

RESEARCH ARTICLE

Phylogenetic correlations can suffice to infer protein partners from sequences

Guillaume Marmier^{1,2}, Martin Weigt², Anne-Florence Bitbol^{1*}

1 Sorbonne Université, CNRS, Laboratoire Jean Perrin (UMR 8237), F-75005 Paris, France, **2** Sorbonne Université, CNRS, Institut de Biologie Paris-Seine, Laboratoire de Biologie Computationnelle et Quantitative (LCQB, UMR 7238), F-75005 Paris, France

* anne-florence.bitbol@sorbonne-universite.fr



Abstract

Determining which proteins interact together is crucial to a systems-level understanding of the cell. Recently, algorithms based on Direct Coupling Analysis (DCA) pairwise maximum-entropy models have allowed to identify interaction partners among paralogous proteins from sequence data. This success of DCA at predicting protein-protein interactions could be mainly based on its known ability to identify pairs of residues that are in contact in the three-dimensional structure of protein complexes and that coevolve to remain physicochemically complementary. However, interacting proteins possess similar evolutionary histories. What is the role of purely phylogenetic correlations in the performance of DCA-based methods to infer interaction partners? To address this question, we employ controlled synthetic data that only involve phylogeny and no interactions or contacts. We find that DCA accurately identifies the pairs of synthetic sequences that share evolutionary history. While phylogenetic correlations confound the identification of contacting residues by DCA, they are thus useful to predict interacting partners among paralogs. We find that DCA performs as well as phylogenetic methods to this end, and slightly better than them with large and accurate training sets. Employing DCA or phylogenetic methods within an Iterative Pairing Algorithm (IPA) allows to predict pairs of evolutionary partners without a training set. We further demonstrate the ability of these various methods to correctly predict pairings among real paralogous proteins with genome proximity but no known direct physical interaction, illustrating the importance of phylogenetic correlations in natural data. However, for physically interacting and strongly coevolving proteins, DCA and mutual information outperform phylogenetic methods. We finally discuss how to distinguish physically interacting proteins from proteins that only share a common evolutionary history.

OPEN ACCESS

Citation: Marmier G, Weigt M, Bitbol A-F (2019) Phylogenetic correlations can suffice to infer protein partners from sequences. *PLoS Comput Biol* 15(10): e1007179. <https://doi.org/10.1371/journal.pcbi.1007179>

Editor: Björn Wallner, Linköping University, SWEDEN

Received: June 10, 2019

Accepted: September 25, 2019

Published: October 14, 2019

Copyright: © 2019 Marmier et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Matlab implementations of the MI-IPA, the DCA-IPA and the Mirrortree-IPA on our standard HK-RR dataset are freely available at <https://doi.org/10.5281/zenodo.1421781>, <https://doi.org/10.5281/zenodo.1421861> and <https://doi.org/10.5281/zenodo.3377592> respectively.

Funding: The authors thank Institut de Biologie Paris-Seine (IBPS) at Sorbonne Université for funding via a Collaborative Grant (Action Incitative) to AFB and MW. MW also acknowledges funding by the EU H2020 research and innovation program

Author summary

Many biologically important protein-protein interactions are conserved over evolutionary time scales. This leads to two different signals that can be used to computationally predict interactions between protein families and to identify specific interaction partners. First, the shared evolutionary history leads to highly similar phylogenetic relationships between

MSCA-RISE-2016 under grant agreement No. 734439 InferNet. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

interacting proteins of the two families. Second, the need to keep the interaction surfaces of partner proteins biophysically compatible causes a correlated amino-acid usage of interface residues. Employing simulated data, we show that the shared history alone can be used to detect partner proteins. Similar accuracies are achieved by algorithms comparing phylogenetic relationships and by methods based on Direct Coupling Analysis (DCA), which are primarily known for their ability to detect the second type of signal. Using natural sequence data, we show that in cases with shared evolutionary history but without known physical interactions, both methods work with similar accuracy, while for some physically interacting systems, DCA and mutual information outperform phylogenetic methods. We propose methods allowing both to predict interactions between protein families and to find interacting partners among paralogs.

Introduction

The vast majority of cellular processes are carried out by interacting proteins. Functional interactions between proteins allow multi-protein complexes to properly assemble, and ensure the specificity of signal transduction pathways. Hence, mapping functional protein-protein interactions is a crucial question in biology. Since high-throughput experiments remain challenging [1], an attractive alternative is to exploit the growing amount of sequence data in order to identify functional protein-protein interaction partners.

The amino-acid sequences of interacting proteins are correlated, both because of evolutionary constraints arising from the need to maintain physico-chemical complementarity among contacting amino-acids, and because of shared evolutionary history. The first type of correlations has recently received substantial interest, both within single proteins and across protein partners. Global statistical models [2, 3] built from the observed sequence correlations using the maximum entropy principle [2, 4], and assuming pairwise interactions, known as Direct Coupling Analysis (DCA) [5], have been used with success to determine three-dimensional protein structures from sequences [6–8], to analyze mutational effects [9–12] and conformational changes [13, 14], to find residue contacts between known interaction partners [5, 15–21], and most recently to predict interaction partners among paralogs from sequence data [22, 23]. Similar global statistical models have also revealed functional relationships in other contexts [24, 25]. The success of DCA-based approaches at predicting protein-protein interactions [22, 23] could originate only from correlations between residues that are in direct contact in the three-dimensional protein complex structure, thus needing to maintain physico-chemical complementarity. However, additional correlations arise in protein sequences due to their common evolutionary history, i.e. phylogeny [26–28], even in the absence of structural constraints. Functionally related protein families [29], especially interacting ones [30] tend to have similar phylogenies, and methods directly based on phylogeny and on sequence similarity [31–35], in particular the Mirrortree method [31, 34, 35] allow to predict which protein families interact. The similarities in the phylogenies of interacting protein families can arise from the coevolution of residues in structural contact, but also from more global shared evolutionary pressures, resulting in similar evolutionary rates [36–40], and from shared evolutionary history unrelated to constraints, including common timing of speciation and gene duplication events [39]. While being detrimental to the identification of contacting residues by DCA [5, 6, 28], these additional sources of signal can aid the identification of interaction partners. Accordingly, a method based on mutual information (MI) was recently shown to slightly

outperform the DCA-based one [41]. MI includes all types of statistical dependence between the sequences of interacting partners.

To what extent do purely phylogenetic correlations contribute to the prediction of interaction partners from sequences by DCA? If sequences only share a common evolutionary history, i.e. in the absence of functional constraints, how do DCA-based methods compare to phylogenetic methods? Answering these questions is important to understand the reasons of the success of DCA-based methods, and will open the way to developing new methods that combine useful information from both phylogeny and contacts. To address these questions, we generate controlled synthetic data that only involve phylogeny, in the absence of functional constraints. Our DCA-based method correctly identifies pairs of synthetic “sequences” that share evolutionary history, even without any training set, thanks to an Iterative Pairing Algorithm (IPA). (Strikingly, this high predictive power is obtained in the absence of real couplings from interactions, from purely phylogenetic correlations). On this synthetic dataset, we find that the DCA-based IPA and a phylogeny-based IPA reach similar performances, with DCA slightly outperforming the phylogenetic method for large training sets. We then show examples of natural proteins without known direct physical interactions but with shared evolutionary history that can be accurately paired by our various methods, thus illustrating the importance of phylogenetic correlations in real data. For a pair of actually interacting and strongly coevolving protein families, we find that DCA and MI substantially outperform phylogenetic methods. Finally, we propose methods to predict protein-protein interactions from the level of protein families to that of paralogs.

Methods

Synthetic data generation

We generate controlled synthetic data where “sequences” are modeled as strings of binary variables (bits) taking values 0 or 1. In real protein sequences, each site can feature 21 states (20 amino acids, plus the alignment gap), but binary models where the consensus or reference amino acid is denoted by 0 and mutant states by 1 retain all conceptual ingredients, and have proved useful to identify sectors of collectively correlated amino acids [27, 42], as well as to predict fitness landscapes from sets of closely related proteins [43]. Our synthetic sequences are evolved along a phylogenetic tree represented by a branching process with random mutations, in the absence of any constraint stemming from interactions or function. Hence, all correlations in this synthetic data arise from shared evolutionary history (and finite-size noise). The data generation process is illustrated in Fig 1.

Specifically, we consider perfect binary trees for simplicity. The ancestral chain, composed of uniformly randomly distributed bits, is duplicated, giving rise to two chains, and mutations are performed independently in the two duplicate chains: each mutation changes the state of one uniformly randomly chosen bit of the chain. Then, the new chains are duplicated again, and so on. We employ two different models for the occurrence of mutations. In the simplest model, a fixed number of mutations per total chain length is performed along each branch of the tree, i.e. between two duplication steps. In a more realistic model, the number of mutations per branch is drawn in a Poisson distribution with fixed mean. After a given number n of duplication steps (“generations”, representing ancestry in terms of speciation or gene duplication events), a final dataset of 2^n chains is obtained (see Fig 1). In practice, we take $n \leq 12$ to have tractable datasets of a few thousands of final chains. Throughout, we perform inference using these final chains, which correspond to contemporary sequences in a natural dataset. In each of these chains, the state of each bit is uniformly distributed. However, correlations exist between chains due to their shared evolutionary history. The strength of these correlations




Fig 1. Construction of a synthetic dataset of chains sharing evolutionary history. Starting from a random ancestral chain AB of bits whose two halves A and B are shaded in blue and red, a series of n duplication and mutation steps (“generations”, here $n = 3$) are performed (bold: mutated bits; here 2 bits per chain are mutated at each step), resulting in $2^n = 8$ chains. Species are then constructed randomly, here with $m = 2$ chains per species. Some species are considered as the training set (green), and the other ones constitute the testing set (pink), where the pairings between each chain A and each chain B will be blinded.

<https://doi.org/10.1371/journal.pcbi.1007179.g001>

depends on the number of mutations per branch and on how close the chains are along the phylogenetic tree.

In order to introduce the notion of species in a minimal way, we randomly group chains into sets of equal size m , each representing a species. The m different chains within a species can then be thought of as paralogs, i.e. homologs sharing common ancestry and present in the same genome. In reality, different correlations are expected between the paralogs present in a given species and the orthologs present across species. Later on, we therefore also consider another type of phylogeny that accounts for these effects, and assess the robustness of our conclusions to this variant. Note that the present minimal model with random species is realistic in the case where exchange between species (i.e. horizontal gene transfer) is sufficiently frequent.

We finally cut each chain of the final dataset in two halves of equal length. These halves, denoted by chain A and chain B, thus represent a pair of proteins that possess the same evolutionary history. Next, we blind the pairings for the chains A and B from some species (testing set) and ask whether DCA-, MI- and similarity-based methods are able to pair each A chain with its “evolutionary partner”, namely with the B chain that possesses the same evolutionary history, starting from the known pairs (training set).

Inference methods

We test several inference methods to predict pairings between chains A and B in our synthetic datasets. For each of them, performance is assessed both with a training set and without a training set. In the first case, the parameters defining scores are computed using the training set and employed to pair data in the testing set [3, 15]. In the second case, we employ the Iterative Pairing Algorithm (IPA) developed in Refs. [22, 41] to bootstrap the predictions starting from initial random within-species pairings. Below, we present the various inference methods assuming that there is a training set. The extension to the training set-free case is then performed exactly as described in Refs. [22, 41]. Matlab implementations of the MI-IPA, the

DCA-IPA and the Mirrortree-IPA on our standard HK-RR dataset are freely available at <https://doi.org/10.5281/zenodo.1421781>, <https://doi.org/10.5281/zenodo.1421861> and <https://doi.org/10.5281/zenodo.3377592> respectively.

Training set statistics. To describe the statistics of a training set of synthetic paired chains AB, of total length $2L$ (where L is the length of a chain A or B), we employ the empirical one-site frequencies of each state $\sigma_i \in \{0, 1\}$ at each site $i \in \{1, \dots, 2L\}$, denoted by $f_i(\sigma_i)$, and the two-site frequencies of occurrence of each ordered pair of states (σ_i, σ_j) at each ordered pair of sites (i, j) , denoted by $f_{ij}(\sigma_i, \sigma_j)$. Correlations are then computed as $C_{ij}(\sigma_i, \sigma_j) = f_{ij}(\sigma_i, \sigma_j) - f_i(\sigma_i) f_j(\sigma_j)$.

Pseudocount. When dealing with real protein sequences, pseudocounts are often introduced to avoid mathematical issues such as divergences due to amino-acid pairs that never appear, both with DCA [5–7, 15] and with MI [41]. Introducing a pseudocount weight Λ , which effectively corresponds to adding a fraction Λ of chains with uniformly distributed states, the corrected one-body frequencies read $\tilde{f}_i(\sigma_i) = \Lambda/2 + (1 - \Lambda)f_i(\sigma_i)$. Similarly, the corrected two-body frequencies read $\tilde{f}_{ij}(\sigma_i, \sigma_j) = \Lambda/4 + (1 - \Lambda)f_{ij}(\sigma_i, \sigma_j)$ if $i \neq j$ and $\tilde{f}_{ii}(\sigma_i, \sigma_j) = \delta_{\sigma_i \sigma_j} \Lambda/2 + (1 - \Lambda)f_{ii}(\sigma_i, \sigma_j) = \delta_{\sigma_i \sigma_j} \tilde{f}_i(\sigma_i)$, where $\delta_{\sigma_i \sigma_j} = 1$ if $\sigma_i = \sigma_j$ and 0 otherwise. We investigated the impact of varying Λ on the performance of pairing prediction on our synthetic data using DCA and MI. S1 Fig shows that small nonzero values of Λ perform best for MI while larger ones improve DCA performance. Therefore, in what follows, we always took $\Lambda = 0.015$ for MI and $\Lambda = 0.5$ for DCA, which is the typical value used when applying DCA to real proteins [6, 7].

DCA-based method. In DCA [5–7, 44], one starts from the empirical covariances $C_{ij}(\sigma_i, \sigma_j)$ between all pairs of sites (i, j) , computed on the training set. Importantly, here, we are considering paired chains AB, and i and j range from 1 to the total length $2L$ of such a chain. DCA is based on building a global statistical model from these covariances (and the one-body frequencies) [5–7, 44], through the maximum entropy principle [4]. This results in a $2L$ -body probability distribution P of observing a given sequence $(\sigma_1, \dots, \sigma_{2L})$ that reads $P(\sigma_1, \dots, \sigma_{2L}) = \exp[\sum_{i < j} e_{ij}(\sigma_i, \sigma_j) + \sum_{i=1}^{2L} h_i(\sigma_i)]/Z$, where Z is a normalization constant: this corresponds to the Boltzmann distribution associated to a Potts model with couplings $e_{ij}(\sigma_i, \sigma_j)$ and fields $h_i(\sigma_i)$ [44]. Inferring the couplings and the fields that appropriately reproduce the empirical covariances is a difficult problem, known as an inverse statistical physics problem [45]. Note that these parameters are not all independent due to the gauge degree of freedom, so one can set e.g. $h_i(0) = 0$ and $e_{ij}(0, \sigma_j) = e_{ij}(\sigma_i, 0) = 0$ for all i, j and σ_i, σ_j , thus leaving only $h_i(1)$ and $e_{ij}(1, 1)$ to determine. Within the mean-field approximation, which will be employed throughout, these coupling strengths can be approximated by $e_{ij}(1, 1) = -C_{ij}^{-1}(1, 1)$ [6, 7, 46]. We then transform to the zero-sum (or Ising) gauge, yielding $e_{ij}(0, 1) = e_{ij}(1, 0) = -e_{ij}(0, 0) = -e_{ij}(1, 1) = C_{ij}^{-1}(1, 1)/4$. The interest of this gauge is that it attributes the smallest possible fraction of the energy to the couplings, and the largest possible fraction to the fields [5, 47]. Note that a fully equivalent approach is to consider sequences of Ising spins instead of bits, and to employ an Ising model. Here, we have chosen the Potts model formalism for consistency with protein sequence analysis by DCA.

The effective interaction energy E_{AB} of each possible pair AB in the testing set, constructed by concatenating a chain A and a chain B, can then be assessed via

$$E_{AB} = - \sum_{i=1}^L \sum_{j=L+1}^{2L} e_{ij}(\sigma_i^A, \sigma_j^B). \tag{1}$$

In real proteins, approximately minimizing such a score has proved successful at predicting interacting partners [15, 22]. Note that we only sum over inter-chain pairs (i.e. pairs of sites involving one site in A and one in B) because we are interested in interactions between A and B.

Importantly, DCA was designed to infer actual interactions between contacting amino acids through the couplings e_{ij} [5–7, 44]. By contrast, in the present synthetic data, there are no such interactions, and all correlations have their origin in phylogeny (or finite-size noise). Nevertheless, the DCA-based interaction energy in Eq 1 contains information about these correlations, and we will investigate how well it captures them.

MI-based method. Our method based on Mutual Information (MI) was introduced in [41]. As with DCA, we start by describing the statistics of the training set, which is composed of complete chains AB. For this, we employ the single-site frequencies $\tilde{f}_i(\sigma_i)$ and the two-site frequencies $\tilde{f}_{ij}(\sigma_i, \sigma_j)$ (see above). The pointwise mutual information (PMI) of a pair of states (σ_i, σ_j) at a pair of sites (i, j) is defined as [48–50]:

$$\text{PMI}_{ij}(\sigma_i, \sigma_j) = \log \left[\frac{\tilde{f}_{ij}(\sigma_i, \sigma_j)}{\tilde{f}_i(\sigma_i)\tilde{f}_j(\sigma_j)} \right]. \quad (2)$$

Averaging this quantity over all possible pairs of states yields an estimate of the mutual information (MI) between sites i and j [51]: $\text{MI}_{ij} = \sum_{\sigma_i, \sigma_j} \tilde{f}_{ij}(\sigma_i, \sigma_j) \text{PMI}_{ij}(\sigma_i, \sigma_j)$.

Next, we define a pairing score S_{AB} for each possible pair AB of chains from the testing set as the sum of the PMIs of the inter-chain pairs of sites of this concatenated chain AB (i.e. those that involve one site in chain A and one site in chain B):

$$S_{AB} = \sum_{i=1}^L \sum_{j=L+1}^{2L} \text{PMI}_{ij}(\sigma_i^A, \sigma_j^B). \quad (3)$$

In real proteins, approximately maximizing such a score has proved successful at predicting interacting partners, slightly outperforming DCA [41].

Mirrortree-based method. Methods based only on phylogeny and sequence similarity have been developed to predict protein-protein interactions. In particular, the Mirrortree method quantifies the similarities of distance matrices between the proteins of two families to determine whether they interact [31, 34, 35], and has allowed the successful prediction of protein-protein interactions. This method generally relies on finding one ortholog of the proteins of interest in each species and does not address the question of which paralog of family A interacts with which paralog of family B. However, related approaches have tackled this problem [33, 52–56], which was subsequently directly addressed by the DCA- and MI-based methods of Refs. [22, 23, 41].

We introduce an approach close to the original Mirrortree algorithm [31, 34, 35] that addresses the paralog pairing problem. Specifically, let $\{A_1B_1, \dots, A_MB_M\}$ be the training set, which contains M known pairs of chains. For each chain A of the testing set, we compute the vector $d_A = (d(A, A_1), \dots, d(A, A_M))$ of Hamming distances between A and each chain A of the training set. We also compute an analogous vector d_B for each chain B of the training set. Next, we define a pairing score M_{AB} for each possible pair AB of chains A and B from the testing set as the Pearson correlation ρ between d_A and d_B :

$$M_{AB} = \rho(d_A, d_B). \quad (4)$$

This score thus assigns high values to pairs AB that have highly similar phylogenetic relationships to the training set, hinting towards substantial shared evolutionary history between A and B. It can be used for predicting partnerships exactly as the DCA- and MI-based scores in Eqs 1 and 3. Note that one then aims to maximize M_{AB} , just like S_{AB} , while DCA effective interaction energies E_{AB} should be minimized.

Other methods based on sequence similarity. Because there are many ways to exploit sequence similarity in order to assess shared evolutionary history, we also consider variants beyond our Mirrortree-based method. Specifically, we present results obtained using orthology between pairs, defined as reciprocal best hits in terms of Hamming distances, as well as results obtained by simply employing the Hamming distance of each possible AB pair of the testing set to its closest AB pair in the training set as a pairing score. These methods are detailed and studied in S4 Fig.

Pairing prediction. We employed two different approaches to predict pairings from each of the three scores defined in Eqs 1, 3 and 4. In the first approach, for each chain A, we simply picked the chain B within the same species that optimizes the pairing score. Note that this simple method is asymmetric and allows multiple chains A to be matched with the same chain B. In the second approach, we used the Hungarian algorithm (also known as the Munkres algorithm) [57–59] to find the one-to-one association of each chain A with a chain B that optimizes the sum of the pairing scores within each species.

Results

DCA accurately identifies pairs of chains that only share a common evolutionary history

First, we set out to assess whether DCA can identify pairs of chains AB that only share a common evolutionary history. For this, we generated chains of bits employing a branching process with random mutations, in the absence of any interaction or functional constraint (see [Methods](#) and [Fig 1](#)). The only correlations present among these chains thus arise from shared evolutionary history (and finite-size noise). We first ask whether DCA pairing scores ([Eq 1](#)) learned on a training set of complete chains allow to correctly predict pairs of evolutionary partners in a testing set of chains separated into half chains A and B.

Specifically, we generated data using a phylogenetic tree of 10 generations, with 5 mutations per branch, out of 100 bits in each complete chain AB, thus yielding 1024 chains AB. Given the relatively small number of mutations per branch, many of the resulting chains AB possess substantial similarities arising from their common evolutionary history. Specifically, [S2A Fig](#) shows the histogram of Hamming distances between all B chains in the dataset, featuring a typical fraction 0.3 of sites with different states. The degree of similarity between two given chains arises from their relatedness along the phylogenetic tree used to generate the data. We ordered the chains employing this phylogenetic tree, so that sister chains are closest to one another etc. (see [Fig 2A](#)).

Next, we randomly picked 75% of the chains to form a substantial training set, and inferred a DCA model from this training set (see [Methods](#)). We employed the inferred couplings $e_{ij}(\sigma_i, \sigma_j)$ to compute effective interaction energies E_{AB} ([Eq 1](#)) between all chains A and all chains B of the remaining 25% of the dataset, which constitutes our testing set. The effective interaction energies obtained are shown in [Fig 2B](#). Importantly, the diagonal of the matrix, corresponding to actual evolutionary partners, features small energies. Furthermore, a nested block structure is apparent in the matrix, reflecting the phylogenetic tree (recall that chains A and B are both ordered according to the tree as shown in [Fig 2A](#)). Specifically, for 22% of chains A in the testing set, the smallest DCA effective interaction energy E_{AB} is obtained with their evolutionary




Fig 2. Pairs of chains with common evolutionary history have small DCA effective interaction energies. **A:** Chains are numbered according to the phylogenetic tree representing the branching process used for data generation (see Fig 1). The same numbering is employed for chains A and for chains B that possess the same evolutionary history. **B:** Matrix of DCA effective interaction energies E_{AB} (Eq 1) for all pairs AB made from a chain A and a chain B of the testing set, numbered according to phylogeny as illustrated in panel A. Data was generated using a tree of 10 generations, with exactly 5 mutations per branch, out of 200 bits in each chain AB, thus yielding 1024 chains AB. Next, 75% of them were randomly selected to form the training set employed to build the DCA model, while the remaining 25% constitute the testing set.

<https://doi.org/10.1371/journal.pcbi.1007179.g002>

partner (corresponding to the diagonal in Fig 2B). In S2 Fig, we further demonstrate that those chains B that have smaller E_{AB} with a chain A than its evolutionary partner B are very similar to that chain B and strongly related to it. Furthermore, if the dataset is divided into random species with 4 chains AB each (see Methods), the smallest E_{AB} for a chain A within its species accurately identifies its evolutionary partner B for 93% of chains A of the testing set. Hence, with a large training set, DCA is able to learn phylogenetic correlations, and to identify evolutionary partners. Recall that the usual goal of DCA is to infer couplings $e_{ij}(\sigma_i, \sigma_j)$ stemming from actual interactions, which do not exist in our synthetic data.

Impact of key parameters and comparison to other methods

Let us investigate the robustness of the ability of DCA to identify pairs of evolutionary partners, and compare it to other methods. First, we ask how large a training set is necessary to learn the correlations arising from phylogeny. Fig 3A shows that a sufficiently large training set is required for DCA to accurately identify evolutionary partners within each species, in line with previous results about DCA-based predictions of protein-protein interactions [22, 23] and three-dimensional protein structures [5–7] from real sequences. Furthermore, similar trends are observed both when employing the Mutual Information (MI) based score S_{AB} (Eq 3), consistently with [41], and when using the Mirrortree-inspired score M_{AB} (Eq 4) that only relies on sequence similarity. All these methods predict pairings much better than the chance expectation (yellow) and reach very high fractions of true positives for training sets larger than ~ 100 pairs AB (see Fig 3A). Better performance is obtained when pairings are predicted using the Hungarian algorithm, which finds a global optimal one-to-one matching within each species (see Methods) than when simply picking for each chain A the optimal partner B within its species [41]. Fig 3A further shows that with the first approach, the Mirrortree method performs better for small training sets, while DCA and MI outperform it for larger datasets. These differences become almost negligible when using the Hungarian algorithm.




Fig 3. Performance of pairing prediction versus training set size and number of pairs per species. **A:** Fraction of pairs correctly identified (TP fraction) versus training set size, for DCA-, MI-, and Mirrortree-based methods. The three pairing scores corresponding to each of these three methods are employed in two ways: either within each species we find the chain B with optimal pairing score with each chain A (dashed lines), or within each species we employ the Hungarian matching algorithm to find the one-to-one pairing of chains A and B that optimizes the sum of the pairing scores (solid lines). Each species comprises 4 chains AB. **B:** Fraction of pairs correctly identified (TP fraction) versus number of pairs per species, employing the same methods (and same colors) as in panel A, and a training set of 50% of the total dataset. In both panels, data was generated using a tree of 10 generations, with exactly 5 mutations per branch, out of 200 bits in each chain AB, thus yielding 1024 chains AB. Species were built randomly, and some of them were chosen randomly to build the training set, the remaining ones making up the testing set. Yellow curves show the chance expectation, i.e. the average TP fraction obtained for random within-species pairings. Results are averaged over 100 replicates in panel A and 20 replicates in panel B, each corresponding to a different realization of the branching process used for data generation. The standard deviation of the TP fraction is 2-3% for large training sets with 4 chains per species (panel A).

<https://doi.org/10.1371/journal.pcbi.1007179.g003>

Since the pairing task becomes harder when the number of pairs per species increases, we next studied how performance is affected by this important parameter. Fig 3B, which employs a substantial training set, shows that the performance of all three pairing scores decays as species contain more pairs AB, as expected. However, this decay is far slower than for the chance expectation (yellow), which highlights the robustness of our methods. Here, DCA reaches the highest performance, followed by MI and then by Mirrortree, in line with the results obtained on Fig 3A for large training sets. The good performance of the Mirrortree approach, which just relies on sequence similarities, arises from the fact that a possible pair AB that is very similar to correct pairs tends to be a correct pair too, as evidenced in S3 Fig. Indeed, pairs that are very similar to correct pairs tend to be their close “relatives” along the tree. Other variants based on sequence similarity can thus be constructed. In S4 Fig, we present two such variants: one employs as a pairing score the smallest Hamming distance from a possible pair AB of the testing set to its closest neighbor in the training set, and the second one is based on the notion of orthologous pairs. Both of them perform very well with large training sets, but are less robust than our other methods to decreasing training set size.

Because the ability of our methods to predict pairings relies on the shared evolutionary history of chains A and B, it is crucial to understand how the mutation rate affects performance. Let $\mu = n\mu_0$ denote the average total number of mutations between the ancestral complete chain AB and any complete chain AB at a leaf of the phylogenetic tree, where n represents the number of generations and μ_0 the average number of mutations per generation (see Fig 1). If the maximum number of differences between two complete final chains AB, namely 2μ , becomes larger than the total length $2L$ of a complete chain AB, then correlations are lost between these two least-related chains AB. Thus, we expect the performance of our pairing prediction methods to decay for $\mu \gtrsim L$. This constitutes a lower bound of the actual number μ




Fig 4. Performance of DCA- and Mirrortree-based predictions for various data parameters. The fraction of pairs AB correctly predicted (TP fraction) is shown versus the average total number of mutations μ per chain AB and the length L of a single chain A or B for DCA (panel A) and Mirrortree (panel B). The Hungarian algorithm was employed to predict pairings. For each μ and L , data was generated using a tree of $n = 10$ generations, thus yielding 1024 chains AB, and random species, each comprising 4 chains AB, were constructed. Half of the species were chosen to form a training set of 512 pairs, and predictions were made on the remaining species, which form the testing set. Here the chance expectation of TP fraction, obtained for random within-species pairings, is 0.25.

<https://doi.org/10.1371/journal.pcbi.1007179.g004>

of mutations causing performance to substantially drop, because (i) we have considered the two least-related chains along the trees, and chains that diverged upon later duplication steps are more correlated, and (ii) since each mutation affects a random site, and each site can mutate several times, some sites may never mutate even when $\mu \geq L$, and thus some correlations can survive in this regime even between the most distant sequences. Similarly, if $\mu_0 \geq L$, i.e. $\mu \geq nL$, then even sister complete final chains AB lose correlation, which gives an upper bound for the number of mutations causing performance to drop.

Fig 4 shows heatmaps of the performance of DCA- and Mirrortree-based pairing predictions versus the total number μ of mutations per chain AB and the single chain length L , with a substantial training set. For both methods, performance is very good and robust over a large range of values of L and μ . In addition, in both cases, a clear transition between good and poor performance is visible as μ is increased at each L . We observe that this transition occurs along a line, such that good performance is obtained for $\mu \lesssim 3.6L - 72$. This linear behavior and its slope are consistent with our predictions above. We also observe that performance drops if there are extremely few mutations, because chains are too conserved and remain almost all the same, and if chains are too short, because there is too much redundancy. Another important parameter is the total number n of generations in the phylogenetic tree, which sets the total number of chains (2^n). We found that varying n from 8 to 12 yielded no significant change the heatmaps of Fig 4. Finally, the DCA- and Mirrortree-based methods perform very similarly over the whole range of parameters studied in Fig 4: specifically, the mean difference of the TP fractions obtained using the two methods is 4×10^{-3} . Despite their conceptual differences, these methods rely on learning phylogenetic correlations from the training set, and thus have similar dependences on evolutionary parameters such as the mutation rate.

Evolutionary partners can be predicted without a training set

In Refs. [22, 41], it was shown that DCA- and MI-based approaches allow to predict interacting partners among the paralogs of actual interacting proteins from their sequences without any training set, i.e. without any prior knowledge of interacting pairs, thanks to an Iterative Pairing Algorithm (IPA). In this approach, at the first iteration, (usually poor) predictions are

made employing pairing scores learned on random within-species pairings. At each subsequent iteration $n > 1$, the predictions from the previous iteration that are deemed most reliable [22, 41] are progressively incorporated. Specifically, pairing scores are re-learned on the $(n - 1)N_{\text{increment}}$ top-ranked predicted pairs from the previous iteration, where $N_{\text{increment}}$ represents the increment step: hence, the number of predicted pairs that are employed to make the next predictions increases by $N_{\text{increment}}$ at each iteration (see Refs. [22, 41] for details, and Ref. [23] for alternative iterative approaches). This iterative strategy gradually improves predictive power and has yielded accurate predictions of interacting partners in ubiquitous prokaryotic protein families [22, 41]. Here, we employed the IPA on our synthetic data where correlations arise only from shared evolutionary history. For comparison, we also developed and studied a variant of the IPA based on the Mirrortree approach, which employs the pairing score M_{AB} (Eq 4), instead of the DCA-based effective interaction energy E_{AB} (Eq 1) [22] or MI score S_{AB} (Eq 3) [41].

Fig 5A shows the TP fraction obtained for different values of the increment step $N_{\text{increment}}$, both for the DCA-IPA and for the Mirrortree-IPA, at the first and last iterations. Overall, it shows that the iterative approach allows to make very accurate pairing predictions in the absence of a training set. With both algorithms, a strong improvement of predictive power is observed at the last iteration, compared to the first iteration and to the random expectation. Furthermore, the iterative method performs best for small increment steps $N_{\text{increment}}$, which highlights the interest of the iterative approach. We emphasize that the high final TP fractions are attained without any prior knowledge of pairings. As discussed in Refs. [22, 41], an important ingredient for the IPA to bootstrap its way toward high predictive power is that among pairs AB comprising a chain A and a chain B from the same species, correct pairs of partners possess more neighbors in terms of sequence similarity, quantified by the Hamming distance,




Fig 5. Pairing prediction without any training set. **A:** The fraction of pairs AB correctly predicted (TP fraction) is shown versus the increment step $N_{\text{increment}}$ of the iterative process, for the DCA-IPA and the Mirrortree-IPA, at the first and last iterations. Data was generated using a tree of 10 generations, with 20 mutations per branch on average, out of 200 bits in each chain AB, thus yielding 1024 chains AB, and random species with 4 pairs AB each were constructed. **B:** TP fraction versus number of pairs per species, for the DCA-IPA and the Mirrortree-IPA, at the first and last iterations, as well as for the Switch-IPA, which uses the Mirrortree pairing score for the first half of iterations and then switches to the DCA pairing score (the first iteration is thus the same for the Switch-IPA as for the Mirrortree-IPA). An increment step $N_{\text{increment}} = 100$ was used. Data was generated using a tree of 10 generations, with 5 mutations per branch on average, out of 200 bits in each chain AB, thus yielding 1024 chains AB, and random species with the same number of chains AB each were constructed. Colors are the same as in panel A, with the addition of the Switch-IPA. In both panels, predictions were made without any training set, and the first iteration employed random within-species pairings to compute the initial pairing scores. The Hungarian algorithm was employed to predict pairings. Results are averaged over 20 replicates in panel A and 100 in panel B, each corresponding to a different realization of the branching process used for data generation. Error bars represent 95% confidence intervals (in many cases, markers are larger than error bars).

<https://doi.org/10.1371/journal.pcbi.1007179.g005>

than incorrect pairs. Ref. [22] called this the *Anna Karenina effect*, referring to the first sentence of Tolstoy's novel. We studied the Anna Karenina effect in our synthetic dataset. S3 Fig shows that correct pairs have closer correct neighbors than incorrect pairs: for instance, employing a threshold Hamming distance of 0.15 to define neighbors (see S3A Fig), correct pairs have 6.2 neighbors on average, of which 90% are correct pairs, while incorrect pairs have 0.63 neighbors on average, of which 33% are correct pairs (see S3B Fig). In this case, correct pairs thus have almost 10 times more neighbors than incorrect ones, demonstrating a strong Anna Karenina effect. This favors correct pairs in the IPA, especially at early iterations [22].

For the parameters used in Fig 5A, the Mirrortree-IPA performs slightly better than the DCA-IPA, and this difference exists right from the first iteration. How does the performance of these two methods depend on parameters characterizing the dataset? S5 Fig shows heatmaps of the performance of the DCA-IPA and of the Mirrortree-IPA as a function of the total number μ of mutations per chain AB and of the single chain length L , without any training set, at the first and last iterations. It shows that the Mirrortree-IPA performs better than the DCA-IPA at the first iteration, especially as mutation rates become larger (S5A and S5C Fig). Recall that at the first iteration, pairing scores are calculated on random within-species pairings, where most pairs (75% on average for species with 4 pairs each) are incorrect. Taken together with our earlier observation that Mirrortree outperforms DCA for small training sets (see Fig 3A), it means that DCA requires a substantial and accurate training set to properly learn correlations and reach good performance, as is already known in the case of real protein sequences [5–7]. Nevertheless, the IPA allows DCA to robustly reach high predictive power at the last iteration over a broad range of values of μ and L (S5B Fig). This range is almost as large as for the Mirrortree-IPA (S5D Fig), despite the initial lower performance of DCA. At the last iteration, there is only a rather narrow band, close to the transition line from large to small TP fractions, where the DCA-IPA is outperformed by the Mirrortree-IPA. Note that the parameters employed in Fig 5A are in this region. Importantly, the parameter range where the DCA-IPA and the Mirrortree-IPA perform well without any training set is very similar to that obtained in Fig 4 for DCA and Mirrortree predictions from a large training set. This result illustrates the power of the iterative approach, which truly allows to bypass the need for a training set.

In Fig 5B, we further investigate the impact of the number of pairs AB per species on performance of the DCA-IPA and of the Mirrortree-IPA without any training set. Very good performance is obtained for species comprising up to 8 pairs AB, and then we observe a decay, which is steeper than with a large training set (Fig 4). We further observe that the DCA-IPA reaches higher performance than the Mirrortree-IPA for small numbers of pairs per species, while the opposite is true for larger numbers of pairs per species. Again, the Mirrortree-IPA performs better than the DCA-IPA at the first iteration, and this might explain the decreased final performance of the DCA-IPA for large numbers of pairs per species: confronted with many pairing possibilities, the DCA-IPA may not be able to recover from the large amount of noise in the initial matches.

An interesting question is whether one can improve predictive power by combining DCA and Mirrortree approaches. It is all the more attractive that the predictions made using the two scores are almost independent (S6A Fig), in contrast to those made using DCA and MI [41] (S6B Fig). Because Mirrortree performs better than DCA at the first iteration while DCA becomes better for larger and more accurate training sets, we devised a version of the IPA that uses the Mirrortree pairing score for the first half of iterations and then switches to the DCA pairing score. The final TP fractions obtained with this Switch-IPA are shown in Fig 5B. We find that it performs as well as the best among the Mirrortree-IPA and the DCA-IPA, which should make it more broadly applicable. However, it does not yield further improvements.

Extension to other phylogenies and application to real proteins

So far, we have considered a minimal model for species, where chains are randomly grouped into sets of equal size m , each representing a species. Such a model would be realistic in the case where exchange between species (i.e. horizontal gene transfer) is very frequent. But in other evolutionary regimes, different correlations are expected between the chains present in a given species (paralogs) and between the most closely related chains across species (orthologs), due to the fact that species are evolutionary units. In practice, paralogous and orthologous pairs respectively arise from duplication and speciation events. In a duplication event, a chain from a given species gives rise to 2 paralogous chains within this species. Note that loss events can also occur, thus decreasing the number of chains within a species. In a speciation event, all chains are duplicated to give rise to 2 distinct species. In order to assess the robustness of our results to these effects, we now consider a phylogeny model that explicitly accounts for duplication-loss and speciation events, without any exchange within species (see [S7A Fig](#)). For simplicity, we assume that duplication and loss always happen together, so that the number of pairs per species remains constant.

[S7B Fig](#) shows the performance of DCA and Mirrortree scores at predicting pairs of evolutionary partners from a substantial training set, versus the fraction of species that undergo a duplication-loss event upon speciation. Overall, performance is good, but it decreases when the frequency of duplication-loss events increases. When there are no duplication-losses, this model features m distinct phylogenies ($m = 4$ in [S7 Fig](#)), one for each ancestral chain, and similarities between the chains of the testing set and of the training set that belong to the same phylogeny allow to predict evolutionary partners. Conversely, when there are many duplication-losses, chains from one single phylogeny will end up fixing in each species, analogously to asexual birth-death population genetics models at fixed population size where all individuals are descended from a single ancestor after a sufficient time [60]. Moreover, chains resulting from recent duplication events will be very hard to distinguish, resulting in pairing ambiguities. [S7B Fig](#) also shows that DCA outperforms Mirrortree, consistently with other cases with a substantial training set. Finally, in [S7C Fig](#), we consider the case without a training set. We find that the first iteration of the Mirrortree-IPA performs much better than that of the DCA-IPA, while things are more even after the iterative process, consistently with our previous results. Furthermore, [S7C Fig](#) shows that in the absence of duplication-loss, final performance is close to the random expectation of 0.25, which stands in contrast with the case with a training set. This is because in the absence of a training set, incorrect pairs comprising A and B chains from different phylogenies cannot be distinguished from those coming from the same phylogeny. Duplication-loss events allow to break this symmetry thanks to the fixation process whereby possible cross-phylogeny pairs become rare within each species. However, when duplication-loss events are too frequent, performance decays again, for the same reason as with a training set: pairs resulting from recent duplications are hard to distinguish. This trade-off yields an optimal performance for an intermediate frequency of duplication-loss events. Overall, our DCA and Mirrortree-based pairing predictions are robust to modifications of the phylogeny used to generate the data, as long as evolutionary correlations exist.

Since the various methods presented here allow to reliably pair synthetic chains on the basis of shared evolutionary history only, it should also be the case for real proteins. While it is difficult to be certain that two real protein families share evolutionary history but do not bear common functional evolutionary constraints, we chose two pairs of protein families that are generally encoded in close proximity on prokaryotic genomes but that do not have known direct physical interactions [61], namely the *Escherichia coli* protein pairs LOLC-MACA and ACRE-ENVR and their homologs. (Note that ENVR has regulatory roles on ACRE expression




Fig 6. Pairing predictions for real pairs of protein families. The final fraction of protein pairs correctly predicted (TP fraction) obtained without a training set by the MI-IPA, the DCA-IPA and the Mirrortree-IPA is shown versus the increment step $N_{\text{increment}}$ of the iterative process. **A, B:** Pairs of protein families with no known direct physical interaction but that are encoded closely on the genome; datasets include ~ 2000 homologous pairs. **C, D, E:** Pairs of protein families with known direct physical interactions; for large families (C and D), datasets of ~ 5000 pairs comprising only full species were extracted from the larger complete datasets, while in E, the full dataset of ~ 2000 pairs was used. In all panels, the mean number $\langle m \rangle$ of pairs per species is indicated, and yellow lines represent the average TP fraction obtained for random within-species pairings. All results are averaged over 50 replicates that differ in their initial random within-species pairings. Datasets were constructed as described in [22, 41], starting from the P2CS database [67, 68] for histidine kinase-response regulator (HK-RR) dataset (C) and using a method adapted from [17] that relies on finding homologs of *Escherichia coli* protein pairs in all other cases.

<https://doi.org/10.1371/journal.pcbi.1007179.g006>

[62]). Indeed, we expect proteins encoded close to one another to share common evolutionary history, because they tend to be horizontally transferred together, and to have similar levels of expression and evolutionary rates [36, 40]. Fig 6A and 6B shows that the DCA-, MI- and Mirrortree-IPA are all able to reliably pair LOLC-MACA and ACRE-ENVR homologs that are encoded close to one another on genomes, despite the absence of known direct physical interactions between these protein families. We further compared these results to three pairs of protein families with known direct physical interactions [41]. Fig 6C, 6D and 6E shows that the Mirrortree-IPA performs less well than the DCA and MI-based versions in these cases, especially for the dataset of histidine kinases and response regulators (Fig 6C). These proteins feature a strongly coevolving interaction interface [63], diversified across many paralogs per species (average number of pairs per species in our dataset: $\langle m \rangle = 11$) to avoid unwanted cross-talk. This argues in the favor of the DCA-IPA and the MI-IPA rather than Mirrortree-based methods in order to predict partnership among physically interacting partners. Interestingly, the MI-IPA slightly outperforms the DCA-IPA for all the real datasets in Fig 6, while this is not the case for synthetic data comprising only phylogenetic correlations (see Fig 3).

Distinguishing physically interacting proteins from proteins only sharing evolutionary history

We have shown that the IPA, which aims to find partners among the paralogs of two protein families, accurately identifies pairs of proteins even if they only share a common evolutionary history, in the absence of direct interactions. In order to computationally predict protein-protein interactions from sequence data, it is in addition necessary to distinguish the pairs of protein families that physically interact from those that only share evolutionary history.

First, interacting protein families often share more evolutionary history than non-interacting ones. This idea is at the heart of the Mirrortree approach [31, 34, 35], which is usually implemented on alignments of orthologs, thus allowing to predict which protein families interact, but not addressing the paralog pairing problem. We have applied the original Mirrortree [31, 34] and pMirrortree [35] methods to the pairs of protein families studied in Fig 6, focusing on the orthologs of the reference *E. coli* protein pairs (see S1 Table). The two pairs without known physical interactions (LOLC-MACA and ACRE-ENVR) feature the smallest Mirrortree scores among the five pairs considered, and they also have non-significant pMirrortree scores, while two out of the three pairs with known direct physical interactions (HK-RR and XDHA-XDHC) possess significant pMirrortree scores. These rather encouraging results confirm the power of the Mirrortree method, and support the hypothesis that our two pairs without known physical interactions are really non-interacting.

Another way to distinguish physically interacting pairs of protein families from non-interacting ones is to leverage the DCA scores of amino-acid pairs. Indeed, strong DCA scores tend to correspond to contacting amino-acid pairs [5–7], and thus, their presence can reveal actual interacting partners [17, 21, 22]. Fig 7 demonstrates that outliers in DCA scores [22] exist for the three pairs with known direct physical interactions studied in Fig 6 (HK-RR, MALG-MALK and XDHA-XDHC), while they are absent for the two pairs without known physical interactions (LOLC-MACA and ACRE-ENVR). Therefore, in these examples, outliers in DCA scores successfully allow to distinguish pairs of protein families that physically interact from those that only share a similar evolutionary history. Importantly, the results on Fig 7 were obtained on the pairs of protein sequences predicted by the IPA, which means that this method can be directly combined with the IPA. Note that the absence of outliers in DCA scores in the pairs without known physical interactions hints that phylogenetic correlations result in multiple small DCA couplings, while direct physical interactions yield few large DCA couplings. We also checked that our phylogeny-only synthetic data (with the same parameters as in Fig 2) does not yield outliers in DCA scores.

Therefore, combining either traditional Mirrortree approaches or the study of outliers in DCA scores with the IPA can allow both to predict interacting pairs of protein families and to solve the paralog pairing problem, starting from two independent multiple sequence alignments of the protein families.

Discussion

Recently, methods relying on pairwise maximum entropy DCA models, originally employed to identify amino acids in contact in the three-dimensional structure of proteins and multi-protein complexes, have allowed to reliably predict interacting partners among the paralogs of several ubiquitous prokaryotic protein families, starting from sequences only [22, 23]. An important motivation for these methods is that the need to maintain the physico-chemical complementarity of contacting amino acids induces correlations in amino-acid usage between the sequences of interacting partners [5, 63]. However, correlations between the sequences of interacting partners can also arise from their shared evolutionary history [39].




Fig 7. Outliers in DCA score provide evidence of protein-protein interactions. DCA scores (APC-corrected Frobenius norms [17, 47]) were evaluated for each pair of amino-acid sites at the final IPA iteration. Their averages over 500 IPA replicates that differ in their initial random pairings of sequences are shown ranked by decreasing value. Datasets correspond to the same protein family pairs as in Fig 6, and the IPA was run with $N_{\text{increment}} = 50$. Outliers in DCA score appear for all three pairs of protein families with known direct physical interactions (BASS-BASR, MALG-MALK, XDHA-XDHC), but not for the other ones (LOLC-MACA, ACRE-ENVR).

<https://doi.org/10.1371/journal.pcbi.1007179.g007>

In the present work, employing controlled synthetic data, we demonstrated that DCA is able to accurately identify partners that only share evolutionary history, in the absence of functional and structural constraints. This result holds even in the absence of a training set, thanks to the Iterative Pairing Algorithm (IPA) [22]. Because our controlled synthetic data only comprises signal from phylogeny, we compared our DCA-based approach to methods that explicitly rely on phylogeny, through sequence similarity. Specifically, we proposed a method based on the Mirrortree approach [31, 34, 35] to predict pairs among the paralogs of two protein families. We obtained similar performances, with DCA slightly outperforming Mirrortree when a substantial correct training set is available. We also considered a method based on Mutual Information (MI) [41], yielding similar performances for our synthetic data, with DCA often slightly outperforming MI, while MI tends to outperform DCA for natural sequence data [41]. The robustness of MI to finite dataset size effects [41] and its ability to quantify statistical dependence whatever its origin might make it most appropriate for complex natural data.

Finally, we applied the DCA-IPA, the MI-IPA and the Mirrortree-IPA to natural sequence data from several pairs of prokaryotic protein families, with or without known direct physical interactions, but always with genome proximity and thus significant shared evolutionary history. We obtained accurate predictions for all these datasets. Therefore, correlations from evolutionary history can play an important part in the performance of these algorithms in the case of natural sequence data. Interestingly, we found that the Mirrortree-IPA performs significantly less well than the MI-IPA and DCA-IPA on the histidine kinase-response regulator dataset, which features large numbers of paralogs per species and is known to possess strongly coevolving contacts [5, 63], and to which DCA is thus particularly well-suited. This points to the complementarity of these methods. In addition, we showed that pairs of protein families

with known direct physical interactions can be distinguished from those without known direct physical interactions, either by employing the original Mirrortree approach involving orthologs only [31, 34, 35], or by studying the amino-acid pairs that are outliers in DCA score [22]. Hence, combining these methods with the IPA allows to both predict interacting protein families and to find interacting partners among paralogs.

The ability of DCA to identify evolutionary partners in the absence of functional and structural constraints can be viewed as surprising. Indeed, DCA models are mainly known for their ability to identify structural contacts, and they emphasize small-eigenvalue modes of the covariance matrix of sequences [64, 65], as illustrated by the inversion of the covariance matrix involved in the mean-field approximation (see [Methods](#)), while important signal from phylogeny lies in the large-eigenvalue modes of the covariance matrix [26–28]. In addition, phylogenetic correlations are often considered deleterious to structure prediction [2, 5, 6, 28], which is one of the major applications of DCA. Nevertheless, a DCA model is fundamentally a global statistical model that aims to faithfully reproduce the empirical pairwise correlations observed in the data [44, 45]. As such, it should also encode phylogenetic correlations, thus rationalizing our result. Our findings demonstrate that DCA models capture coevolution in a broad sense, not limited to contacting pairs of residues, but also including shared evolutionary history unrelated to functional constraints. Importantly, while phylogenetic correlations are often viewed as a confounding factor for structure prediction [2, 5, 6, 28], our work shows that they are actually useful in order to address the paralog pairing problem by DCA. This is consistent with previous work directly exploiting phylogenetic correlations to predict protein-protein interactions [29, 30, 32, 34, 35]. Interesting further directions include analyzing the contributions of phylogeny topology and evolutionary rates [36–40] in the phylogenetic signal that is useful for DCA-based pairing predictions, analyzing the impact of variations of the evolutionary rate, and addressing how signals from phylogeny and from amino-acid contacts combine together.

Here, our natural data examples were pairs of proteins colocalized on genomes, allowing us to easily evaluate performance via the fraction of correctly predicted protein pairs. However, computational interaction prediction and paralog pairing methods are particularly important in order to discover interaction partners that are not encoded in close genomic locations. This corresponds to the general case in eukaryotes, which makes this problem extremely relevant. In prokaryotes too, important interactions between very different cellular processes exist across operons [66]. It will be of great interest to apply the IPA to such cases. One restrictive assumption we have made so far is that interactions are one-to-one, which is appropriate for strongly specific protein-protein interactions, but not for cases involving promiscuity, cross-talk, or non-interacting orphan proteins. Hence, an important direction for future work is to generalize the IPA to allow for multiple partners.

Supporting information

S1 Fig. Impact of the pseudocount weight on performance. The fraction of pairs AB correctly predicted (TP fraction) is shown versus the pseudocount weight Λ both for DCA- and for MI-based pairing methods. The Hungarian algorithm was employed to predict pairings. **A:** Full range of variation of Λ (excluding $\Lambda = 1$, where no more signal comes from the actual dataset, and accordingly, predictions follow the chance expectation of 0.25 TP fraction). **B:** Zoom over small values of Λ . In both panels, data was generated using a tree of 10 generations, with exactly 5 mutations per branch, out of 200 bits in each chain AB, thus yielding a total of 1024 chains AB. Random species with 4 chains AB each were constructed; 25 of them were randomly selected to form a training set of 100 pairs AB employed to build the DCA model,

and the rest constitutes the testing set, for which we calculated DCA pairing scores (Eq 1). Results are averaged over 100 realizations corresponding to different random partitions of the dataset into training and testing sets. Error bars represent 95% confidence intervals.

(EPS)

S2 Fig. Pairs with smaller DCA effective interaction energy than evolutionary partners are similar and closely related to them.

A: Blue: Histogram of Hamming distances between all chains B from the testing set. Red: Histogram of Hamming distances between each chain B from the testing set and those with smaller DCA effective interaction energy E_{AB} (Eq 1) than B itself with the evolutionary partner A of B. **B:** Blue: Histogram of the last common ancestor index between all chains B from the testing set. Red: Histogram of the last common ancestor index between each chain B from the testing set and those with smaller DCA effective interaction energy E_{AB} (Eq 1) than B itself with the evolutionary partner A of B. The last common ancestor index is defined according to the phylogenetic tree representing the branching process used for data generation (see Fig 1): it is 1 for chains that are evolutionarily closest (“sisters”), 2 for the next level (“cousins”), etc. In both panels, data was generated using a tree of 10 generations, with exactly 5 mutations per branch, out of 200 bits in each chain AB, thus yielding 1024 chains AB. Next, 75% of them were randomly selected to form the training set employed to build the DCA model and compute E_{AB} values, while the remaining 25% constitute the testing set.

(EPS)

S3 Fig. Neighbors of correct pairs tend to be correct pairs.

A: Schematic illustrating the notion of neighbor pairs. On the left, each chain A is represented as a blue dot, and distances between dots correspond to Hamming distances, i.e. fractions of sites with different states. A threshold Hamming distance t is chosen to define the notion of neighbors. The dots inside the circle with center A_1 and radius t represent the neighbors of A_1 . A similar schematic is shown on the right for the B chains. Two pairs A_1B_1 and A_2B_2 are considered neighbors if A_1 and A_2 are neighbors, and B_1 and B_2 are neighbors too. Thus, in our schematic, A_2B_2 is a neighbor of A_1B_1 , but A_3B_3 is not. In practice, only pairs where A and B belong to the same species are considered, because they are the only possible correct pairs of evolutionary partners. **B:** The number of within-species pairs AB that possess neighbors that are correct pairs of evolutionary partners is shown versus the Hamming distance threshold t defining the notion of neighbor (see panel A). The correct pairs, i.e. those are evolutionary partners, and the incorrect within-species pairs, made of a chain A and a chain B that are not evolutionary partners, are distinguished. For intermediate distance thresholds, there are far more correct pairs than incorrect pairs that possess correct neighbors. For instance, for a Hamming distance threshold of 0.1, 22 incorrect pairs have correct neighbors, while all 1024 correct pairs have correct neighbors.

Data was generated using a tree of 10 generations, with exactly 5 mutations per branch, out of 200 bits in each chain AB, thus yielding 1024 chains AB, and random species with 4 chains AB each were constructed.

(EPS)

S4 Fig. Performance of pairing prediction versus training set size and number of pairs per species for various methods.

A: Fraction of pairs correctly identified (TP fraction) versus training set size, for DCA- and MI-based methods, and for methods directly based on sequence similarity. “Distance” employs as a pairing score the Hamming distance between a possible pair AB of the testing set and its closest neighbor in the training set. In “Orthology”, one finds reciprocal closest neighbors in terms of Hamming distance (“orthologs”) for each possible within-species pair AB of the testing set, among the correct and incorrect within-

species pairs of the training set. Specifically, a pair P (from the testing set) and a pair P' (from the training set) are considered as reciprocal closest neighbors if (i) P is the the closest neighbor of P' among the pairs of the species of P and (ii) P' is the closest neighbor of P among all the pairs from the training set. Next, one ranks possible pairs AB of the testing set using this notion of orthology: the pairs whose orthologs are all correct pairs from the training set come first, ranked by decreasing number of orthologs. Then, the pairs of the testing set that have both correct and incorrect orthologs are ranked by decreasing fraction of correct pairs among their orthologs. Finally, the pairs whose orthologs are all incorrect are ranked by increasing number of number of orthologs. In case of equality, pairs are ranked by decreasing distance to the closest correct pair from the training set. We employ the rank of a pair AB in this scheme as a pairing score. The four pairing scores corresponding to each of the four methods are employed in two ways: either within each species we find the B chain with optimal pairing score with each A chain (dashed lines), or within each species we employ the Hungarian matching algorithm to find the one-to-one pairing of A and B chains that optimizes the sum of the pairing scores (“Hung.”, solid lines). Each species comprises 4 chains AB . **B:** Fraction of pairs correctly identified (TP fraction) versus number of pairs per species, employing the same methods (and same colors) as in panel A, and a training set of 50% of the total dataset. For the Orthology method, results are shown only up to 8 pairs per species, because computations become lengthy. In both panels, and as in Fig 3, data was generated using a tree of 10 generations, with exactly 5 mutations per branch, out of 200 bits in each chain AB , thus yielding 1024 chains AB . Species were built randomly, and some of them were chosen randomly to build the training set, the remaining ones making up the testing set. The chance expectation, corresponding to random within-species pairings, is shown for comparison. Results are averaged over 100 replicates in panel A and 20 replicates in panel B, each corresponding to a different realization of the branching process used for data generation.

(EPS)

S5 Fig. Performance of the DCA-IPA and of the Mirrortree-IPA. The fraction of pairs AB correctly predicted (TP fraction) is shown versus the total number μ of mutations per chain AB and the length L of a single chain A or B for the DCA-IPA at the first (A) and last iteration (B), and similarly for the Mirrortree-IPA (C and D). Predictions were made without any training set, and the first iteration employs random within-species pairings to compute the initial pairing scores. An increment step $N_{\text{increment}} = 100$ was used. The Hungarian algorithm was employed to predict pairings. Data was generated using a tree with 10 generations, with a variable average number of mutations per branch and a variable chain length, thus yielding 1024 chains AB , and random species with 4 chains AB each were constructed.

(EPS)

S6 Fig. Correlation of the predictions from different methods. A: The fraction of pairs AB correctly predicted (TP fraction) is shown versus the number of pairs in the training set, for DCA- and Mirrortree-based pairing methods. The fraction of pairs correctly predicted by both methods is also shown, as well as its expectation if the two methods were fully independent. B: Similar plot, comparing DCA and MI instead of DCA and Mirrortree. In both panels, data was generated using a tree of 10 generations, with 5 mutations per branch on average, out of 200 bits in each chain AB , thus yielding 1024 chains AB . Random species with 4 chains AB each were constructed. Results are averaged over 100 replicates, each corresponding to a different realization of the branching process used for data generation.

(EPS)

S7 Fig. Extension to different phylogenies. A: Starting from one ancestral species with 4 random chains whose two halves A and B are shaded in blue and red, a series of speciation and diversification steps are performed (each sequence of these steps is called a “generation”). Upon speciation, each species is duplicated. Then, each of the two resulting species has a certain probability to undergo a duplication-loss step where one chain is eliminated and replaced by a copy of another chain from the same species. Next, mutations occur independently in each species (light green; here 2 bits per chain AB are mutated at each generation). Hence, the number of species doubles at each generation. **B:** Predictions with a training set. The fraction of pairs AB correctly predicted (TP fraction) is shown versus the fraction of species that undergo duplication-loss at each generation, for DCA- and Mirrortree-based pairing methods. Two versions are shown for each method: either within each species we find the B chain with optimal pairing score with each A chain (dashed lines), or within each species we employ the Hungarian matching algorithm to find the one-to-one pairing of A and B chains that optimizes the sum of the pairing scores (solid lines). Here, 25 species were randomly selected to form a training set of 100 pairs AB employed to build the pairing scores, and the rest constitutes the testing set. **C:** Prediction without a training set. The fraction of pairs AB correctly predicted (TP fraction) is shown versus the fraction of species that undergo duplication-loss at each generation, for the DCA-IPA and the Mirrortree-IPA, at the first and last iterations. An increment step $N_{\text{increment}} = 100$ is used. The Hungarian algorithm is employed to predict pairings. In panels B and C, data was generated using a tree of 8 generations, with 5 mutations per branch on average, out of 200 bits in each chain AB. Because here we start from one ancestral species with 4 random chains and not from 1 single random chain, this yields 1024 chains AB, already separated in species. Results are averaged over 20 replicates, each corresponding to a different realization of the branching process used for data generation. Error bars represent 95% confidence intervals. (EPS)

S1 Table. Mirrortree results for the protein pairs considered in Fig 6. For each pair of *E. coli* proteins, we applied the original Mirrortree and pMirrortree methods, as described in Refs. [31, 34] and [35], respectively. The resulting scores quantify the similarity of the phylogenetic trees constructed from the orthologs of the two proteins considered. Note that paralogs are not included in this analysis, in contrast to the one presented in our paper. The Mirrortree score is a Pearson correlation between the sets of distances of the two protein families considered, and a large value is indicative of coevolution [31, 34]. Among the five pairs considered, LOLC-MACA and ACRE-ENVR have the smallest Mirrortree scores, in agreement with the fact that they are the only pairs that possess no known direct physical interactions. The pMirrortree score is a p-value that provides an assessment of the confidence in the tree similarity scores: a small pMirrortree score is indicative of coevolution [35]. Considering the standard significance threshold 0.05, BASS-BASR (HK-RR) and XDHA-XDHC have significant coevolution, while LOLC-MACA and ACRE-ENVR do not. This is in good agreement with the fact that BASS-BASR and XDHA-XDHC possess known direct physical interactions, while LOLC-MACA and ACRE-ENVR do not. Surprisingly, MALG-MALK has a large pMirrortree score while possessing a known direct physical interaction, but this is mitigated by the fact that the close paralogs POTI-POTA have significant coevolution according to their pMirrortree score. (PDF)

Acknowledgments

AFB thanks Ned S. Wingreen and Yaakov Kleorin for inspiring discussions.

Author Contributions

Conceptualization: Martin Weigt, Anne-Florence Bitbol.

Funding acquisition: Martin Weigt, Anne-Florence Bitbol.

Investigation: Guillaume Marmier, Anne-Florence Bitbol.

Methodology: Guillaume Marmier, Martin Weigt, Anne-Florence Bitbol.

Supervision: Martin Weigt, Anne-Florence Bitbol.

Visualization: Guillaume Marmier, Anne-Florence Bitbol.

Writing – original draft: Anne-Florence Bitbol.

Writing – review & editing: Martin Weigt.

References

1. Rajagopala SV, Sikorski P, Kumar A, Mosca R, Vlasblom J, Arnold R, et al. The binary protein-protein interaction landscape of *Escherichia coli*. *Nat Biotechnol*. 2014; 32(3):285–290. <https://doi.org/10.1038/nbt.2831> PMID: 24561554
2. Lapedes AS, Giraud BG, Liu L, Stormo GD. Correlated mutations in models of protein sequences: phylogenetic and structural effects. In: *Statistics in molecular biology and genetics—IMS Lecture Notes—Monograph Series*. vol. 33; 1999. p. 236–256.
3. Burger L, van Nimwegen E. Accurate prediction of protein-protein interactions from sequence alignments using a Bayesian method. *Mol Syst Biol*. 2008; 4:165. <https://doi.org/10.1038/msb4100203> PMID: 18277381
4. Jaynes ET. Information Theory and Statistical Mechanics. *Phys Rev*. 1957; 106(4):620–630. <https://doi.org/10.1103/PhysRev.106.620>
5. 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. <https://doi.org/10.1073/pnas.0805923106> PMID: 19116270
6. 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. <https://doi.org/10.1371/journal.pone.0028766> PMID: 22163331
7. 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. <https://doi.org/10.1073/pnas.1111471108> PMID: 22106262
8. Sułkowska JI, Morcos F, Weigt M, Hwa T, Onuchic JN. Genomics-aided structure prediction. *Proc Natl Acad Sci USA*. 2012; 109(26):10340–10345. <https://doi.org/10.1073/pnas.1207864109> PMID: 22691493
9. 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. <https://doi.org/10.1534/genetics.113.155861> PMID: 23934888
10. 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. <https://doi.org/10.1073/pnas.1323734111> PMID: 24449878
11. 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. <https://doi.org/10.1093/molbev/msw188> PMID: 27604223
12. 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. <https://doi.org/10.1093/molbev/msv211> PMID: 26446903
13. Morcos F, Jana B, Hwa T, Onuchic JN. Coevolutionary signals across protein lineages help capture multiple protein conformations. *Proc Natl Acad Sci USA*. 2013; 110(51):20533–20538. <https://doi.org/10.1073/pnas.1315625110> PMID: 24297889
14. 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. <https://doi.org/10.1371/journal.pcbi.1004262> PMID: 26046683

15. 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. <https://doi.org/10.1371/journal.pone.0019729> PMID: 21573011
16. 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. <https://doi.org/10.1371/journal.pone.0092721> PMID: 24663061
17. Ovchinnikov S, Kamisetty H, Baker D. Robust and accurate prediction of residue-residue interactions across protein interfaces using evolutionary information. *Elife*. 2014; 3:e02030. <https://doi.org/10.7554/eLife.02030> PMID: 24842992
18. Hopf TA, Scharfe CP, Rodrigues JP, Green AG, Kohlbacher O, Sander C, et al. Sequence co-evolution gives 3D contacts and structures of protein complexes. *Elife*. 2014; 3:e03430. <https://doi.org/10.7554/eLife.03430>
19. Tamir S, Rotem-Bamberger S, Katz C, Morcos F, Hailey KL, Zuris JA, et al. Integrated strategy reveals the protein interface between cancer targets Bcl-2 and NAF-1. *Proc Natl Acad Sci USA*. 2014; 111(14):5177–5182. <https://doi.org/10.1073/pnas.1403770111> PMID: 24706857
20. dos Santos RN, Morcos F, Jana B, Andricopulo AD, Onuchic JN. Dimeric interactions and complex formation using direct coevolutionary couplings. *Sci Rep*. 2015; 5:13652. <https://doi.org/10.1038/srep13652> PMID: 26338201
21. Feinauer C, Szurmant H, Weigt M, Pagnani A. Inter-Protein Sequence Co-Evolution Predicts Known Physical Interactions in Bacterial Ribosomes and the Trp Operon. *PLoS ONE*. 2016; 11(2):e0149166. <https://doi.org/10.1371/journal.pone.0149166> PMID: 26882169
22. Bitbol AF, Dwyer RS, Colwell LJ, Wingreen NS. Inferring interaction partners from protein sequences. *Proc Natl Acad Sci USA*. 2016; 113(43):12180–12185. <https://doi.org/10.1073/pnas.1606762113> PMID: 27663738
23. 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. <https://doi.org/10.1073/pnas.1607570113> PMID: 27729520
24. Lezon TR, Banavar JR, Cieplak M, Maritan A, Fedoroff NV. Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns. *Proc Natl Acad Sci USA*. 2006; 103(50):19033–19038. <https://doi.org/10.1073/pnas.0609152103> PMID: 17138668
25. Jiang XL, Martinez-Ledesma E, Morcos F. Revealing protein networks and gene-drug connectivity in cancer from direct information. *Sci Rep*. 2017; 7(1):3739. <https://doi.org/10.1038/s41598-017-04001-3> PMID: 28623316
26. Casari G, Sander C, Valencia A. A method to predict functional residues in proteins. *Nat Struct Biol*. 1995; 2(2):171–178. <https://doi.org/10.1038/nsb0295-171> PMID: 7749921
27. Halabi N, Rivoire O, Leibler S, Ranganathan R. Protein sectors: evolutionary units of three-dimensional structure. *Cell*. 2009; 138(4):774–786. <https://doi.org/10.1016/j.cell.2009.07.038> PMID: 19703402
28. Qin C, Colwell LJ. Power law tails in phylogenetic systems. *Proc Natl Acad Sci USA*. 2018; 115(4):690–695. <https://doi.org/10.1073/pnas.1711913115> PMID: 29311320
29. Fryxell KJ. The coevolution of gene family trees. *Trends Genet*. 1996; 12(9):364–369. [https://doi.org/10.1016/s0168-9525\(96\)80020-5](https://doi.org/10.1016/s0168-9525(96)80020-5) PMID: 8855667
30. Goh CS, Bogan AA, Joachimiak M, Walther D, Cohen FE. Co-evolution of proteins with their interaction partners. *J Mol Biol*. 2000; 299(2):283–293. <https://doi.org/10.1006/jmbi.2000.3732> PMID: 10860738
31. Pazos F, Valencia A. Similarity of phylogenetic trees as indicator of protein–protein interaction. *Protein Eng Des Sel*. 2001; 14(9):609–614. <https://doi.org/10.1093/protein/14.9.609>
32. Jothi R, Kann MG, Przytycka TM. Predicting protein-protein interaction by searching evolutionary tree automorphism space. *Bioinformatics*. 2005; 21 Suppl 1:i241–250. <https://doi.org/10.1093/bioinformatics/bti1009> PMID: 15961463
33. Bradde S, Braunstein A, Mahmoudi H, Tria F, Weigt M, Zecchina R. Aligning graphs and finding substructures by a cavity approach. *EPL*. 2010; 89(3). <https://doi.org/10.1209/0295-5075/89/37009>
34. Ochoa D, Pazos F. Studying the co-evolution of protein families with the Mirrortree web server. *Bioinformatics*. 2010; 26(10):1370–1371. <http://csbg.cnb.csic.es/mtserver>. <https://doi.org/10.1093/bioinformatics/btq137> PMID: 20363731
35. Ochoa D, Juan D, Valencia A, Pazos F. Detection of significant protein coevolution. *Bioinformatics*. 2015; 31(13):2166–2173. <http://csbg.cnb.csic.es/pMT/>. <https://doi.org/10.1093/bioinformatics/btv102> PMID: 25717190
36. Hakes L, Lovell SC, Oliver SG, Robertson DL. Specificity in protein interactions and its relationship with sequence diversity and coevolution. *Proc Natl Acad Sci USA*. 2007; 104(19):7999–8004. <https://doi.org/10.1073/pnas.0609962104> PMID: 17468399

37. Juan D, Pazos F, Valencia A. Co-evolution and co-adaptation in protein networks. *FEBS Lett.* 2008; 582(8):1225–1230. <https://doi.org/10.1016/j.febslet.2008.02.017> PMID: 18282476
38. Kann MG, Shoemaker BA, Panchenko AR, Przytycka TM. Correlated evolution of interacting proteins: looking behind the mirrortree. *J Mol Biol.* 2009; 385(1):91–98. <https://doi.org/10.1016/j.jmb.2008.09.078> PMID: 18930732
39. Lovell SC, Robertson DL. An integrated view of molecular coevolution in protein-protein interactions. *Mol Biol Evol.* 2010; 27(11):2567–2575. <https://doi.org/10.1093/molbev/msq144> PMID: 20551042
40. Swapna LS, Srinivasan N, Robertson DL, Lovell SC. The origins of the evolutionary signal used to predict protein-protein interactions. *BMC Evol Biol.* 2012; 12:238. <https://doi.org/10.1186/1471-2148-12-238> PMID: 23217198
41. Bitbol AF. Inferring interaction partners from protein sequences using mutual information. *PLoS Comput Biol.* 2018; 14(11):e1006401. <https://doi.org/10.1371/journal.pcbi.1006401> PMID: 30422978
42. 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. <https://doi.org/10.1073/pnas.1105315108> PMID: 21690407
43. 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. <https://doi.org/10.1371/journal.pcbi.1003776> PMID: 25102049
44. Cocco S, Feinauer C, Figliuzzi M, Monasson R, Weigt M. Inverse statistical physics of protein sequences: a key issues review. *Rep Prog Phys.* 2018; 81(3):032601. <https://doi.org/10.1088/1361-6633/aa9965> PMID: 29120346
45. Nguyen HC, Zecchina R, Berg J. Inverse statistical problems: from the inverse Ising problem to data science. *Advances in Physics.* 2017; 66(3):197–261. <https://doi.org/10.1080/00018732.2017.1341604>
46. Pfleka 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. <https://doi.org/10.1088/0305-4470/15/6/035>
47. Ekeberg M, Lovkvist C, Lan Y, Weigt M, Aurell E. Improved contact prediction in proteins: using pseudo-likelihoods to infer Potts models. *Phys Rev E.* 2013; 87(1):012707. <https://doi.org/10.1103/PhysRevE.87.012707>
48. Fano R. *Transmission of Information: A Statistical Theory of Communications.* Cambridge, MA: The MIT Press; 1961.
49. Church KW, Hanks P. Word association norms, mutual information, and lexicography. *Comput Linguist.* 1990; 16(1):22–29.
50. Role F, Nadif M. Handling the impact of low frequency events on co-occurrence based measures of word similarity. In: Filipe J, Fred A, editors. *Proceedings of the International Conference on Knowledge Discovery and Information Retrieval (KDIR-2011).* Scitepress; 2011. p. 218–223.
51. Cover TM, Thomas JA. *Elements of Information Theory.* New York, NY, USA: Wiley-Interscience; 2006.
52. Ramani AK, Marcotte EM. Exploiting the co-evolution of interacting proteins to discover interaction specificity. *J Mol Biol.* 2003; 327(1):273–284. [https://doi.org/10.1016/s0022-2836\(03\)00114-1](https://doi.org/10.1016/s0022-2836(03)00114-1) PMID: 12614624
53. Gertz J, Elfond G, Shustrova A, Weisinger M, Pellegrini M, Cokus S, et al. Inferring protein interactions from phylogenetic distance matrices. *Bioinformatics.* 2003; 19(16):2039–2045. <https://doi.org/10.1093/bioinformatics/btg278> PMID: 14594708
54. Izarzugaza JM, Juan D, Pons C, Ranea JA, Valencia A, Pazos F. TSEMA: interactive prediction of protein pairings between interacting families. *Nucleic Acids Res.* 2006; 34(Web Server issue):W315–319. <https://doi.org/10.1093/nar/gkl112> PMID: 16845017
55. Izarzugaza JM, Juan D, Pons C, Pazos F, Valencia A. Enhancing the prediction of protein pairings between interacting families using orthology information. *BMC Bioinformatics.* 2008; 9:35. <https://doi.org/10.1186/1471-2105-9-35> PMID: 18215279
56. El-Kebir M, Marschall T, Wohlers I, Patterson M, Heringa J, Schonhuth A, et al. Mapping proteins in the presence of paralogs using units of coevolution. *BMC Bioinformatics.* 2013; 14 Suppl 15:S18. <https://doi.org/10.1186/1471-2105-14-S15-S18> PMID: 24564758
57. Kuhn HW. The Hungarian Method for the assignment problem. *Naval Research Logistics Quarterly.* 1955; 2:83–97. <https://doi.org/10.1002/nav.3800020109>
58. Munkres J. Algorithms for the Assignment and Transportation Problems. *Journal of the Society for Industrial and Applied Mathematics.* 1957; 5(1):32–38. <https://doi.org/10.1137/0105003>

59. Buehren M. Functions for the rectangular assignment problem; 2011. Available from: <https://fr.mathworks.com/matlabcentral/fileexchange/6543-functions-for-the-rectangular-assignment-problem>.
60. Ewens WJ. *Mathematical Population Genetics*. Springer-Verlag; 1979.
61. Szklarczyk D, Gable AL, Lyon D, Junge A, Wyder S, Huerta-Cepas J, et al. STRING v11: protein-protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets. *Nucleic Acids Res*. 2019; 47(D1):D607–D613. <https://doi.org/10.1093/nar/gky1131> PMID: 30476243
62. Hirakawa H, Takumi-Kobayashi A, Theisen U, Hirata T, Nishino K, Yamaguchi A. AcrS/EnvR represses expression of the acrAB multidrug efflux genes in *Escherichia coli*. *J Bacteriol*. 2008; 190(18):6276–6279. <https://doi.org/10.1128/JB.00190-08> PMID: 18567659
63. Laub MT, Goulian M. Specificity in two-component signal transduction pathways. *Annu Rev Genet*. 2007; 41:121–145. <https://doi.org/10.1146/annurev.genet.41.042007.170548> PMID: 18076326
64. Rivoire O. Elements of coevolution in biological sequences. *Phys Rev Lett*. 2013; 110(17):178102. <https://doi.org/10.1103/PhysRevLett.110.178102> PMID: 23679784
65. 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. <https://doi.org/10.1371/journal.pcbi.1003176> PMID: 23990764
66. Peters JM, Colavin A, Shi H, Czarny TL, Larson MH, Wong S, et al. A Comprehensive, CRISPR-based Functional Analysis of Essential Genes in Bacteria. *Cell*. 2016; 165(6):1493–1506. <https://doi.org/10.1016/j.cell.2016.05.003> PMID: 27238023
67. Barakat M, Ortet P, Jourlin-Castelli C, Ansaldi M, Mejean V, Whitworth DE. P2CS: a two-component system resource for prokaryotic signal transduction research. *BMC Genomics*. 2009; 10:315. <https://doi.org/10.1186/1471-2164-10-315> PMID: 19604365
68. Ortet P, Whitworth DE, Santaella C, Achouak W, Barakat M. P2CS: updates of the prokaryotic two-component systems database. *Nucleic Acids Res*. 2015; 43(Database issue):D536–541. <https://doi.org/10.1093/nar/gku968> PMID: 25324303