of ICOD and to compare to SCA over a range of selection biases we focused on binary sequences. To obtain the average curve for single selections in Fig. 3(a), we first generated 100 distinct synthetic $\vec{\Delta}$s, one for each sector size from $n = 1$ to 100, where sector sites are defined as those with large mutational effects. To this end, the mutational effects of the sector sites and the non-sector sites were sampled, respectively, from zero-mean Gaussian distributions with standard deviations 20 and 1. For each sector size and each selection bias we generated a sequence ensemble of 50,000 random sequences and weighted each sequence according to the distribution

$$P(\vec{S}) = \frac{\exp(w(\vec{S}))}{\sum_{\vec{S}} \exp(w(\vec{S}))}\,, \tag{S45}$$

where $w(\vec{S})$ is the fitness of sequence $\vec{S}$, given by the single selection formula Eq. 6. In general, we wish to employ a selection window whose width in energy (or any other selected variable) scales with the overall width of the unselected distribution. Hence, as mentioned in the main text, we perform all selections with a strength

$$\kappa = \frac{10}{\sum_l \Delta_l^2}\,. \tag{S46}$$

Then, for each method (ICOD or SCA), performance as measured by Recovery of $\vec{\Delta}$ by the first eigenvector was averaged over the 100 different sector sizes.

As an aside, Fig. S3 demonstrates that the performance of ICOD and SCA is robust to varying selection strength κ, as long as $\kappa \sum_l \Delta_l^2 \gg 1$. (A small value of $\kappa \sum_l \Delta_l^2$ implies weak selection, where most random sequences pass selection and the resulting ensemble does not significantly reflect the constraint.)

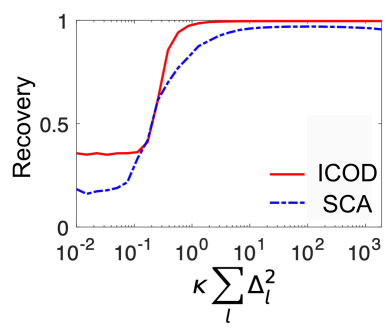

Fig S3. Impact of selection strength κ on the performance of ICOD and SCA on synthetic data. Results obtained on binary synthetic sequences with $L = 100$, selected using a synthetic $\vec{\Delta}$ where the first 20 and the other 80 mutational effects are, respectively, sampled from Gaussian distributions with variances of 20 and 1. Selection is performed on ensembles of 50,000 random sequences, and each data point is obtained by averaging over 100 realizations. The relative bias is $\gamma = 0.5$.

Similarly, to obtain the average curve for double selection in Fig. 3(b), we generated 100 distinct pairs of $\vec{\Delta}_1$s and $\vec{\Delta}_2$s, one pair for each sector size from $n = 1$ to 100. Specifically, the sector for $\vec{\Delta}_1$ consisted of the first n sites, while the sector for $\vec{\Delta}_2$ corresponded to the last n sites, so that the two sectors overlap for $n > 50$. As for the single selections, the mutational effects of the sector sites and the non-sector sites were sampled, respectively, from

Gaussian distributions with standard deviations 20 and 1. As an example, two synthetic $\vec{\Delta}$s for $n = 20$ are shown in Fig. S4. Again, for each sector size and each selection bias, we generated an ensemble of 50,000 random sequences and weighted them according to Eq. S45 along with the double selection formula Eq. S42 (i.e. Eq. 11 in the main text). The performance of ICOD as measured by Recovery of $\vec{\Delta}_1$ and $\vec{\Delta}_2$ by the first two eigenvectors was averaged over the 100 different sector sizes. In Fig. 3(b) we also reported the performance of ICOD for two non-overlapping sectors, each with 20 sites, and for two fully overlapping sectors, each with 100 sites. We followed a protocol similar to that described above, but in each of these cases, we averaged Recovery over 100 realizations using distinct pairs of $\vec{\Delta}_1$ and $\vec{\Delta}_2$.

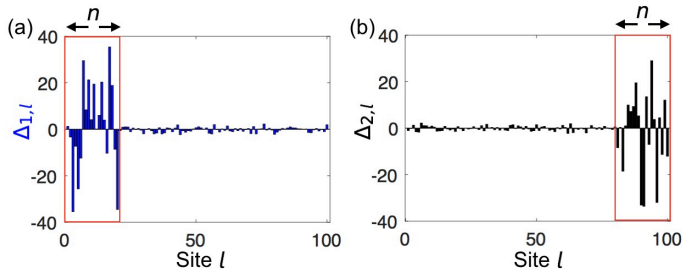

Fig S4. Example of two synthetic $\vec{\Delta}$s generated for the double selection in Fig. 3(b). (a) Generation of $\vec{\Delta}_1$, where the mutational effects for the first 20 sites and for the last 80 sites are sampled, respectively, from zero-mean Gaussian distributions with a standard deviation of 20 and 1. (b) Generation of $\vec{\Delta}_2$, where the mutational effects for the last 20 sites and for the first 80 sites are sampled, respectively, from zero-mean Gaussian distributions with a standard deviation of 20 and 1.

Unless otherwise stated, data for other plots were generated in the same way, i.e. using 50,000 random sequences, sequence length $L = 100$, selection strength κ in Eq. S46, and standard deviation 20/1 of Δ_l in the sector/non-sector sites.

Note that to improve Recovery in the case of double selection, we applied Independent Component Analysis (ICA) [7,53,54] to the first two eigenvectors in order to disentangle the contributions coming from the two constraints. In general, we expect that the first N eigenvectors of the ICOD matrix $\tilde{C}^{-1}$ will report N constraints. However, each of these N eigenvectors is likely to include a mixture of contributions from different constraints. Applying ICA to the first N eigenvectors to recover the individual constraints amounts to assuming that all the constraints are statistically independent. As an example, in Fig. S5, we consider the case of two selections targeting a different set of sites and with different selection windows (one biased, one non-biased). In this case, ICOD plus ICA yields excellent Recovery (Fig. S5). Without ICA, the results are noticeably worse (Fig. S6). Moreover, Fig. 3(b) demonstrates that ICOD plus ICA can achieve a high Recovery for a broad range of overlaps between two sectors.

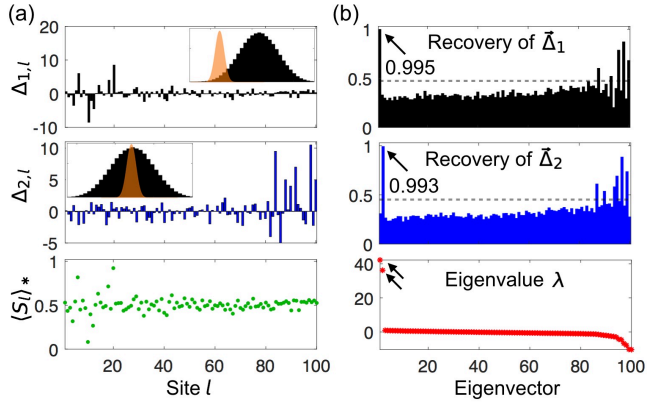

Fig S5. ICOD method for simultaneous selection on two traits. (a) Upper panels: Components at each site l of two synthetically generated mutational-effect vectors, with insets showing biased selection around T_1^* for $\vec{\Delta}_1$ and neutral selection around T_2^* for $\vec{\Delta}_2$. Lower panel: average mutant fraction $\langle S_l \rangle_*$ at site l after selection on both traits. (b) Performance of ICOD method. Recovery of $\vec{\Delta}_1$ and $\vec{\Delta}_2$ for all eigenvectors (upper) and corresponding eigenvalues (lower). The gray dashed line indicates the random expectation of Recovery (Eq. S11).

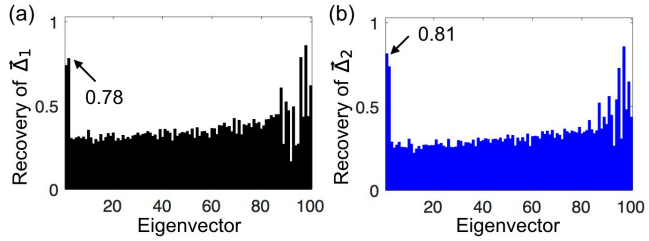

Fig S6. Performance of ICOD for the two-sector case in Fig. S5, without applying ICA.

In Fig. 3(b), one observes a slight decrease of performance of ICOD plus ICA for double selection with overlapping sectors. Does this arise from increasing sector size or from increasing overlap? As expected from Eqs. 8 and 9, Fig. S7(a) shows that Recovery does not fall off with increased sector size. Thus, we tested whether larger sector overlaps could reduce Recovery. Fig. S7(b) shows that this is indeed the case for sequence ensembles subject to two selections each with a fixed sector size of 20, but with different numbers of overlapping sites. However, the reduction of Recovery is quite modest, as even for 100% overlap, Recovery remains above 0.9. It is interesting to note that, independent of sector size and overlap, Recovery decreases faster for double selection than for single selection at large relative biases (see Figs. 3 and S7).

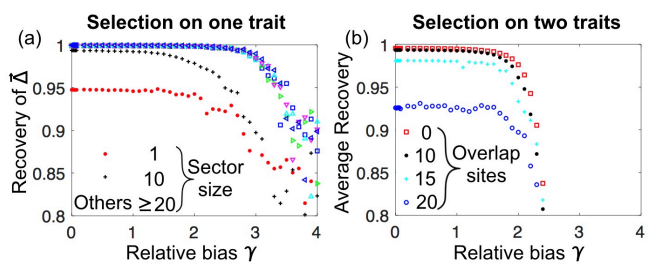

Fig S7. Performance of ICOD for different sector sizes and sector overlaps. (a) Selection on a single trait with varying sector size. Recovery is shown as a function of relative selection bias $\gamma \equiv (T^* - \langle T \rangle)/\sqrt{\langle (T - \langle T \rangle)^2 \rangle}$ for sectors of size 1, 10, 20, 40, 60, 80, and 100 out of 100 sequence sites (*cf.* Fig. 3(a)). Recovery is almost perfect for sectors of size larger than 10, but is substantially lower for sector size 1, which violates the criteria $\Delta_l \ll \sqrt{\sum_{l'} \Delta_{l'}^2}$. (b) Simultaneous selection on two traits with different degrees of sector overlap. For each selection, the sector size is 20 out of 100 sequence sites, and the overlap varies from 0 to 20 sites. The average Recovery for $\vec{\Delta}_1$ and $\vec{\Delta}_2$ is shown as a function of relative selection bias. The data in (b) is averaged over 20 realizations of $\vec{\Delta}$s.

4.2 Robustness of functional sectors to different forms of selection

To assess the robustness of physical sectors to forms of selection other than the simple Gaussian selection window of Eqs. 2-3, we generated ensembles of 50,000 random binary sequences as above, and used synthetically generated mutational effects, with 20 sector sites out of $L = 100$ total sites. As before, the mutational effects of the sector sites and the non-sector sites were sampled, respectively, from zero-mean Gaussian distributions with standard deviations 20 and 1.

We first addressed selection for sequences with an additive trait T above a threshold T_t. We thus considered the selected ensembles of sequences such that the value of the trait $T(\vec{S}) = \vec{S} \cdot \vec{\Delta}$ is larger than a threshold T_t, and we varied this threshold. Fig. 4 in the main text demonstrates that the corresponding sectors are identified by ICOD just as well as those resulting from our initial Gaussian selection window.

We also successfully applied ICOD to various other forms of selection. In Fig. S8, we used the quartic fitness function:

$$w(\vec{S}) = -\frac{\kappa_1}{4}\left(\sum_{l=1}^{L} \Delta_l S_l - T^*\right)^4 , \tag{S47}$$

with $\kappa_1 = (10/\sum_l \Delta_l^2)^2$, instead of our initial quadratic fitness function (see Eq. S14 and Eq. 3 in the main text) and we weighted sequences using the Boltzmann distribution in Eq. 2. Finally, in Fig. S9, we considered the selected

ensembles of sequences such that the value of the trait $T(\vec{S}) = \vec{S} \cdot \vec{\Delta}$ is between $T^* - \eta/2$ and $T^* + \eta/2$, where η is the width of the selection window. In Fig. S9, we used $\eta = 0.6\sqrt{\sum_l \Delta_l^2}$.

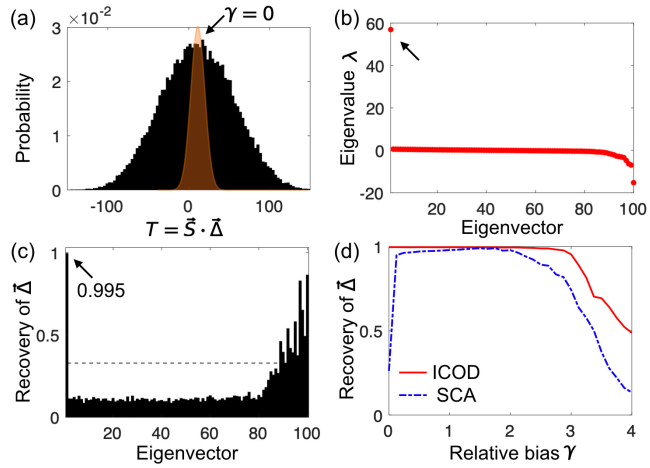

Fig S8. Identification of sectors that result from quartic selection. (a) Histogram of the additive trait $T(\vec{S}) = \vec{S} \cdot \vec{\Delta}$ for randomly sampled sequences where 0 and 1 are equally likely at each site. Sequence length is $L = 100$, mutational effects are synthetically generated with 20 sector sites. Sequences are selectively weighted using a quartic window (orange) around T^*. Selection is shown for $T^* = \langle T \rangle$, or equivalently $\gamma = 0$, in terms of the relative selection bias $\gamma \equiv (T^* - \langle T \rangle)/\sqrt{\langle (T - \langle T \rangle)^2 \rangle}$. (b) Eigenvalues of the ICOD-modified inverse covariance matrix $\tilde{C}^{-1}$ (Eq. 10) of the selected sequences for $\gamma = 0$. (c) Recovery of $\vec{\Delta}$ for all eigenvectors of $\tilde{C}^{-1}$ for $\gamma = 0$. Gray dashed line: random expectation of Recovery. (d) Recovery of $\vec{\Delta}$ for ICOD and for SCA as functions of the relative selection bias γ. The data in (d) is averaged over 100 realizations of $\vec{\Delta}$.

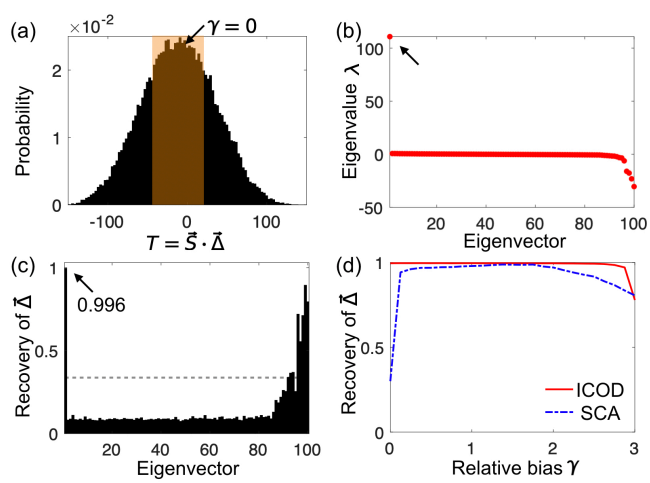

Fig S9. Identification of sectors that result from rectangular-window selection. (a) Histogram of the additive trait $T(\vec{S}) = \vec{S} \cdot \vec{\Delta}$ for randomly sampled sequences where 0 and 1 are equally likely at each site. Sequence length is $L = 100$, mutational effects are synthetically generated with 20 sector sites. Sequences are selected if they have a trait value $T^* - \eta/2 < T(\vec{S}) < T^* + \eta/2$ (orange shaded region). Selection is shown for $T^* = \langle T \rangle$, or equivalently $\gamma = 0$, in terms of the relative selection bias $\gamma \equiv (T^* - \langle T \rangle)/\sqrt{\langle (T - \langle T \rangle)^2 \rangle}$. (b) Eigenvalues of the ICOD-modified inverse covariance matrix $\tilde{C}^{-1}$ (Eq. 10) of the selected sequences for $\gamma = 0$. (c) Recovery of $\vec{\Delta}$ for all eigenvectors of $\tilde{C}^{-1}$ for $\gamma = 0$. Gray dashed line: random expectation of Recovery. (d) Recovery of $\vec{\Delta}$ for ICOD and for SCA as functions of the relative selection threshold γ. The data in (d) is averaged over 100 realizations of $\vec{\Delta}$.

These results confirm the robustness of our approach to different plausible forms of selection.

4.3 Multiple states per site and alternative gauge choice

In Section 3.2 above, we described how to generalize from binary sequences to sequences with q possible states at each site. Correspondingly, we now generalize the ICOD method to higher values of q. Since we are interested in extracting the single-site mutational effects $\Delta_l(\alpha_l)$ with respect to a reference state at each site, we can simply set to zero the diagonal blocks of C^{-1} in Eq. S44, yielding the modified inverse covariance matrix

$$\tilde{C}^{-1}_{kl}(\alpha_k, \alpha_l) = (1 - \delta_{kl}) \sum_{i=1}^{N} \kappa_i \Delta_{i,k}(\alpha_k) \Delta_{i,l}(\alpha_l) \,, \tag{S48}$$

for the case of multiple selections, or more simply for a single selection

$$\tilde{C}^{-1}_{kl}(\alpha_k, \alpha_l) = (1 - \delta_{lk})\, \kappa \Delta_k(\alpha_k) \Delta_l(\alpha_l). \tag{S49}$$

This equation generalizes Eq. 10 obtained for $q = 2$ in the main text. As in that case, the first eigenvector of $\tilde{C}^{-1}$ (associated with the largest eigenvalue) should accurately report the single-site mutational effects $\Delta_k(\alpha_k)$. Indeed, Fig. S10 in the main text shows that this generalized version of ICOD performs very well on synthetic data generated for the case $q = 21$ relevant to real protein sequences. Note that in the reference-sequence gauge, Recovery generalizes naturally to the q-state model as

$$\text{Recovery} = \frac{\sum_{l,\alpha_l} |\nu_l(\alpha_l) \Delta_l(\alpha_l)|}{\sqrt{\sum_{l,\alpha_l} \nu_l(\alpha_l)^2} \sqrt{\sum_{l,\alpha_l} \Delta_l(\alpha_l)^2}}, \tag{S50}$$

where the sums over states α_l do not include the reference state at each site.

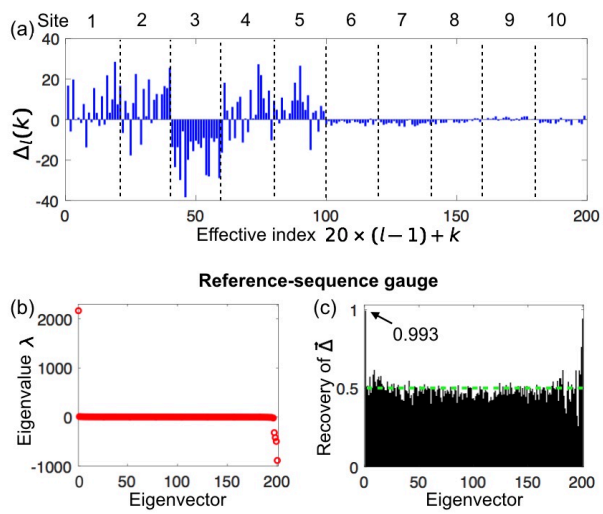

Fig S10. Performance of ICOD on synthetic sequence data with $q = 21$ possible states at each site. (a) Mutational effects $\Delta_l(k)$ with respect to a reference sequence, chosen to be state 21 at every site. The mutational effect at $q = 21$ is not shown. Note that while mutational effects are initially generated from a Gaussian distribution, *relative* mutational effects (calculated with respect to the reference sequence) can have a systematic bias at each site l. (b) Eigenvalues of the ICOD-modified inverse covariance matrix $\tilde{C}^{-1}$ defined in Eq. S49. (c) Recovery of $\vec{\Delta}$ (see Eq. S50). The green dashed line indicates the random expectation of Recovery (Eq. S11).

While the reference-sequence gauge is convenient and allows a clear interpretation of the mutational effects, other gauge choices are possible. For instance, in the DCA literature, the zero-sum (or Ising) gauge is often employed [61, 62]. In this gauge, the couplings satisfy

$$\sum_{\alpha} e_{ij}(\alpha, \beta) = \sum_{\beta} e_{ij}(\alpha, \beta) = 0 \,, \tag{S51}$$

Qualitatively, the gauge degree of freedom means that contributions to the Hamiltonian in Eq. S32 can be shifted between the fields and the couplings [36]. In DCA, the focus is on identifying the dominant two-body interactions,

so one does not want the couplings to include contributions that can be accounted for by the one-body fields [60]. The zero-sum gauge satisfies this condition because it minimizes the Frobenius norms of the couplings

$$\|e_{ij}\| = \sqrt{\sum_{\alpha,\beta=1}^{q} \left[e_{ij}(\alpha,\beta)\right]^2}\,. \tag{S52}$$

Hence, the zero-sum gauge attributes the smallest possible fraction of the energy in Eq. S32 to the couplings, and the largest possible fraction to the fields [36, 61]. In order to transform to the zero-sum gauge defined in Eq. S51, each coupling $e_{ij}(\alpha,\beta)$ is replaced by

$$\tilde{e}_{ij}(\alpha,\beta) = e_{ij}(\alpha,\beta) - \langle e_{ij}(\zeta,\beta)\rangle_\zeta - \langle e_{ij}(\alpha,\eta)\rangle_\eta + \langle e_{ij}(\zeta,\eta)\rangle_{\zeta,\eta}\,, \tag{S53}$$

where $\langle . \rangle_\zeta$ denotes an average over $\zeta \in \{1, ..., q\}$ [61].

Shifting from the reference-sequence gauge where one state (in our derivations, state q) is taken as a reference at each site to the zero-sum gauge requires the replacement

$$\tilde{\Delta}_l(\alpha) = \Delta_l(\alpha) - \frac{1}{q}\sum_{\beta=1}^{q} \Delta_l(\beta), \tag{S54}$$

The new reference-state-free mutational effects satisfy $\sum_{\beta=1}^{q} \tilde{\Delta}_l(\beta) = 0$, and the associated couplings $\tilde{e}_{lk}(\alpha_l,\alpha_k) = -\kappa\tilde{\Delta}_l(\alpha_l)\tilde{\Delta}_k(\alpha_k)$ (see Eq. S33) are related to the initial ones $e_{lk}(\alpha_l,\alpha_k)$ through Eq. S53.

Importantly, these reference-state-free mutational effects can be used to assess the overall importance of mutations at each particular site in the sequence. To this end, let us introduce the Frobenius norm of the reference-state-free mutational effects:

$$||\Delta_l|| = \sqrt{\sum_{\beta=1}^{q} \left(\tilde{\Delta}_l(\beta)\right)^2}. \tag{S55}$$

This quantity, which we refer to as the "site significance", measures the overall importance of mutational effects at site l. In order to assess site significances from an ensemble of sequences, without investigating the impact of each particular mutation at each site, one can apply the zero-sum gauge to the ICOD-modified inverse covariance matrix (see Eq. S49), and compute the Frobenius norm of each 20×20 block associated to each pair of sites (i, j) according to Eq. S52. The first eigenvector of this compressed $L \times L$ matrix accurately reports the mutational significance of each site, as illustrated in Fig. S11. Specifically, it yields a high Recovery of site significances as defined in Eq. S55 (see Fig. S11(c)), and it successfully predicts the most important sites, i.e. the sector sites, in our synthetic data (see Fig. S11(d)).

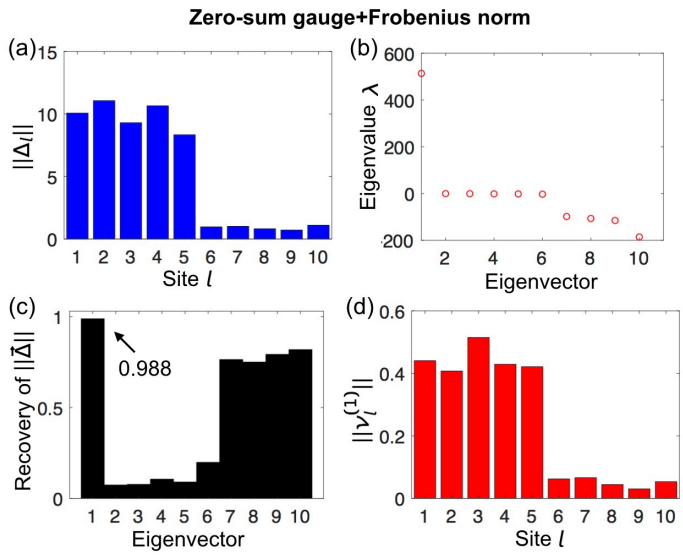

Fig S11. Assessing site significance for synthetic sequence data. The same synthetic data as in Fig. S10 (with $q = 21$ possible states at each site) is used. (a) Significance $||\Delta_l||$ of each site l, computed directly by applying Eqs. S54 and S55 to the mutational effects $\Delta_l(k)$ shown in Fig. S10(a). (b) Eigenvalues of the compressed $(L \times L)$ ICOD-modified inverse covariance matrix, calculated by applying the zero-sum gauge to the ICOD-modified inverse covariance matrix (see Eq. S49), and by computing the Frobenius norm of each 20×20 block associated to each pair of sites (i, j) according to Eq. S52. (c) Recovery of site significances $||\vec{\Delta}||$ from each eigenvector of the compressed ICOD-modified inverse covariance matrix (see panel (a) and Eq. S50) (d) Estimated site significances computed from the first eigenvector $\vec{\nu}^{(1)}$ of the compressed ICOD-modified inverse covariance matrix.

4.4 Pseudocounts

As pseudocounts are often necessary to regularize real sequence data, and as a high fraction of pseudocounts is generally used in DCA, we consider here whether the ICOD method is robust to the addition of pseudocounts.

Until now, we used only raw empirical frequencies obtained from sequence data. For instance, one-body frequencies were obtained by counting the number of sequences where a given state occured at a given site and dividing by the total number M of sequences in the ensemble. Covariances were computed from the empirical single-site frequencies of occurrence of each state α at each site i, denoted by $f_i^e(\alpha)$, and the empirical two-site frequencies of occurrence of each ordered pair of states (α, β) at each ordered pair of sites (i, j), denoted by $f_{ij}^e(\alpha, \beta)$. Specifically, we obtained the covariance matrix as $C_{ij}(\alpha, \beta) = f_{ij}^e(\alpha, \beta) - f_i^e(\alpha) f_j^e(\beta)$ [36].

To avoid issues arising from limited sample size, such as states that never appear at some sites (which present mathematical difficulties, e.g. a non-invertible covariance matrix [37]), one can introduce pseudocounts via a parameter Λ [36–38, 63]. The one-site frequencies f_i then become

$$f_i(\alpha) = \frac{\Lambda}{q} + (1 - \Lambda) f_i^e(\alpha)\,, \tag{S56}$$

where q is the number of states per site. Similarly, the two-site frequencies f_{ij} become

$$f_{ij}(\alpha, \beta) = \frac{\Lambda}{q^2} + (1 - \Lambda) f_{ij}^e(\alpha, \beta) \text{ if } i \neq j\,, \tag{S57}$$

$$f_{ii}(\alpha, \beta) = \frac{\Lambda}{q} \delta_{\alpha\beta} + (1 - \Lambda) f_{ii}^e(\alpha, \beta) = f_i(\alpha) \delta_{\alpha\beta}\,. \tag{S58}$$

These pseudocount corrections are uniform (i.e. they have the same weight $1/q$ for all states), and their influence relative to the raw empirical frequencies can be tuned through the parameter Λ. In DCA, a high value of f Λ has been found to improve contact prediction: typically $\Lambda \approx 0.5$ [37, 38, 51]. Note that the correspondence of Λ with the parameter λ in Refs. [37, 38, 63] is obtained by setting $\Lambda = \lambda/(\lambda + M)$.

From these quantities, we define the pseudocount-corrected covariances

$$C'_{ij}(\alpha, \beta) = f_{ij}(\alpha, \beta) - f_i(\alpha) f_j(\beta)\,. \tag{S59}$$

We show in Fig. S12 that adding pseudocounts as high as $\Lambda = 0.3$ still allows for accurate extraction of mutational effects (Recovery 0.96) and provides a reliable prediction of sector sites.

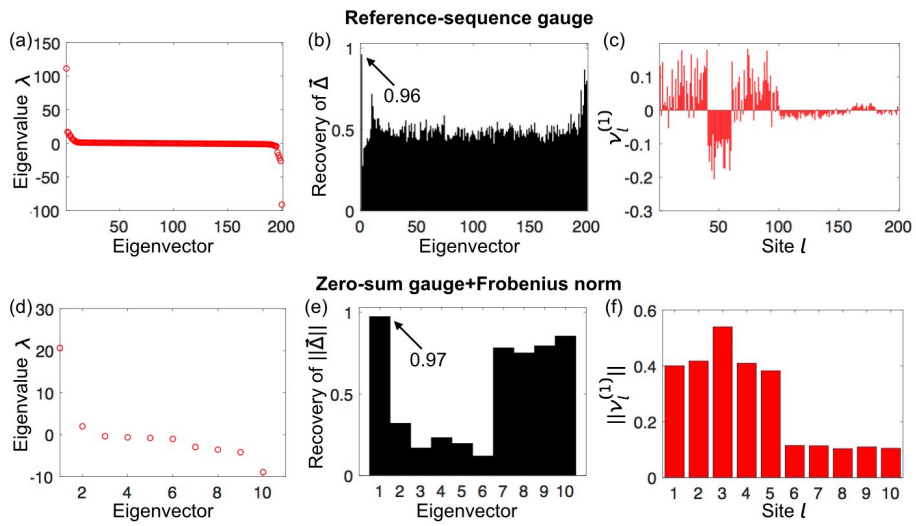

Fig S12. Effect of pseudocounts on ICOD performance on synthetic sequence data with $q = 21$ possible states at each site. The same synthetic data as in Fig. S10 and S11 is used, but here pseudocounts are employed, with weight $\Lambda = 0.3$. (a-c) Similar analysis as in Fig. S10: (a) Eigenvalues of the ICOD-modified inverse covariance matrix. (b) Recovery of $\vec{\Delta}$ from each eigenvector of the ICOD-modified inverse covariance matrix. (c) First eigenvector of the ICOD-modified inverse covariance matrix. (d-f) Similar analysis as in Fig. S11: (d) Eigenvalues of the compressed ICOD-modified inverse covariance matrix. (e) Recovery of site significances $||\vec{\Delta}||$ from each eigenvector of the compressed ICOD-modified inverse covariance matrix. (f) Estimated site significances computed from the first eigenvector of the compressed ICOD-modified inverse covariance matrix.

5 Performance of SCA

5.1 Analytical estimates for $\langle S_l \rangle_*$ and $C_{ll'}$ for a single selection with binary sequences

Protein sectors were first discovered from sequence data using a PCA-based method called Statistical Coupling Analysis (SCA) [4, 7]. Interestingly, in SCA, sectors are found from the eigenvectors associated to the largest eigenvalues, while in ICOD they are found from the (modified) eigenvectors associated to the smallest eigenvalues. This difference stems from the fact that SCA and ICOD do not start from the same matrix. For binary sequences, SCA uses the absolute value of a conservation-weighted covariance matrix, $\tilde{C}^{(\mathrm{SCA})}_{ll'} = |\phi_l C_{ll'} \phi_{l'}|$ (see main text and Ref. [4]). When all amino-acid states are accounted for, SCA compresses each block of the conservation-weighted matrix corresponding to two sites to obtain one positive value, e.g. the Frobenius norm of the block [7]. Conversely, ICOD employs the regular covariance matrix, suppressing the diagonal blocks of its inverse at the last step before diagonalization. To better understand the performance of SCA in recovering the site-dependent mutational effects associated with a selective constraint, it is helpful to have analytical estimates for the average mutant fraction $\langle S_l \rangle_*$ at each site l and the covariance matrix $C_{ll'}$ for an ensemble of binary sequences obtained from a single selection using vector of mutational effects $\vec{\Delta}$. To this end, we provide the following two ansatzes:

$$\langle S_l \rangle_* - \langle S_l \rangle \approx (T^* - \langle T \rangle) \frac{\Delta_l}{\sum_l \Delta_l^2}, \tag{S60}$$

$$C_{ll'} \approx \begin{cases} -\frac{\Delta_l \Delta_{l'} \sigma_l^2 \sigma_{l'}^2}{\sum_l \Delta_l^2 \sigma_l^2}, & l \neq l' \\ \sigma_l^2, & l = l', \end{cases} \tag{S61}$$

where $\sigma_l^2 = \langle S_l^2 \rangle_* - \langle S_l \rangle_*^2 = \langle S_l \rangle_* (1 - \langle S_l \rangle_*)$ represents the variance of S_l. Recall that $S_l \in \{0, 1\}$, where 0 is the reference state and 1 the mutant state, and that $\langle \cdot \rangle_*$ denotes ensemble averages over the selectively weighted subset of sequences, while $\langle \cdot \rangle$ denotes averages over the unselected (unweighted) ensemble.

Although we have not proven these two ansatzes, numerical tests (Fig. S13) have verified these two relations for ensembles generated from a $\vec{\Delta}$ with multiple sites of comparably large mutational effects so as not to be dominated by a single site, i.e., $\Delta_l/\sqrt{\sum_{l'}\Delta_{l'}^2} \ll 1$ for any l. As a counterexample, the $\vec{\Delta}$ from our elastic network model does not satisfy this condition.

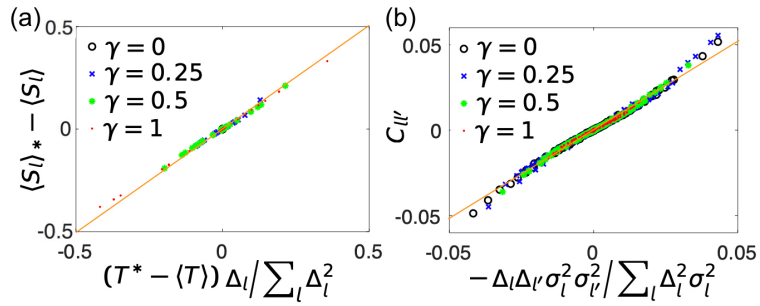

Fig S13. Numerical verification of the ansatzes in Eq. S60 and Eq. S61. We generate a sequence ensemble by considering four values of relative selection bias $\gamma \equiv (T^* - \langle T\rangle)/\sqrt{\langle (T-\langle T\rangle)^2\rangle} = 0, 0.25, 0.5, 1$ and for each case we use a synthetic $\vec{\Delta}$ with a sector size of 20. (a) Numerically computed average bias of the mutant fractions $\langle S_l\rangle_* - \langle S_l\rangle$. Here, $\langle S_l\rangle = 0.5$ for the unselected ensemble. (b) Numerically computed covariances $C_{ll'}$. The results in (a,b) compare well with the analytical predictions (orange lines), provided that $\Delta_l/\sqrt{\sum_l' \Delta_l'^2} \ll 1$ for any l. For each case, 10^6 random sequences were generated to minimize noise from sampling.

5.2 Analysis of the SCA method

Here, we provide a detailed analysis of the SCA method from Refs. [4,7]. Following these references, the reweighting factor is chosen to be

$$\phi_l = \frac{\partial D\left(\langle S_l\rangle_*, \langle S_l\rangle\right)}{\partial \langle S_l\rangle_*}, \tag{S62}$$

where, for each site l, $D\left(\langle S_l\rangle_*, \langle S_l\rangle\right)$ is the Kullback-Leibler divergence between the distribution of mutant fractions for the selected sequences and the background distribution:

$$D\left(\langle S_l\rangle_*, \langle S_l\rangle\right) = \langle S_l\rangle_* \log\frac{\langle S_l\rangle_*}{\langle S_l\rangle} + \left(1-\langle S_l\rangle_*\right)\log\frac{1-\langle S_l\rangle_*}{1-\langle S_l\rangle}. \tag{S63}$$

In our case, the background distribution is obtained from the unselected sequence ensemble, for which $\langle S_l\rangle = 0.5$. Hence, we have

$$\phi_l = \log\left[\frac{\langle S_l\rangle_*\left(1-\langle S_l\rangle\right)}{\langle S_l\rangle\left(1-\langle S_l\rangle_*\right)}\right], \tag{S64}$$

as illustrated in Fig. S14(a). In the regime of relatively weak conservation, i.e. when $\langle S_l\rangle$ is not close to 0 or 1, and $|\langle S_l\rangle_* - \langle S_l\rangle| \ll \langle S_l\rangle$, a first-order expansion yields

$$\phi_l \approx \frac{\langle S_l\rangle_* - \langle S_l\rangle}{\langle S_l\rangle(1-\langle S_l\rangle)}, \tag{S65}$$

as shown in Fig. S14(b). Employing the ansatz (S60) in this regime, we obtain

$$\phi_l \propto (T^* - \langle T\rangle)\Delta_l. \tag{S66}$$

This relation is verified in Fig. S14(c) for a sequence ensemble generated with a synthetic Δ_l. Hence, the SCA reweighting factor carries information about Δ_l as long as $T^* \neq \langle T\rangle$. In this regime, information about conservation (namely ϕ_l) should thus be sufficient to recover mutational effects and sectors. This was indeed found to be the case for some real proteins with a single sector [52]. However, when the selection bias, $T^* - \langle T\rangle$, is small, random noise due to finite sampling will typically swamp this relationship.

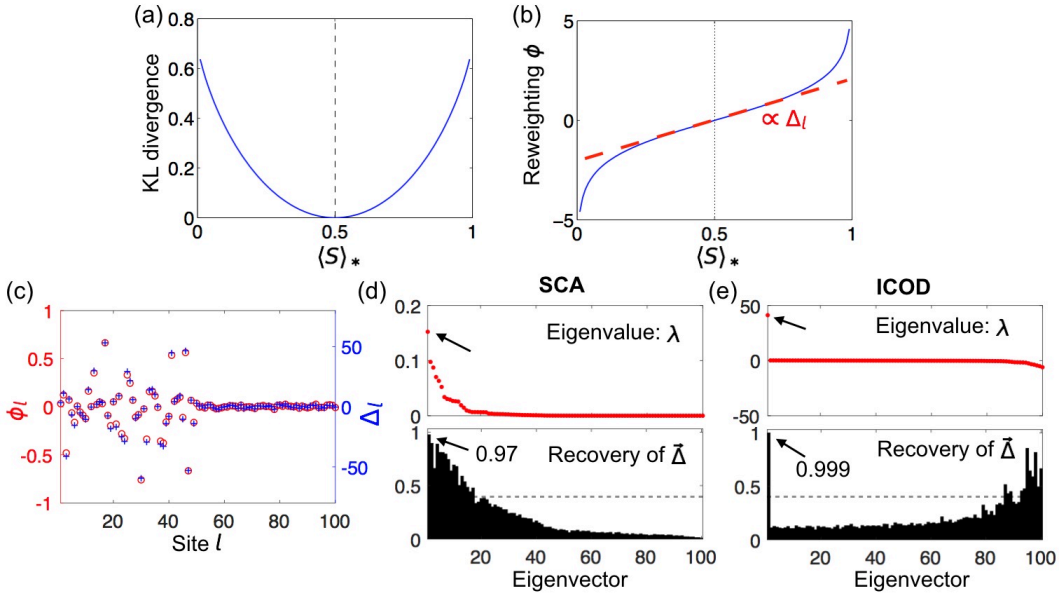

Fig S14. Underpinnings of Recovery of mutational effect vector $\vec{\Delta}$ by SCA. (a) Kullback-Leibler divergence versus mutant fraction $\langle S \rangle_*$ for background mutant fraction $\langle S \rangle = 0.5$. (b) Reweighting factor ϕ as a function of mutant fraction $\langle S \rangle_*$ for background mutant fraction $\langle S \rangle = 0.5$. (c) Reweighting factor ϕ_l and synthetic Δ_l for an ensemble of sequences generated with a single selection at relative selection bias $\gamma = 1$. $\vec{\Delta}$ was generated with the first 50 sites as sector sites, and 50,000 sequences were employed, as in most of our examples using ICOD (see above). (d-e) Performance of SCA and ICOD for this ensemble, respectively. In computing Recovery using SCA, we use the normalized vector $\sqrt{\nu_l^{(j)}}$ to predict $\vec{\Delta}$. The gray dashed lines in (d) and (e) indicate the random expectation of Recovery (Eq. S11).

In Refs. [4,7], the first eigenvectors of the conservation-reweighted SCA covariance matrix, $\tilde{C}^{(\mathrm{SCA})}_{ll'} = |\phi_l C_{ll'} \phi_{l'}|$, were used to find sectors from sequence data. How does the first eigenvector of $\tilde{C}^{(\mathrm{SCA})}$ relate to the mutational effect vector $\vec{\Delta}$? Utilizing both Eq. S61 and Eq. S66, and assuming $T^* \neq \langle T \rangle$, we obtain

$$\tilde{C}^{(\mathrm{SCA})}_{ll'} \propto \begin{cases} \Delta_l^2 \Delta_{l'}^2 \sigma_l^2 \sigma_{l'}^2, & l \neq l' \\ \Delta_l^2 \sigma_l^2, & l = l'. \end{cases} \tag{S67}$$

Apart from the diagonal, the matrix is approximately proportional to the tensor product of $\Delta_l^2 \sigma_l^2$ with itself. If we neglect the contribution from the diagonal elements of $\tilde{C}^{(\mathrm{SCA})}$, the first eigenvector $\vec{\nu}^{(1)}$ satisfies

$$\nu_l^{(1)} \propto \Delta_l^2 \sigma_l^2. \tag{S68}$$

Eq. S68 explains why $\sqrt{\nu_l^{(1)}}$ carries information about Δ_l. In Fig. S14(d), Recovery using SCA (and Eq. 8 with $\sqrt{\nu_l^{(1)}}$ instead of $\nu_l^{(1)}$) is 0.97, which remains lower than Recovery using ICOD, which is 0.999 here. Besides, Fig. S15 illustrates that Recovery of $\vec{\Delta}$ by SCA is much better using $\sqrt{\nu_l^{(1)}}$ than $\nu_l^{(1)}$.

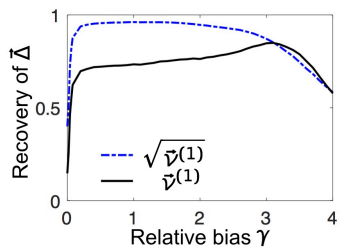

Fig S15. Recovery of $\vec{\Delta}$ from the first SCA eigenvector using $\vec{\nu}^{(1)}$ or $\sqrt{\vec{\nu}^{(1)}}$. The sequence data are the same as used for the blue curves in Fig. 3(a). As suggested by Eq. S68, use of the square root of $\vec{\nu}^{(1)}$ significantly improves Recovery.

5.3 Comparison between ICOD and SCA

In the main text, we compared the performance of ICOD and SCA with respect to Recovery of mutational-effect vectors $\vec{\Delta}$ in synthetic data (see Fig. 3). We found that ICOD performs well over a broader range of relative biases γ than SCA. The failure of SCA at biases close to zero can be explained by the fact that the conservation weights ϕ_l then vanish (see Eq. S66). A further example of the failure of SCA for non-biased selections is given by the case studied in Fig. S5, where we considered two selections, a biased one associated to $\vec{\Delta}_1$ and a non-biased one associated to $\vec{\Delta}_2$. Fig. S16 shows that SCA recovers $\vec{\Delta}_1$ well, but performs badly for $\vec{\Delta}_2$, while ICOD recovers both of them very well (see Fig. S5).

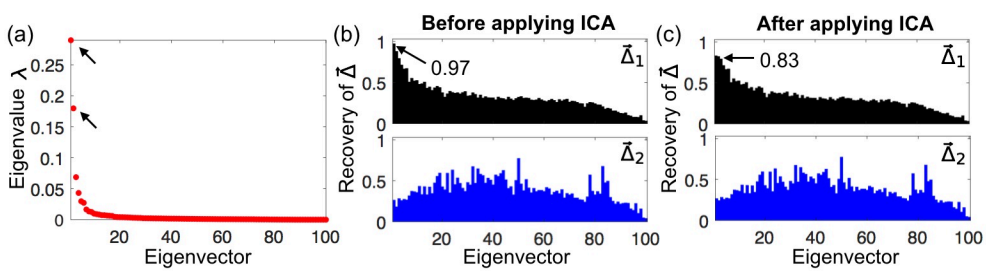

Fig S16. Performance of SCA for the double selection from Fig. S5. (a) Eigenvalues. (b) Before applying ICA, the first eigenvector has high Recovery of $\vec{\Delta}_1$, but no eigenvector has substantial Recovery of $\vec{\Delta}_2$. This difference matches our observation that SCA performs well for selections of intermediate bias, but not for unbiased selections. (c) Applying ICA on the first two eigenvectors does not improve Recovery.

While the comparison of Recovery favors ICOD, SCA was originally used to identify sectors (in our model, sites with important mutational effects under a given selection) rather than to recover complete mutational effect vectors

$\vec{\Delta}$. Hence, in Fig. S17, we compare the ability of ICOD and SCA to predict the n sites with the largest mutational effects. Note that this comparison is independent of whether we use $\vec{\nu}^{(1)}$ or $\sqrt{\vec{\nu}^{(1)}}$ as the predictor in SCA, since the square-root function is increasing and preserves order. Using this criterion, we again find that ICOD performs well over a broad range of relative biases γ, while SCA only works well for sequences selected under moderate biases.

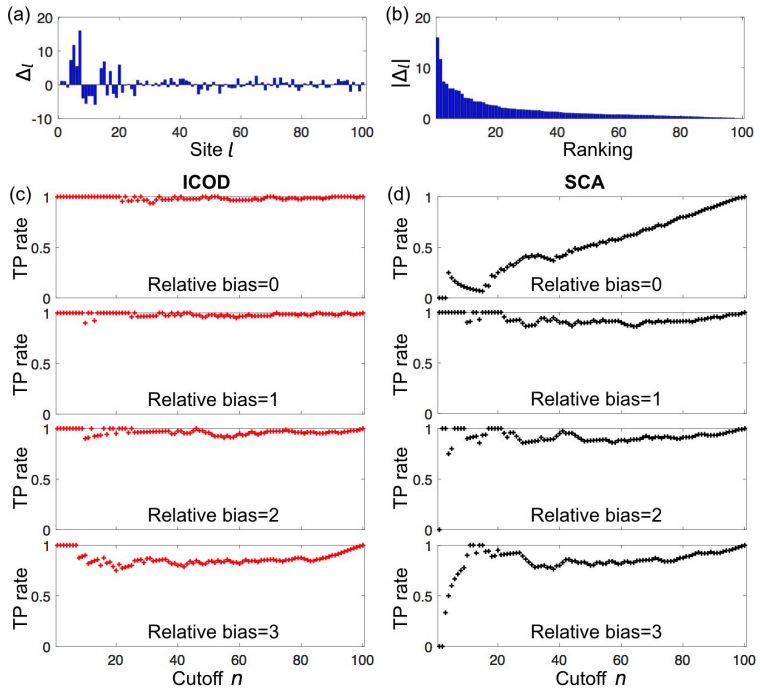

Fig S17. Comparison of sector-site identification by ICOD and SCA (see also Fig. 3). We use the synthetic $\vec{\Delta}$ in (a) to selectively weight 5,000 random sequences at four relative bias values $\gamma \equiv (T^* - \langle T \rangle)/\sqrt{\langle (T - \langle T \rangle)^2 \rangle} = 0, 1, 2, 3$ and test the ability of ICOD or SCA to correctly predict the sites with the n largest mutational effects. (b) Magnitudes of mutational effects of $\vec{\Delta}$ by rank. (c-d). True Positive (TP) rates obtained by taking the first eigenvector $\vec{\nu}^{(1)}$ from either ICOD or SCA, generating a ranked list of sites of descending $|\nu_l^{(1)}|$ at each site l, and computing the fraction of the top n sites in this predicted ordering that are also among the top n sites of the actual ordering of mutational effect magnitudes $|\Delta_l|$. The effect of relative bias γ on Recovery is shown in Fig. 3. (c) As expected, the prediction of ICOD is very good under all relative biases. (d) On the other hand, SCA does not perform well at the smallest or largest relative biases.

6 Performance of a method based on the generalized Hopfield model

As mentioned in the main text, we also compared ICOD with another PCA-based approach developed in Ref. [34], which employs an inference method specific to the generalized Hopfield model. For L Ising spins ($s_l \in \{-1, 1\}$ for $1 \leq l \leq L$), the Hamiltonian of the generalized Hopfield model reads (see Eq. 6 in Ref. [34])

$$H(\vec{s}) = -\sum_{l=1}^{L} h_l \, s_l - \frac{1}{2L} \sum_{i=1}^{N} \left(\sum_{l=1}^{L} \xi_{i,l} \, s_l \right)^2 + \frac{1}{2L} \sum_{i=1}^{N'} \left(\sum_{l=1}^{L} \xi'_{i,l} \, s_l \right)^2 , \tag{S69}$$

where h_l is the local field at site l, while $\vec{\xi}_i = (\xi_{i,1}, \ldots, \xi_{i,L})$ is an attractive pattern and $\vec{\xi'}_i = (\xi'_{i,1}, \ldots, \xi'_{i,L})$ is a repulsive pattern. Here there are N attractive patterns and N' repulsive ones. In our model, in the single-selection case, the fitness of a sequence $\vec{s}$ in the Ising representation reads (see above, Sec. 3.1.1, Eq. S15)

$$w(\vec{s}) = -\frac{\kappa}{2} \left(\sum_{l=1}^{L} D_l s_l - \alpha \right)^2 = -\frac{\kappa}{2} \left[\left(\sum_{l=1}^{L} D_l s_l \right)^2 - 2\alpha \sum_{l=1}^{L} D_l s_l + \alpha^2 \right] , \tag{S70}$$

with $D_l = \Delta_l/2$ and $\alpha = T^* - \sum_l D_l$. Recalling that fitnesses and Hamiltonians have opposite signs, a comparison of Eqs. S69 and S70 shows that $\vec{\Delta}$ plays the part of a repulsive pattern in the two-body coupling terms, with the exact

correspondence given by $\vec{\xi}' = \vec{\Delta}\sqrt{\kappa L}/2$. Note that in our model the local fields are proportional to the components of $\vec{\Delta}$.

Ref. [34] proposed a method to infer attractive and repulsive patterns from data generated using a generalized Hopfield model Eq. S69. Introducing the correlation matrix G, which is related to the covariance matrix C through

$$G_{ll'} = \frac{C_{ll'}}{\tilde{\sigma}_l \tilde{\sigma}_{l'}}\,, \tag{S71}$$

where $\tilde{\sigma}_l^2 = \langle s_l^2 \rangle_* - \langle s_l \rangle_*^2 = 1 - \langle s_l \rangle_*^2$. Ref. [34] found, to lowest order, the following approximation for a single repulsive pattern $\vec{\xi}'$ (see Eq. 9 in Ref. [34]):

$$\xi'_l \approx \sqrt{L\left(\frac{1}{\lambda^{(L)}} - 1\right)}\,\frac{\nu_l^{(L)}}{\tilde{\sigma}_l}, \tag{S72}$$

where $\lambda^{(L)}$ is the smallest (last) eigenvalue of the correlation matrix G and $\nu_l^{(L)}$ is the associated eigenvector. This yields

$$\Delta_l \propto \frac{\nu_l^{(L)}}{\tilde{\sigma}_l}. \tag{S73}$$

Inference of $\vec{\Delta}$ based on Eq. S73 is referred to as GHI (for Generalized Hopfield Inference) below.

GHI performs very well for the sequence ensembles from the elastic network model used in Fig. 1 and Fig. 2 (Fig. S18). Importantly, just as for simple PCA and for ICOD (see main text), the top Recovery is obtained for the (modified) bottom eigenvector of the covariance matrix, consistently with $\vec{\Delta}$ being a repulsive pattern, but the large-eigenvalue modes also contain some information about $\vec{\Delta}$ (Fig. S18).

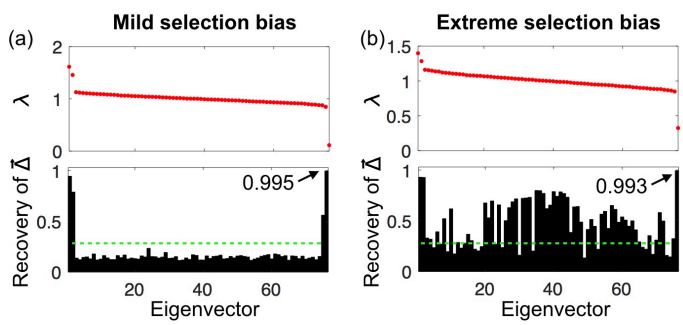

Fig S18. Performance of GHI on sequence ensembles generated with our elastic-network $\vec{\Delta}$. (a) Eigenvalues of G and Recovery under mild selection bias, as in Fig. 1 in the main text. (b) Eigenvalues of G and Recovery under extreme selection bias, as in Fig. 2 in the main text. The green dashed lines in (a,b) indicate the random expectation of Recovery (Eq. S11).

In Fig. S19, we systematically compare all methods discussed in our work to recover $\vec{\Delta}$ from sequence data under various selection biases, using different sector sizes, for selectively weighted ensembles of 50,000 random sequences. We focus on the case of a single selection and compare Recovery of $\vec{\Delta}$ according to:

- ICOD, using the first eigenvector of the modified inverse covariance matrix $\tilde{C}^{-1}$ (see main text, Eq. 10)
- PCA, using the last principal component of the data (last eigenvector of the covariance matrix, see main text)
- SCA, using the first eigenvector of the absolute value of a conservation-weighted covariance matrix, $\tilde{C}_{ll'}^{(\mathrm{SCA})} = |\phi_l C_{ll'} \phi_{l'}|$ (see main text and Ref. [4])
- GHI, using the reweighted last eigenvector of the correlation matrix (see Eqs. S71 and S73).

Overall, ICOD and GHI perform best. For small selection biases, all methods perform accurately, except SCA, which fails when selection bias vanishes, as explained above. When the sector size is small compared to the sequence length L (Fig. S19 (a-d)), GHI performs a little bit better than ICOD for relatively small selection biases (however Recovery remains $\gtrsim 95\%$ with ICOD). Conversely, GHI is significantly outperformed by ICOD for relatively large selection bias, and the performance of PCA and SCA falls off quite rapidly in this regime. The performances of ICOD, PCA, and GHI become similar when the sector size becomes comparable to the sequence length (Fig. S19 (e, f)).

We further find that GHI is more sensitive to the size of the sequence ensemble than ICOD, although it becomes the most accurate for very large dataset sizes (see Fig. S20). The performance of ICOD is quite robust to dataset size. Note that PCA outperforms other methods when the data size becomes very small (Fig. S20, number of sequences = 500).

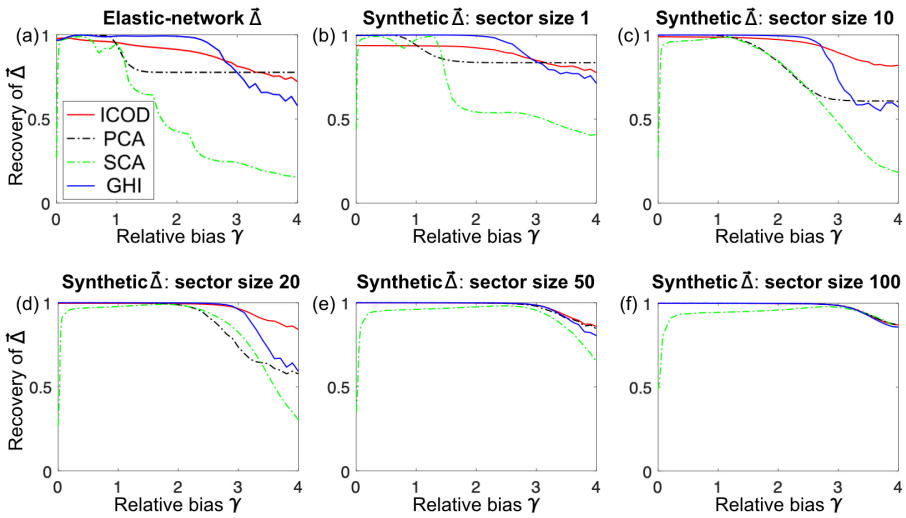

Fig S19. Comparing Recovery of different methods for various $\vec{\Delta}$s. Here, GHI refers to inference based on Eq. S73. Curves are obtained by averaging over 100 realizations, each for an ensemble of 50,000 random sequences. For synthetic $\vec{\Delta}$s, each realization corresponds to a new $\vec{\Delta}$.

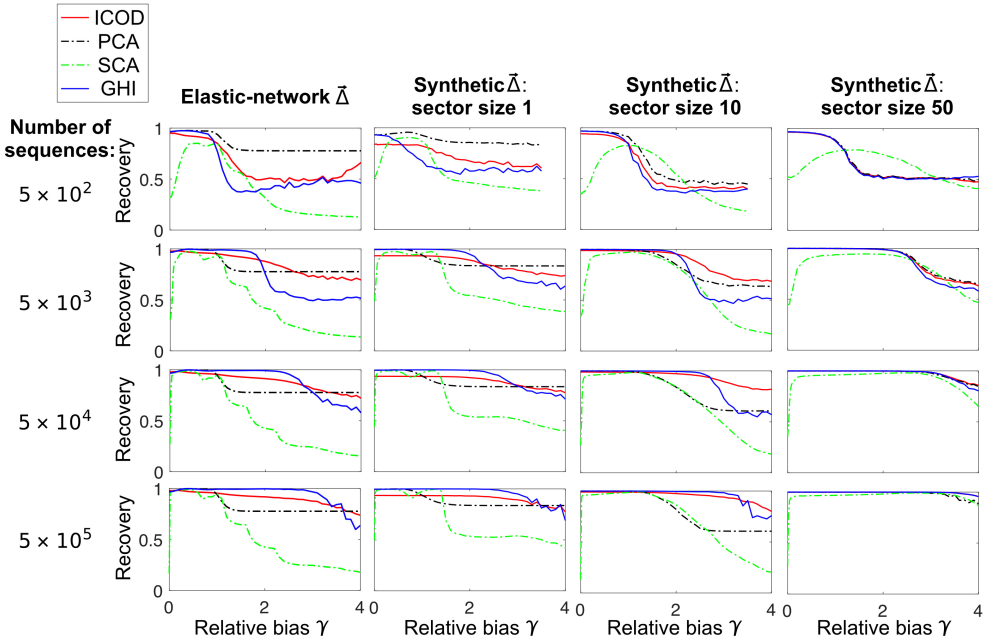

Fig S20. Effect of dataset size on Recovery of $\vec{\Delta}$. Selectively reweighted ensembles of 5×10^2, 5×10^3, 5×10^4, and 5×10^5 random sequences are generated for the elastic-network $\vec{\Delta}$ and synthetic $\vec{\Delta}$s with sector sizes 1, 10, and 50. All results are averaged over 100 realizations, except those using 5×10^5 sequences, where only 5 realizations were used. For synthetic $\vec{\Delta}$s, each realization employs a different $\vec{\Delta}$ with the same sector size. For the case of 500 sequences, some Recoveries were not computed at high biases due to numerical instabilities.

Overall, we find that GHI is very well suited to infer $\vec{\Delta}$ from very large synthetic datasets. However, ICOD is more robust to variation of dataset size and to selection bias, which should be an advantage in the application to real protein data.

7 Application of ICOD to a multiple sequence alignment of PDZ domains

Our general physical model for sectors provides insights into the statistical signatures of sectors in sequence data. In particular, we have found that the primary signature of physical sectors lies in the modes associated with the smallest eigenvalues of the covariance matrix, even though there is often additional signal from these sectors in the large eigenvalue modes, as studied more conventionally, e.g. in SCA. The success of ICOD on synthetic data demonstrates that information about sectors can indeed be extracted from the small eigenvalue modes of the covariance matrix.

How well does ICOD perform on real sequence data? Here, we apply ICOD to an actual alignment of sequences of PDZ domains from the Pfam database (`https://pfam.xfam.org/`) containing 24,934 sequences of length $L = 79$ (corresponding to sites 313-391 in the numbering in Fig. 2 of Ref. [6]). In Ref. [6], sites important for the specific binding of PDZ to peptide ligands were identified experimentally via complete single-site mutagenesis. In particular, 20 sites showing particularly high mutational effects were deemed functionally significant [6]. It was further shown that 15 among the 20 sector amino acids found by SCA (i.e. 75%) were also functionally significant sites.

In order to compute the empirical covariance matrix of the data, we first removed sites with more than 15% gaps (11 sites out of 79). To eliminate the confounding effects of very rare residues at particular sites, we used a pseudocount weight $\Lambda = 0.02$.

Next, we performed both SCA and ICOD using this empirical covariance matrix:

- For SCA, we computed the conservation reweighting factors as in Refs. [4, 7], using the background frequency values from Ref. [4]. We compressed the conservation-reweighted covariance matrix using the Frobenius norm, and we focused on the first eigenvector of this reweighted and compressed covariance matrix in order to predict sector sites. Finally, we took the square root of each component of this eigenvector to predict the mutational effect at each site (see above, Section 5.2, and Ref. [7]).
- For ICOD, we inverted the covariance matrix and set its diagonal blocks to zero, thus obtaining the ICOD-modified inverse covariance matrix (see Eq. S49). Next, we computed the Frobenius norm of each 20×20 block associated to each pair of sites (i, j) according to Eq. S52. The magnitude of the l-th component of the first eigenvector $\vec{\nu}^{(1)}$ of this compressed $L \times L$ matrix, denoted by $||\nu_l^{(1)}||$, is the ICOD prediction of the overall mutational effect at site l (see above, Section 4.3, especially Fig. S12). Since mutational effects were experimentally measured with respect to the wild-type residues [6], we used as reference the wild-type sequence of the PDZ domain employed in Fig. 1 and retained this reference-sequence gauge to perform ICOD, thus allowing direct comparison to experiments.

We then assessed the ability of SCA and ICOD to predict experimentally-measured mutational effects [6]. Specifically, we compared SCA and ICOD predictions to the overall mutational effects corresponding to the Frobenius norm of the experimentally-measured residue-specific mutational effects $\Delta_l(\alpha)$ with $\alpha \in \{1, \ldots, 20\}$:

$$||\Delta_l|| = \sqrt{\sum_{\alpha=1}^{q} (\Delta_l(\alpha))^2}, \tag{S74}$$

which is the counterpart in the reference-sequence gauge of the "site significance" introduced in the zero-sum gauge in Eq. S55. The ability of SCA and ICOD to identify the sites with the experimentally most important mutational effects is shown in Fig. 5 in the main text. Here, we discuss the impact of parameters on these results. Fig. S21 shows the effect of varying the cutoff for removal of sites with a large proportion of gaps. As illustrated in panel (a), sites with a fraction of gaps larger than a cutoff are discarded. Many of these sites are on the edges of the PDZ domain, and tend to be less conserved. Fig. S21(b) shows that ICOD performance is robust to variations of this cutoff within a reasonable range. We have chosen a cutoff of 15% in the rest of this analysis.

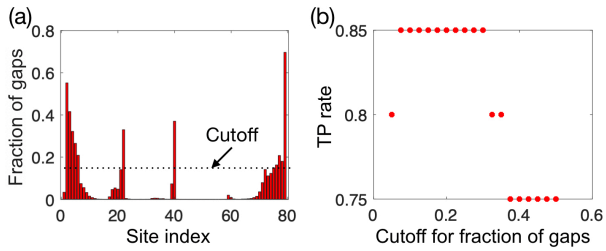

Fig S21. Impact of varying the cutoff for removal of sites with a large fraction of gaps. (a) Fraction of sequences that have a gap at each site. Sites with a fraction of gaps larger than the cutoff shown by the dashed line are discarded in the rest of our analysis. (b) Impact of varying the gap-fraction cutoff on the performance of ICOD. A pseudocount weight of $\Lambda = 0.02$ is used.

Fig. S22 shows the effect of varying the pseudocount weight, both for ICOD and for SCA. Panels (a) and (b) show the TP rate, defined as the fraction of the top 20 predicted sites that are among the 20 sites with the largest experimentally-determined mutational effects. Panels (c) and (d) show the Pearson correlation between ICOD or SCA predictions of mutational effects and the corresponding experimental measurements. Both ICOD and SCA identify experimentally-important sites significantly better than random expectation over the whole range of pseudocounts shown. However, ICOD performs best with small but nonzero pseudocount weights (panels (a) and (c)), while the performance of SCA is more robust to changing the pseudocount weight (panels (b) and (d)).

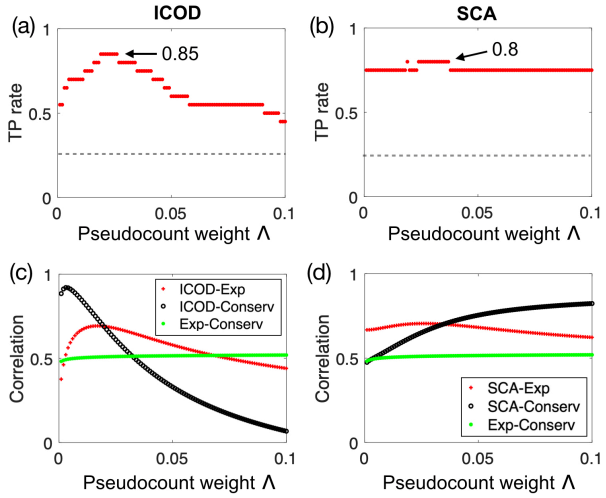

Fig S22. Impact of the pseudocount weight Λ on the performance of ICOD and SCA. (a) Fraction of the 20 top sites predicted by ICOD that are among the 20 sites with the largest experimentally-determined mutational effects ("TP rate") versus pseudocount weight. The TP rate definition is the same as that shown in Fig. 5. Gray dashed line: random expectation for the TP rate, namely 20/68=0.29 (68 sites are left after removing those with a gap fraction larger than the cutoff). (b) Counterpart of (a) for SCA. (c) Pearson correlation: between mutational effects predicted by ICOD and those measured experimentally at each site (ICOD-Exp; data from Fig. 5(c) for $\Lambda = 0.02$); between mutational effects predicted by ICOD and conservation scores ϕ_l (ICOD-Conserv; see Ref. [7] and Eq. S62 in the binary case); and between experimentally measured mutational effects and conservation scores ϕ_l (Exp-Conserv). (d) Counterpart of (b) for SCA. In all panels, a gap-fraction cutoff of 15% is used.

Since residue conservation plays a very important part in the PDZ sector [52], we compared prediction based simply on conservation to those of SCA and ICOD. We employed the conservation scores ϕ_l used in SCA [7], which are a generalization of Eq. S62 to 21 states. Conservation alone identifies 70% of the 20 sites with largest experimentally-determined mutational effects, versus 85% for ICOD (for $\Lambda = 0.02$) and 75% for SCA (see Fig. 5). In addition, the Pearson correlation between conservation scores and experimentally-measured mutational effects is significant (see Fig. S23(c)), even though it is smaller than between ICOD or SCA scores and experimentally-measured mutational effects (see Fig. S23(a-b)). In fact, both ICOD and SCA scores are significantly correlated with conservation scores (see Fig. S23 (d-e)). In the case of SCA, this is not surprising given that conservation scores

are explicitly used to weight the covariance matrix. Interestingly, ICOD naturally identifies these conserved sites as being important. This correlation between ICOD and conservation highlights the ability of ICOD to identify functionally important amino acids in a principled way that only relies on covariance.

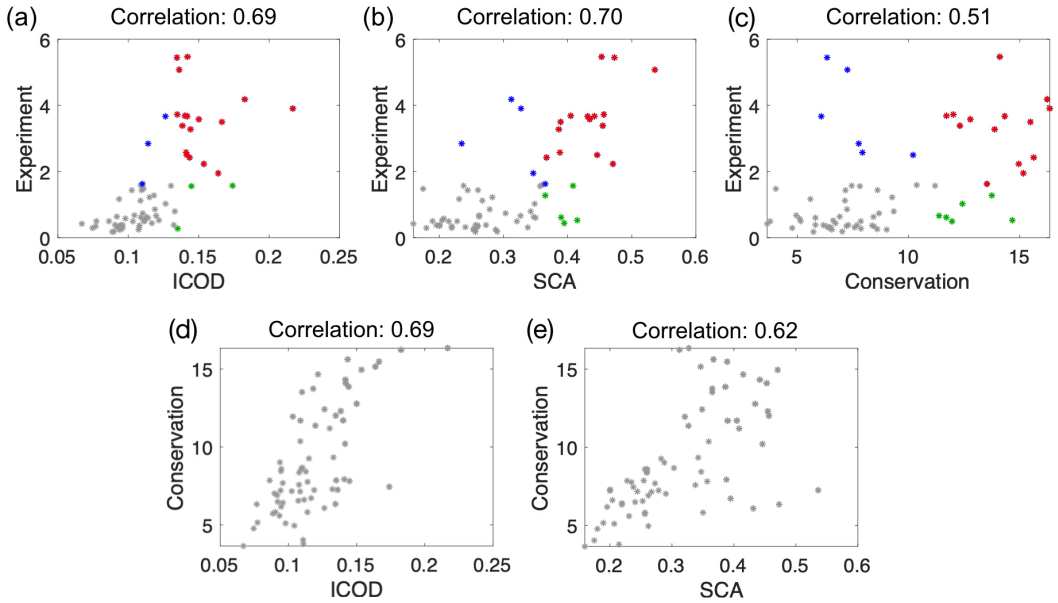

Fig S23. Predicting experimentally-measured mutational effects using ICOD, SCA, or conservation. (a) Experimentally-measured mutational effect versus mutational effect predicted by ICOD for each site of the PDZ sequence. (b) Experimentally-measured mutational effect versus mutational effect predicted by SCA for each site of the PDZ sequence. (c) Experimentally-measured mutational effect versus Conservation score ϕ_l for each site of the PDZ sequence. In panels (a, b, c), to highlight the matches between the top 20 predictions and the top 20 experimentally important sites [6], correct hits are shown in red, false negatives in blue, and false positives in green. (d) Conservation score ϕ_l versus mutational effect predicted by ICOD for each site of the PDZ sequence. (e) Conservation score ϕ_l versus mutational effect predicted by SCA for each site of the PDZ sequence. In all panels, a pseudocount weight $\Lambda = 0.02$ and a gap-fraction cutoff of 15% were used.

References

1. Lockless SW, Ranganathan R. Evolutionarily conserved pathways of energetic connectivity in protein families. Science. 1999;286(5438):295–299.

2. Süel GM, Lockless SW, Wall MA, Ranganathan R. Evolutionarily conserved networks of residues mediate allosteric communication in proteins. Nat Struct Biol. 2003;10(1):59–69.

3. Socolich M, Lockless SW, Russ WP, Lee H, Gardner KH, Ranganathan R. Evolutionary information for specifying a protein fold. Nature. 2005;437(7058):512.

4. Halabi N, Rivoire O, Leibler S, Ranganathan R. Protein sectors: evolutionary units of three-dimensional structure. Cell. 2009;138(4):774–786.

5. Dahirel V, Shekhar K, Pereyra F, Miura T, Artyomov M, Talsania S, et al. Coordinate linkage of HIV evolution reveals regions of immunological vulnerability. Proc Natl Acad Sci USA. 2011;108(28):11530–11535.

6. McLaughlin Jr RN, Poelwijk FJ, Raman A, Gosal WS, Ranganathan R. The spatial architecture of protein function and adaptation. Nature. 2012;491(7422):138.

7. Rivoire O, Reynolds KA, Ranganathan R. Evolution-Based Functional Decomposition of Proteins. PLoS Comput Biol. 2016;12(6):e1004817.

8. DePristo MA, Weinreich DM, Hartl DL. Missense meanderings in sequence space: a biophysical view of protein evolution. Nat Rev Genet. 2005;6(9):678–687.

9. Starr TN, Thornton JW. Epistasis in protein evolution. Protein Sci. 2016;25(7):1204–1218.

10. Otwinowski J, McCandlish DM, Plotkin JB. Inferring the shape of global epistasis. Proc Natl Acad Sci USA. 2018;115(32):E7550–E7558.

11. Bahar I, Lezon TR, Yang LW, Eyal E. Global dynamics of proteins: bridging between structure and function. Annu Rev Biophys. 2010;39:23–42.

12. Cunningham AD, Colavin A, Huang KC, Mochly-Rosen D. Coupling between Protein Stability and Catalytic Activity Determines Pathogenicity of G6PD Variants. Cell Rep. 2017;18(11):2592–2599.

13. Wylie CS, Shakhnovich EI. A biophysical protein folding model accounts for most mutational fitness effects in viruses. Proc Natl Acad Sci USA. 2011;108(24):9916–9921.

14. Kryazhimskiy S, Rice DP, Jerison ER, Desai MM. Microbial evolution. Global epistasis makes adaptation predictable despite sequence-level stochasticity. Science. 2014;344(6191):1519–1522.

15. Posfai A, Zhou J, Plotkin JB, Kinney JB, McCandlish DM. Selection for Protein Stability Enriches for Epistatic Interactions. Genes (Basel). 2018;9(9).

16. Zheng W, Brooks BR, Thirumalai D. Low-frequency normal modes that describe allosteric transitions in biological nanomachines are robust to sequence variations. Proc Natl Acad Sci USA. 2006;103(20):7664–7669.

17. Lukman S, Grant GH. A network of dynamically conserved residues deciphers the motions of maltose transporter. Proteins. 2009;76(3):588–597.

18. Saldano TE, Monzon AM, Parisi G, Fernandez-Alberti S. Evolutionary Conserved Positions Define Protein Conformational Diversity. PLoS Comput Biol. 2016;12(3):e1004775.

19. Granata D, Ponzoni L, Micheletti C, Carnevale V. Patterns of coevolving amino acids unveil structural and dynamical domains. Proc Natl Acad Sci USA. 2017;114(50):E10612.

20. De Los Rios P, Cecconi F, Pretre A, Dietler G, Michielin O, Piazza F, et al. Functional dynamics of PDZ binding domains: a normal-mode analysis. Biophys J. 2005;89(1):14–21.

21. Delarue M, Sanejouand YH. Simplified normal mode analysis of conformational transitions in DNA-dependent polymerases: the elastic network model. J Mol Biol. 2002;320(5):1011–1024.

22. Zheng W, Doniach S. A comparative study of motor-protein motions by using a simple elastic-network model. Proc Natl Acad Sci USA. 2003;100(23):13253–13258.

23. Yan L, Ravasio R, Brito C, Wyart M. Architecture and coevolution of allosteric materials. Proc Natl Acad Sci USA. 2017; p. 201615536.

24. Tlusty T, Libchaber A, Eckmann JP. Physical Model of the Genotype-to-Phenotype Map of Proteins. Phys Rev X. 2017;7:021037.

25. Flechsig H. Design of Elastic Networks with Evolutionary Optimized Long-Range Communication as Mechanical Models of Allosteric Proteins. Biophys J. 2017;113(3):558–571.

26. Dutta S, Eckmann JP, Libchaber A, Tlusty T. Green function of correlated genes in a minimal mechanical model of protein evolution. Proc Natl Acad Sci USA. 2018;115(20):E4559–E4568.

27. Rocks JW, Pashine N, Bischofberger I, Goodrich CP, Liu AJ, Nagel SR. Designing allostery-inspired response in mechanical networks. Proc Natl Acad Sci USA. 2017;114(10):2520–2525.

28. Yan L, Ravasio R, Brito C, Wyart M. Principles for Optimal Cooperativity in Allosteric Materials. Biophys J. 2018;114(12):2787–2798.

29. Bravi B, Ravasio R, Brito C, Wyart M. Direct Coupling Analysis of Epistasis in Allosteric Materials. arXiv:181110480. 2018;.

30. Doyle DA, Lee A, Lewis J, Kim E, Sheng M, MacKinnon R. Crystal structures of a complexed and peptide-free membrane protein–binding domain: molecular basis of peptide recognition by PDZ. Cell. 1996;85(7):1067–1076.

31. Hung AY, Sheng M. PDZ domains: structural modules for protein complex assembly. J Biol Chem. 2002;277(8):5699–5702.

32. Berman HM, Bhat TN, Bourne PE, Feng Z, Gilliland G, Weissig H, et al. The Protein Data Bank and the challenge of structural genomics. Nat Struct Biol. 2000;7 Suppl:957–959.

33. Mann JK, Barton JP, Ferguson AL, Omarjee S, Walker BD, Chakraborty A, et al. The fitness landscape of HIV-1 gag: advanced modeling approaches and validation of model predictions by in vitro testing. PLoS Comput Biol. 2014;10(8):e1003776.

34. Cocco S, Monasson R, Sessak V. High-dimensional inference with the generalized Hopfield model: principal component analysis and corrections. Phys Rev E. 2011;83(5 Pt 1):051123.

35. Cocco S, Monasson R, Weigt M. From principal component to direct coupling analysis of coevolution in proteins: Low-eigenvalue modes are needed for structure prediction. PLOS Comput Biol. 2013;9(8):e1003176.

36. Weigt M, White RA, Szurmant H, Hoch JA, Hwa T. Identification of direct residue contacts in protein-protein interaction by message passing. Proc Natl Acad Sci USA. 2009;106(1):67–72.

37. Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, Sander C, et al. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. Proc Natl Acad Sci USA. 2011;108(49):E1293–E1301.

38. Marks DS, Colwell LJ, Sheridan R, Hopf TA, Pagnani A, Zecchina R, et al. Protein 3D structure computed from evolutionary sequence variation. PLoS ONE. 2011;6(12):e28766.

39. Morcos F, Schafer NP, Cheng RR, Onuchic JN, Wolynes PG. Coevolutionary information, protein folding landscapes, and the thermodynamics of natural selection. Proc Natl Acad Sci USA. 2014;111(34):12408–12413.

40. Dwyer RS, Ricci DP, Colwell LJ, Silhavy TJ, Wingreen NS. Predicting functionally informative mutations in Escherichia coli BamA using evolutionary covariance analysis. Genetics. 2013;195(2):443–455.

41. Cheng RR, Morcos F, Levine H, Onuchic JN. Toward rationally redesigning bacterial two-component signaling systems using coevolutionary information. Proc Natl Acad Sci USA. 2014;111(5):E563–571.

42. Cheng RR, Nordesjo O, Hayes RL, Levine H, Flores SC, Onuchic JN, et al. Connecting the Sequence-Space of Bacterial Signaling Proteins to Phenotypes Using Coevolutionary Landscapes. Mol Biol Evol. 2016;33(12):3054–3064.

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

44. Barton JP, Goonetilleke N, Butler TC, Walker BD, McMichael AJ, Chakraborty AK. Relative rate and location of intra-host HIV evolution to evade cellular immunity are predictable. Nat Commun. 2016;7:11660.

45. Hopf TA, Ingraham JB, Poelwijk FJ, Scharfe CP, Springer M, Sander C, et al. Mutation effects predicted from sequence co-variation. Nat Biotechnol. 2017;35(2):128–135.

46. Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, Sander C, et al. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. Proc Natl Acad Sci USA. 2011;108(49):E1293–1301.

47. Malinverni D, Marsili S, Barducci A, De Los Rios P. Large-Scale Conformational Transitions and Dimerization Are Encoded in the Amino-Acid Sequences of Hsp70 Chaperones. PLoS Comput Biol. 2015;11(6):e1004262.

48. Bitbol AF, Dwyer RS, Colwell LJ, Wingreen NS. Inferring interaction partners from protein sequences. Proc Natl Acad Sci USA. 2016;113(43):12180–12185.

49. Gueudre T, Baldassi C, Zamparo M, Weigt M, Pagnani A. Simultaneous identification of specifically interacting paralogs and interprotein contacts by direct coupling analysis. Proc Natl Acad Sci USA. 2016;113(43):12186–12191.

50. Plefka T. Convergence condition of the TAP equation for the infinite-ranged Ising spin glass model. J Phys A: Math Gen. 1982;15(6):1971–1978.

51. Bitbol AF, Dwyer RS, Colwell LJ, Wingreen NS. Inferring interaction partners from protein sequences. Proc Natl Acad Sci USA. 2016;113(43):12180–12185.

52. Teşileanu T, Colwell LJ, Leibler S. Protein sectors: statistical coupling analysis versus conservation. PLoS Comput Biol. 2015;11(2):e1004091.

53. Hyvärinen A, Karhunen J, Oja E. Independent Component Analysis. John Wiley and Sons; 2001.

54. Hansen LK, Larsen J, Kolenda T. Blind Detection of Independent Dynamic Components. In: IEEE International Conference on Acoustics, Speech, and Signal Processing 2001. vol. 5; 2001. p. 3197–3200.

55. Salinas VH, Ranganathan R. Coevolution-based inference of amino acid interactions underlying protein function. Elife. 2018;7.

56. Casari G, Sander C, Valencia A. A method to predict functional residues in proteins. Nat Struct Biol. 1995;2(2):171–178.

57. Qin C, Colwell LJ. Power law tails in phylogenetic systems. Proc Natl Acad Sci USA. 2018;115(4):690–695.

58. Rivoire O. Elements of coevolution in biological sequences. Phys Rev Lett. 2013;110(17):178102.

59. Thouless DJ, Anderson PW, Palmer RG. Solution of'solvable model of a spin glass'. Philos Mag. 1977;35(3):593–601.

60. Ekeberg M, Hartonen T, Aurell E. Fast pseudolikelihood maximization for direct-coupling analysis of protein structure from many homologous amino-acid sequences. J Comput Phys. 2014;276:341–356.

61. Ekeberg M, Lövkvist C, Lan Y, Weigt M, Aurell E. Improved contact prediction in proteins: using pseudolikelihoods to infer Potts models. Phys Rev E. 2013;87(1):012707.

62. Baldassi C, Zamparo M, Feinauer C, Procaccini A, Zecchina R, Weigt M, et al. Fast and accurate multivariate Gaussian modeling of protein families: predicting residue contacts and protein-interaction partners. PLoS ONE. 2014;9(3):e92721.

63. Procaccini A, Lunt B, Szurmant H, Hwa T, Weigt M. Dissecting the specificity of protein-protein interaction in bacterial two-component signaling: orphans and crosstalks. PLoS ONE. 2011;6(5):e19729.