Evolutionary couplings detect side-chain interactions

Adam J. Hockenberry[a,1] **and Claus O. Wilke**[a]

[a]Department of Integrative Biology, The University of Texas at Austin

This manuscript was compiled on December 19, 2018

Patterns of amino acid covariation in large protein sequence alignments can inform the prediction of *de novo* protein structures, binding interfaces, and mutational effects. While algorithms that detect these so-called evolutionary couplings between residues have proven useful for practical applications, less is known about how and why these methods perform so well, and what insights into biological processes can be gained from their application. Evolutionary coupling algorithms are commonly benchmarked by comparison to true structural contacts derived from solved protein structures. However, the methods used to determine true structural contacts are not standardized and different definitions of structural contacts may have important consequences for interpreting the results from evolutionary coupling analyses and understanding their overall utility. Here, we show that evolutionary coupling analyses are significantly more likely to identify structural contacts between side-chain atoms than between backbone atoms. We use both simulations and empirical analyses to highlight that purely backbone-based definitions of true residue–residue contacts (i.e., based on the distance between Cα atoms) may underestimate the accuracy of evolutionary coupling algorithms by as much as 40% and that a commonly used reference point (Cβ atoms) underestimates the accuracy by 10–15%. These findings show that co-evolutionary outcomes differ according to which atoms participate in residue–residue interactions and suggest that accounting for different interaction types may lead to further improvements to contact-prediction methods.

Protein evolution | Structural constraints | Contact prediction

A long-standing problem in physical biology is to predict the structure of a protein based solely on its primary amino acid sequence (1–3). Despite advances in x-ray crystallography, NMR spectroscopy, cryo-electron microscopy, and template-based homology modeling, the pace at which new structural data is generated pales in comparison to the rate at which we are accumulating new genomes and gene sequences (4–9). As a result, many protein families still lack even a single representative structure or set of functional annotations (10).

In recent years, however, several groups have made substantial improvements to *de novo* structure prediction by leveraging co-evolutionary information contained within large sequence alignments of protein homologs (11–17). One key result has been that residues that co-evolve with one another frequently do so as a result of their spatial proximity within the protein structure, i.e. a mutation to one site can be compensated for by subsequent mutations to other directly interacting sites (18–20). By determining an "evolutionary coupling" score for all pairs of amino acid residues within a sequence alignment—and assuming that the highest-scoring residue–residue pairs are in close spatial proximity within the structure—the search space of computational protein folding methods can be constrained resulting in accurate 3D-structure determination (14, 21, 22). Other applications have used evolutionary coupling scores to predict protein binding partners and interfaces (12, 23, 24), as well as to predict the effect of mutations on protein stability and function (25). Many of these approaches have been further improved through the use of machine learning (26–28) and deep neural networks that leverage evolutionary couplings along-side numerous other protein features (29–40).

During the development and testing of evolutionary coupling algorithms, predictions of residue–residue couplings from large multiple-sequence alignments are frequently benchmarked against a set of known protein structures to determine their accuracy in identifying residue–residue contacts (41–43). The large number of protein structures that have been solved at atomic resolution provides a training dataset where intramolecular contacts are known (9). However, even the most high-resolution crystal structures of proteins require extrapolation from the location of particular atoms and residues in order to classify residue–residue contacts (44–47). A commonly used heuristic is to determine that any amino acid residue that lies within some pre-defined physical distance—frequently 8Å—of another amino acid residue is said to be in structural contact (14).

Different reference points can be chosen to use as a reference point in determining the distance between any pair of amino acid residues, and prior research has shown that the number of residue–residue contacts identified via side-chain centers is more closely related to evolutionary rates than similar metrics derived from Cα atoms (48–50). However, the consequences of choosing different reference points to determine the accuracy

Significance Statement

Evolutionary couplings between residues within a protein can provide valuable information about protein structures, protein-protein interactions, and the mutability of individual residues. However, the mechanistic factors that determine whether two residues will co-evolve remains unknown. We show that structural proximity by itself is not sufficient for co-evolution to occur between residues. Rather, evolutionary couplings between residues are specifically governed by interactions between side-chain atoms. By contrast, intramolecular contacts between atoms in the protein backbone display only a weak signature of evolutionary coupling. These findings highlight that different types of stabilizing contacts exist within protein structures and that these types have a differential impact on the evolution of protein structures that should be considered in co-evolutionary applications.

Please provide details of author contributions here.

Please declare any conflict of interest here.

[1]To whom correspondence should be addressed. E-mail: adam.hockenberry@utexas.edu

DRAFT

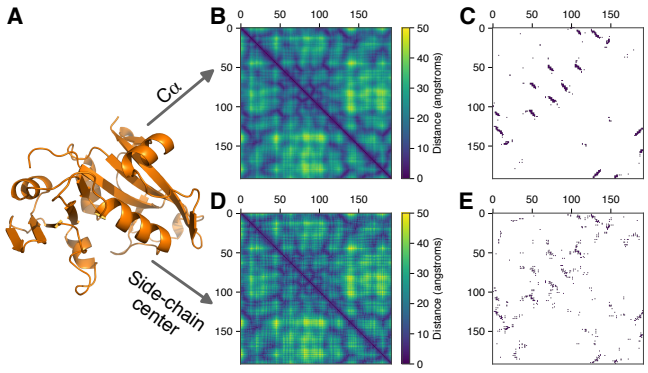

Fig. 1. Constructing contact maps from protein structures. (*A*) An example structure (PDB:1AOE). (*B*) A symmetric distance matrix between all pairs of amino acid residues measured from each residue's Cα atom. (*C*) Medium- to long-range contacts (≥ 12 residues apart along the linear chain are identified using an 8Å cutoff (dark blue). (*D*) and (*E*) Same methodology as depicted in (B) and (C), using the geometric center of each residue's side chain as a reference point for measuring distances.

of modern evolutionary coupling approaches is unknown.

Currently, there are no accepted standards in the field for how to determine a network of residue–residue contacts for a given protein structure. Further, there has yet to be a systematic analysis of whether co-evolutionary signatures are more- or less-closely related to different types of intramolecular contacts that may exist within a protein structure. Here, we systematically test the accuracy of several evolutionary coupling algorithms against true positive contacts defined via Cα, Cβ, or side-chain geometric centers. We find that residue–residue contacts defined according to the distances between side-chain centers are much more accurately predicted by evolutionary couplings. These results imply that the dominant epistatic effects resulting in co-evolutionary signatures arise from interactions between side-chain atoms. Our findings highlight the importance of the choice of contact definition and provide insight into the constraints governing the evolution of protein structures.

Results

Structural contact definitions. Putatively true interactions between amino acid residues within a given protein family are frequently derived from the distance between residues in a representative protein structure. Figure 1 depicts an example protein structure (PDB:1AOE) alongside two symmetric matrices depicting all residue–residue distances (in angstroms, Å) defined according to either the distance between the Cα atoms of individual residues or the distance between the geometric centers of each residue's side chain. From these distance matrices, we created contact matrices by using an 8Å, Cα-based distance threshold to define residue–residue contacts for a given protein. We use the same absolute number of contacts to determine a comparable distance threshold (specific to each protein) to use for side-chain center distances such that an equal number of putatively true contacts were identified regardless of the reference point used to measure distances (SI Fig. S1). Additionally, we note that for most applications the structurally *interesting* contacts are mid- to long-range, which we define here as amino acid pairs separated by a primary chain distance of at least 12 residues; we only consider this subset of possible contacts for the remainder of this manuscript.

While the distance matrices look similar for an example protein when calculated via Cα atoms or side-chain centers (Fig. 1B compared to D), the resulting maps of residue–residue contacts show considerable heterogeneity (Fig. 1C compared to E). More quantitatively, the set of all residue–residue distances measured by Cα atoms are highly correlated with comparable distances measured via side-chain centers (Fig. 2A, left). However, this strong *overall* correlation obscures important differences in *contact* definitions which we observe when focusing within the narrow region where direct amino acid residue contacts are defined (Fig. 2A, right). For 1AOE, we identified a total of 295 contacts according to the 8Å Cα-based distance threshold. Of the shortest 295 contact distances identified via side-chain centers—corresponding to a distance threshold of 7.33Å—the percent of residue–residue pairs that appear in both definitions is only 56%.

To assess the generality of these findings, we applied this analysis to a commonly used benchmark set of 150 proteins (27, 38, 51). Across all of these proteins, we observed a median correlation of 0.97 between residue–residue distances calculated via Cα atoms and side-chain centers (Fig. 2B), but a median overlap of just 63% between contacts defined via Cα and side-chain centers (Fig. 2C).

We focused on comparing the use of Cα atoms and side-chain centers as reference points to determine contacts since these two features represent extreme ends of the spectrum. In practice, Cβ atoms occupy an intermediate location between the backbone and the side-chain center and are one of the most frequently used reference points (44–47). We repeated the primary results of this section by comparing the use of Cβ atoms and side-chain centers as reference points and observed only a 78% overlap in residue–residue contacts identified by

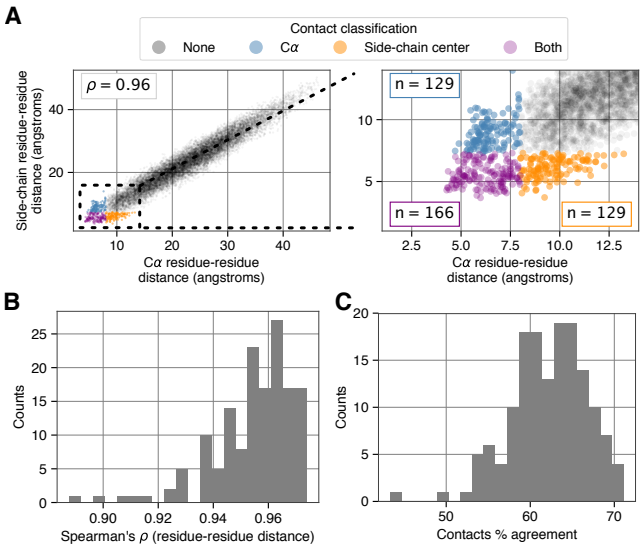

Fig. 2. Relationship between different contact identification methods. (*A*) Correlation between residue–residue distances in PDB:1AOE measured according to Cα atoms and side-chain centers (left). A zoomed in view (right) highlights variably defined residue–residue contacts indicated by the various colors. (*B*) Distribution of Spearman's correlation coefficient values (ρ) between residue–residue distances for 150 different proteins. (*C*) Distribution of the percent agreement for contact definitions for the same set of proteins. (SI Fig.S2 shows a comparable comparison between Cβ and side-chain center distances.)

DRAFT

these two methods (SI Fig. S2). Together, these findings highlight that true contacts vary substantially according to the reference point used to measure residue–residue distances.

Evolutionary couplings reflect side-chain contacts in simulated sequence alignments. While the previous analysis of empirical structures shows that the choice of reference point has important consequences for true contact identification, it is not clear which of the different methods is more biologically correct or practically meaningful. We thus turned our attention to a simplified biophysical system of simulated protein evolution to test the ability of evolutionary coupling analyses to recover intramolecular contacts. We used the ROSETTA modeling software (52–54) to perform all-atom simulations of the evolutionary process (55, 56) while selecting for the maintenance of protein stability (expressed as a fraction of the initial PDB model stability). We simulated thousands of independent evolutionary trajectories, and used the resulting amino acid sequences from these simulations to calculate evolutionary couplings using CCMpred with default parameters (57). Within this defined system, we are able to remove the constraints of phylogenetic biases, limited data availability, homo-oligomerization, insertions/deletions, and changes in evolutionary pressures over time between species—all of which partially limit the power of algorithms to detect true evolutionary couplings in real data (58).

We continued to use 1AOE as an example structure and varied several parameters of our simulation to ensure robust results. We defined true positive residue–residue contacts according to the original PDB structure using residue–residue distances calculated between different quantities for comparison (Cα, Cβ, and side-chain center). To assess the accuracy of evolutionary couplings, we determined the positive predictive value (PPV) of the top $L/2$ couplings—where L is the primary chain length of the protein under investigation, $L = 192$ in the case of 1AOE. For Cα-based contact definitions, we found that the PPV increases rapidly according to the number of independent sequences that we simulated and consequently used as input for evolutionary coupling analyses (Fig. 3A). For these simulations, we ran each independent evolutionary simulation until we accepted a number of point mutations totaling 10 times the length of the protein sequence. Regardless of the selection strength that we imposed on the sequence evolution (colored lines in Fig. 3A), PPV values plateaued at a value below 0.75. This indicates that more than a quarter of the top ranked evolutionary couplings represent incorrect predictions and are not true intramolecular contacts. However, when we analyzed the same evolutionary coupling values but instead used side-chain center distances to define true contacts, we observed that nearly all of the top ranked couplings were defined as true contacts (PPV values approached 1).

We additionally explored how the number of mutations accumulated per sequence affected the ability of evolutionary coupling algorithms to recover intramolecular contacts. We fixed the number of replicate sequences at 3000, and found that PPV values showed minimal variation according to the number of accepted mutations that had accumulated per sequence over the course of our *in silico* evolution (Fig. 3C). As before, however, prediction accuracies were substantially higher when we defined true contacts according to side-chain center distances (Fig. 3D).

These simulation results highlight that—across numerous

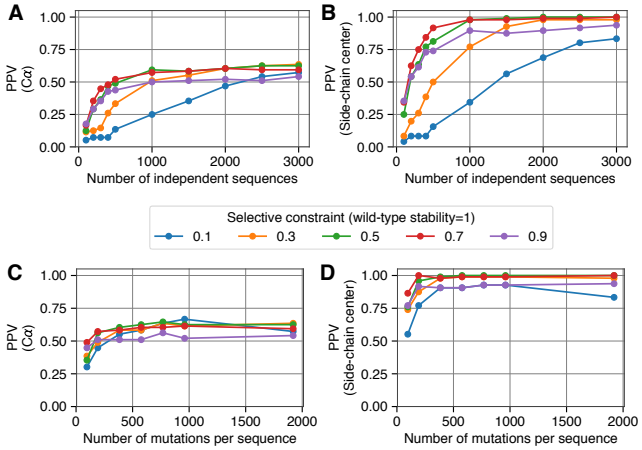

Fig. 3. Comparing simulation-derived evolutionary couplings to different contact definitions. (*A*) For each of 5 separate selection strengths (colored lines), we ran evolutionary simulations until a number of mutations totaling 10 times the length of the protein were accumulated per replicate. We varied the number of independent replicate sequences (x-axis) that were used as input for evolutionary coupling analysis, and found that the resulting evolutionary couplings fail to fully recover Cα defined contacts for PDB:1AOE. (*B*) By contrast, contacts defined via side-chain centers are near-perfectly recovered across a range of simulation parameters. (*C*) and (*D*) Similar to parts (A) and (B), but along the x-axis we now show results from simulations where a different number of accepted mutations were accumulated per sequence. We fixed the number of replicate sequences that were simulated—and used as input for evolutionary coupling analysis—at 3,000 for each of these data points. (Results comparing Cβ and side-chain center contact definitions, can be found in SI Fig. S3.)

parameter combinations—the top $L/2$ evolutionary couplings corresponded to true intramolecular contacts as long as true positives were defined according to side-chain centers and not Cα carbons. Definitions of contacts based on Cβ atoms resulted in intermediate accuracy, plateauing at higher values than Cα but lower than side-chain centers (SI Fig. S3).

Evolutionary couplings reflect side-chain contacts in natural sequence alignments. To see how evolutionary couplings compare to different definitions of true residue–residue contacts in empirical data, we used PHMMER to identify sequence homologs for each of the 150 proteins (see Materials and Methods for details). We assessed the relationship between evolutionary couplings and structural contacts for all proteins by calculating the positive predictive value (PPV) of the highest $L/2$ couplings.

As expected, the PPV between the top $L/2$ evolutionary couplings and Cα-based contacts varied substantially across the 150 structures. This variation may result from a number of different effects, and we observed a clear correlation between PPV values and the number of homologous sequences used to determine evolutionary couplings (SI Fig. S4). Despite the variability in prediction accuracy between proteins, we observed systematic differences in the PPV according to which metric was used to identify true positive contacts from the PDB structure files (Fig. 4A). When compared to Cα based distances, residue–residue distances measured according to Cβ atoms resulted in significantly higher PPVs, and side-chain based contact distances resulted in even further improvements (Wilcoxon signed-rank test, $p < 10^{-20}$ for all comparisons). Further, the magnitude of the effect was substantial: across all 150 proteins the median percent increase in PPV between Cα and side-chain center contact identification methods was 43%

DRAFT

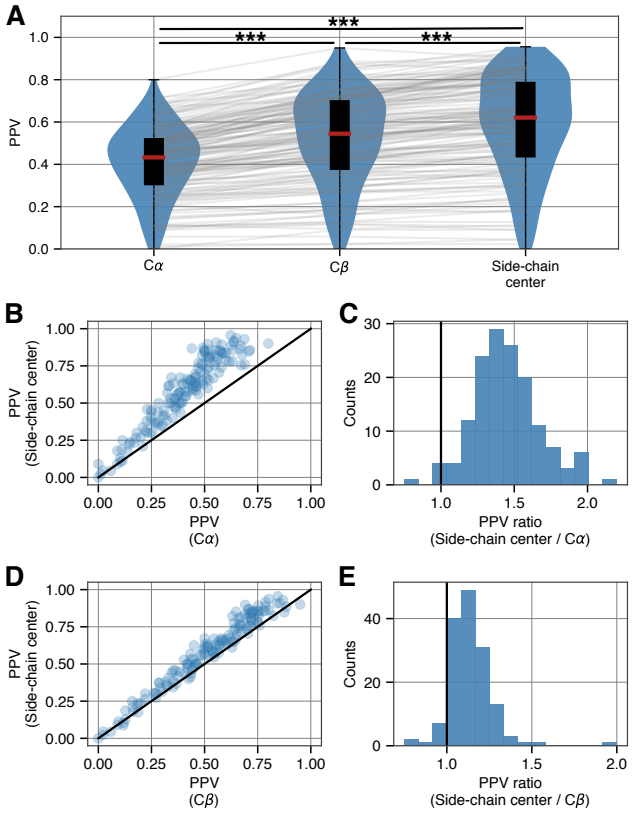

Fig. 4. Accuracy of evolutionary couplings derived from empirical alignments. (*A*) For a diverse set of 150 proteins, the PPV of the top $L/2$ evolutionary coupling scores—derived from empirical sequence alignments—is progressively higher when intramolecular contacts are defined according to Cα atoms, Cβ atoms, and side-chain centers. (*** indicates $p < 10^{-20}$, Wilcoxon signed-rank test) (*B*) Scatter plot of PPVs for each protein according to Cα and side-chain center contact identification methods. (*C*) Histogram of the ratios from the data in (B) indicate that the median percent increase in accuracy is 43%. (*D*) and (*E*) As in (B) and (C), but now comparing Cβ and side-chain center contact identification methods. Results show a median percent increase in contact identification accuracy of 13%. (Results for two other evolutionary coupling algorithm implementations can be found in SI Figs. S5 & S6.)

(Fig. 4B,C). Even between the more similar Cβ and side-chain center methods, the median percent increase in accuracy was 13% (Fig. 4D,E). Both comparisons were highly significant and persisted across the entire range of PPVs represented within our dataset (Fig. 4B,D). Additionally, these results were highly consistent across different evolutionary coupling algorithms (SI Figs. S5 & S6).

Side-chain orientations are important to consider when examining residue–residue interactions. The preceding analyses have shown that for the the exact same evolutionary couplings, PPVs are substantially higher when using side-chain based distances to identify true positive intramolecular contacts than when using either Cα or Cβ-based distances. To look more specifically at why these differences were so pronounced, we decided to investigate the orientation of residue–residue pairs identified by the various criteria. Any two residues considered to be in contact can display one of three distinct orientations of their respective side chains: i) both residues' side chains may point towards one another with the energetic interactions occurring through side-chain atoms, ii) one residue's side chain may point towards the other residue while that residue's side chain points away, or iii) both residues' side chains may point away from one another with energetic interactions occurring between the respective amino acid backbones (Fig. 5A). As we expected, the Cα-based definition yields fractions close to the random expectation for 1AOE: in ∼25% of the cases side chains point towards each other, in another ∼25% of the cases they point away from each other, and in the remaining ∼50% of the cases one side chain points towards the other while the other points away (Fig. 5B). By contrast, side-chain (and to a lesser degree Cβ) based contact definitions enrich for cases where both side chains point towards one another (Fig. 5B).

Across all 150 proteins in our dataset, we calculated the fraction of *all* residue–residue pairs (regardless of distance in 3D space but subject to the same primary chain distance constraints applied throughout this manuscript) where *both* side chains point towards one another and found a median of 0.2 (Fig. 5C, "All pairs"). However, this fraction increases progressively when we consider residue–residue pairs identified as true contacts for each protein according to Cα, Cβ, and side-chain centers—illustrating that the trend observed in (Fig. 5B) applies broadly. If instead we consider residue–residue pairs predicted by the top ranked evolutionary couplings, we observe that a large fraction of the these couplings are between residues whose side chains point towards one another in the reference protein structure. Additionally, this fraction is highest for the most highly ranked evolutionary couplings (Top 0.25L, where L is the length of the protein) and is substantially higher than the proportion identified by Cα-based distances (Fig. 5C).

To further illustrate this point, we turned to an alternative method for determining intramolecular contacts: determining structural contacts based on the minimum distance between any two heavy atoms for each residue–residue pair (22). We implemented two versions of this algorithm, determining the

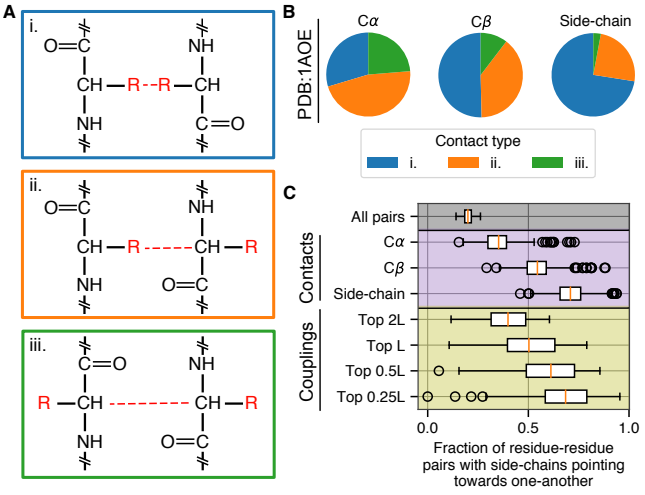

Fig. 5. Different types of residue–residue interactions are possible. (*A*) Two interacting residues may interact via: each residue's side-chain atoms (type i), the side chain of one residue and the backbone of the other (type ii), or the backbone atoms of each residue (type iii). (*B*) For intramolecular contacts identified in PDB:1AOE, the relative proportion of different interaction types varies according to contact identification method: residue–residue contacts defined via side-chain centers are enriched in type i interactions (blue). (*C*) For 150 proteins, the fraction of residue–residue pairs where the side chains point towards one-another is highest in contacts defined via side-chain centers (purple). Further, the top ranked evolutionary couplings are progressively enriched in residue–residue pairs where the side chains point towards one another (yellow).

DRAFT

minimum distance between: i) all heavy atoms within residues and ii) side-chain heavy atoms only. In each case, and to facilitate comparison between methods, we selected the n shortest distances as contacts where n is the number of contacts identified for each protein via the 8Å distance threshold using Cα. For the set of 150 proteins, the resulting PPVs were significantly higher when contacts were defined *only* according to side-chain atoms as opposed to the complete set of backbone and side-chain atoms (SI Fig. S7). Furthermore, PPVs calculated via side-chain center distances were statistically indistinguishable from PPVs derived from the minimum distance between all heavy atoms within side chains.

Taken together, our analyses of side-chain orientations and contacts identified via minimum atomic distances both highlight that evolutionary couplings frequently occur between residues whose side chains point towards one another. Cα (and to a lesser extent Cβ) based contact definitions classify a smaller number of contacts in this orientation, and including backbone atoms in minimum-distance-based contact identification methods actually decreases the accuracy of contact predictions based on evolutionary couplings.

Discussion

The co-evolutionary patterns of amino acid substitutions can provide important information about protein structures. A number of competing methods are currently in use to detect evolutionary couplings between residues, and the ability to recover true residue–residue contacts has been the primary metric employed to assess performance of various methods. However, true structural contacts are ill-defined and variability in contact definitions can prohibit comparison between the efficacy of different methods as well as obscure the biological interpretation of evolutionary constraints. We have shown here that evolutionary couplings are significantly more accurate at detecting true residue–residue contacts based off of side-chain center distances than based off of Cα or Cβ distances. This finding provides important biological insight into protein evolution and epistatic interactions between residues. We posit that although different types of interactions between amino acid residues may stabilize protein structures, evolutionary couplings predominantly represent residues whose contact occurs via interactions between the side-chain atoms of both residues.

While we have shown that using side-chain centers as a reference point to define intramolecular contacts is advantageous compared to Cα or Cβ atoms, it is important to note that the contacts defined via side-chain centers are still likely to be only a rough approximation of reality. Actual contacts between residues within a single structure would ideally be defined via knowledge of the hydrogen bonding, van der Waals forces, and ionic/dipole/hydrophobic interactions that collectively stabilize protein structures (59–61). Defining contacts via these features can be time-consuming and resource intensive and the use of distance-based cutoffs is therefore advantageous under many circumstances. Furthermore, a highly accurate view of residue–residue contacts may still not be particularly useful or informative. For instance, our findings do not suggest that contacts between the backbone atoms of residues do not occur or that they are unimportant for stabilizing protein structures. Rather, we show that contacts between backbone atoms are not likely to be detected by evolutionary coupling analyses. Thus, regardless of the method used for defining residue–residue contacts, different contact types may be of varying importance for different applications.

Amongst distance-based methods, a potential disadvantage of using side-chain centers as reference points is that the geometric center of different amino acid side chains lies at varying absolute distances from the side-chain atoms. Here, we have used arbitrary distance cutoffs to define side-chain center contacts (based on identifying an equal number of contacts per protein as an 8Å Cα cutoff). This facilitates comparison of different contact identification methods on the same scale, but a more rigorous approach would be preferable in practice. The method of evaluating the minimum distance between heavy atoms is attractive in this regard, since it permits using a physically realistic and uniform distance cutoff between atoms regardless of amino acid identity or the protein under consideration. However, our analysis shows that it is critical to only consider side-chain atoms when performing this type of calculation, and we found that results from this method are statistically indistinguishable from using side-chain centers. Nevertheless, it is possible that analysis of still larger datasets might uncover differences in accuracies between these conceptually similar methods.

Leveraging homologous sequence information to predict intramolecular protein contacts has long been a goal of structural biologists, but progress towards this goal has been made possible only in recent years with increasing availability of sequence information and the development and application of Direct Coupling Analysis to isolate directly interacting residue–residue correlations from spurious correlation generated by chains of directly coupled residues (12, 13, 51, 57, 62). However, there are still a number of known limitations to current methods for computing evolutionary couplings as illustrated by the variability in performance that we observe when applying the same basic algorithms to collect and evaluate different homologous protein sets (Fig. 4A). For some protein families, top ranked evolutionary couplings correspond closely to intramolecular contacts within the reference structure whereas for other families this correspondence is close to random chance.

Most evolutionary coupling algorithms require large numbers of sequence homologs to produce accurate predictions (22); we observed this in our own data, where the number of sequences in a multiple sequence alignment is significantly correlated with the PPV of inferred evolutionary couplings (SI Fig. S4). However, the size of the multiple sequence alignment explains only a moderate fraction of the variation in accuracies. Features such as the evolutionary relatedness of sequences and the heterogeneity of substitution rates across sites may impose further constraints on the overall identifiability of evolutionary couplings, but it's not clear whether these effects should systematically vary between different protein families (63).

False positive predictions resulting from evolutionary coupling analyses may arise from a combination of different factors including problems with repeat proteins, homo-oligomerization, and structural variation within protein families (58). Here, we have shown that another source of false positives in evolutionary coupling analyses may simply be ill-defined true positive contacts. We also observed that accuracy improvements stemming from the use of side-chain centers as a reference point were variable across proteins (Fig. 4C,E)—this variability may

DRAFT

be related to features of the proteins themselves in terms of their compactness or other physio-chemical properties.

Practically speaking, contact identification methods and evolutionary couplings are increasingly important for a variety of applications. Frequently, evolutionary couplings are combined with a variety of other features and used as input for machine learning and neural network-based algorithms to predict structural and functional properties of proteins (26–40). Our analysis suggests that there may be biases in the training data—essential to supervised learning techniques—owing to the method used to define true positive contacts. Given that raw evolutionary couplings more accurately predict contacts identified by side-chain centers, we speculate that the accuracy of machine learning and neural network-based algorithms may be similarly improved by training and testing on contacts identified via side-chain distances. Different contact types can be classified according to which atoms interact between residue–residue pairs (Fig. 5A), and it is possible that the accuracy of supervised algorithms in particular may be improved by separating different types of residue–residue contacts according to their atomic interactions, training separate models to detect each type, and integrating the results.

Materials and Methods

Dataset compilation and processing. We downloaded the so-called PSICOV dataset of 150 proteins that have been extensively studied (27, 38, 51). We processed each starting PDB file to select a single chain, ensure a consistent numbering of residues ($1...n$), test for unknown or non-standard residues, select the most likely state for all disordered sequence atoms, and remove all extraneous information (including "HETATM" lines). Next, to ensure that all residues were represented in full and repair those that were not, we used PYROSETTA to read in the ".PDB" files using the 'pose_from_pdb' function and wrote the output as our final clean structure (52).

Determining structural contacts and contact-types. From each cleaned ".PDB" file, we calculated residue–residue distance matrices using custom python scripts (the euclidean distance from 3-dimensional atomic coordinates). All residues contain a Cα atom so this calculation was straightforward. For Cβ calculations, we used the Cβ atom of all residues except glycine, for which we continued to use the Cα atom. For side-chain center calculations, we calculated the geometric center of each residue based on the coordinates of all non-backbone heavy atoms. This calculation included Cβ atoms but excluded Cα atoms for all amino acids except glycine, for which we again continued to use Cα as the reference point.

To calculate minimum atomic distances between two residues, we calculated all pairwise euclidean distances between heavy atoms and selected the minimum distance. In extending this analysis to only consider side-chain atoms, we continued to consider Cβ atoms as part of the side chain but not Cα. Again, we relaxed this restriction for glycine and included Cα as a side-chain atom to permit calculations for this amino acid.

For all methods, contacts were assessed by first removing all residue–residue pairs where the two residues were shorter than 12 amino acids apart in primary chain distance. Contacts were determined throughout this manuscript for each structure according to an 8Å cutoff between Cα atoms. Since accuracy values are partially dependent on the number of true positives that are called, we maintained a constant number of true positive contact classifications throughout to facilitate comparison between methods. For each contact definition (Cβ, side-chain center, minimum atomic distances), we selected n residue–residue pairs with the shortest distances where n is the number of contacts defined according to the aforementioned Cα-based method. Thus, the actual distance threshold used to classify contacts for these other metrics varies slightly from protein to protein (SI Fig. S1).

To classify residue–residue pairs (a and b) via their side-chain orientations, we chose a reference residue (a) and drew two vectors: i) from the Cα atom coordinate to the side-chain center for that residue and ii) from the Cα atom coordinate to the Cα atom coordinate for the other residue in question (b). If the angle between these two vectors was less than $\pi/2$ radians, the side chain of residue a was said to point towards residue b. To determine the residues classification as in Fig.5A, we next repeated the calculation using residue b as the reference and classified the residue–residue pair accordingly.

Evolutionary coupling algorithms. For each of the 150 proteins in our dataset, we followed a principled method to retrieve homologous sequences. We first extracted the primary amino acid sequence from the ".PDB" file. We next used PHMMER to search through progressively larger databases in order to retrieve up to 10,000 homologous sequences (64). To do so, we downloaded local versions of the rp15, rp35, rp55, and rp75 databases (65). We first searched the smallest, least redundant, database for each sequence using an E-value threshold of 0.0001. For any sequence with greater than 10,000 hits we stopped and selected the top scoring 10,000 hits for further analysis. For sequences with fewer than 10,000 hits we moved to the next largest database and repeated the process. Finally for the small number of sequences for which we did not accumulate at least 1,000 sequences in the largest database (rp75), we used the online version of PHMMER to search the UniprotKB database and downloaded the maximum results.

For each protein, we next aligned the hits along-side the reference sequence using MAFFT (v7.380, L-INS-i method with default parameters) (66). Next, we cleaned these results to remove all columns that were gapped in the reference (".PDB") sequence. All other columns and sequences in the sequence alignments were retained regardless of gap coverage.

Using these alignments, we next calculated evolutionary couplings between residue–residue pairs. All results in the main manuscript are displayed using CCMpred with default parameters (0.8 local sequence re-weighting threshold, 0.2 pairwise regularization coefficients, average product correction) (57). We additionally used the 'plmc' method from the EVcouplings framework with default parameters (no average product correction) (67) and PSICOV (default parameters excepting: "-z 50 -r 0.001") (51) to ensure the robustness of our findings.

Except where otherwise noted (Fig. 5), main text results (Fig. 3, Fig. 4) were calculated using the top $L/2$ couplings for each protein where L is the length of the reference amino acid sequence. Positive Predictive Value (PPV) is calculated as the number of classified contacts among these top couplings divided by the total number of top couplings considered.

Evolutionary simulations. For the example protein used throughout the text (PDB:1AOE) we performed mutation accumulation simulations using PYROSETTA (52). We first read in the ".PDB" structure (with disulfide bonds disabled), and minimized it to optimize thermodynamic stability via rotamer selection and backbone movements. We next fixed the backbone, and implemented an expedited evolutionary algorithm to select amino acid point mutations (no insertions or deletions were allowed) according to their predicted impact on structural stability (55, 56). At each step, we selected a random amino acid position, and attempted to mutate it randomly to one of the remaining 19 amino acids. We re-packed the structure within a 12Å radius and determined whether or not to accept the mutation based off of the resulting change in structural stability. Mutations which either did not alter or increased stability (i.e. resulted in a decreased ΔG) were accepted with a probability of 1. Mutations that decreased stability were accepted with a probability proportional to their $\Delta\Delta G$ as in Teufel and Wilke (55). At the end of the evolutionary process, the resulting amino acid sequence was stored for future analysis.

We performed thousands of independent replicates of this expedited evolutionary process where we altered the number of accepted mutations that we accumulated before halting the simulation, the number of replicate evolutionary experiments that we performed, and the fraction of the initial wild-type stability value that we used for our selection criteria. Collections of the resulting sequences were

DRAFT

analyzed via evolutionary coupling algorithms in the same manner as empirical sequences, with no need for sequence alignment.

Data access. All code and data are freely available at: https://github.com/adamhockenberry/side-chain-couplings.

ACKNOWLEDGMENTS. This work was funded by National Science Foundation Cooperative Agreement no. DBI-0939454 (BEACON Center) and National Institutes of Health grant R01 GM088344. The authors acknowledge valuable feedback and support from members of the Wilke lab.

1. Anfinsen CB (1973) Principles that Govern the Folding of Protein Chains. *Science* 181(4096):223–230.
2. Sadowski MI, Jones DT (2009) The sequence-structure relationship and protein function prediction. *Current Opinion in Structural Biology* 19(3):357–362.
3. Marks DS, Hopf TA, Sander C (2012) Protein structure prediction from sequence variation. *Nature Biotechnology* 30(11):1072–1080.
4. Söding J (2005) Protein homology detection by HMM-HMM comparison. *Bioinformatics* 21(7):951–960.
5. Liao M, Cao E, Julius D, Cheng Y (2013) Structure of the TRPV1 ion channel determined by electron cryo-microscopy. *Nature* 504(7478):107–112.
6. Amunts A, et al. (2014) Structure of the Yeast Mitochondrial Large Ribosomal Subunit. *Science* 343:1485–1489.
7. Biasini M, et al. (2014) SWISS-MODEL: Modelling protein tertiary and quaternary structure using evolutionary information. *Nucleic Acids Research* 42:252–258.
8. Punjani A, Rubinstein JL, Fleet DJ, Brubaker MA (2017) CryoSPARC: Algorithms for rapid unsupervised cryo-EM structure determination. *Nature Methods* 14(3):290–296.
9. Rose PW, et al. (2017) The RCSB protein data bank: Integrative view of protein, gene and 3D structural information. *Nucleic Acids Research* 45:D271–D281.
10. Wyman SK, Avila-herrera A, Nayfach S, Id SP (2018) A most wanted list of conserved microbial protein families with no known domains. *PLoS One* 13(10):e0205749.
11. Lapedes AS, Giraud BG, Liu L, Stormo GD (1999) Correlated mutations in models of protein sequences: phylogenetic and structural effects. *Statistics in molecular biology and genetics* 33:236–256.
12. Burger L, Van Nimwegen E (2008) Accurate prediction of protein-protein interactions from sequence alignments using a Bayesian method. *Molecular Systems Biology* 4(165).
13. Burger L, Van Nimwegen E (2010) Disentangling direct from indirect co-evolution of residues in protein alignments. *PLoS Computational Biology* 6(1).
14. Marks DS, et al. (2011) Protein 3D structure computed from evolutionary sequence variation. *PLoS ONE* 6(12).
15. Morcos F, et al. (2011) Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences* 108(49):E1293–E1301.
16. Sulkowska JI, Morcos F, Weigt M, Hwa T, Onuchic JN (2012) Genomics-aided structure prediction. *Proceedings of the National Academy of Sciences* 109(26):10340–10345.
17. Kamisetty H, Ovchinnikov S, Baker D (2013) Assessing the utility of coevolution-based residue-residue contact predictions in a sequence- and structure-rich era. *Proceedings of the National Academy of Sciences* 110(39):15674–15679.
18. Gobel U, Sander C, Schneider R, Valencia A (1994) Correlated Mutations and Residue Contacts in Proteins. *Proteins* 18:309–317.
19. Shindyalov IN, Kolchanov NA, Sander C (1994) Can three-dimensional contacts in protein structures be predicted by analysis of correlated mutations? *Protein Engineering, Design and Selection* 7(3):349–358.
20. Weigt M, White RA, Szurmant H, Hoch JA, Hwa T (2009) Identification of direct residue contacts in protein-protein interaction by message passing. *Proceedings of the National Academy of Sciences* 106(1):67–72.
21. Hopf TA, et al. (2012) Three-dimensional structures of membrane proteins from genomic sequencing. *Cell* 149(7):1607–1621.
22. Ovchinnikov S, et al. (2017) Protein structure determination using metagenome sequence data. *Science* 355(6322):294–298.
23. Hopf TA, et al. (2014) Sequence co-evolution gives 3D contacts and structures of protein complexes. *eLife* 3:1–45.
24. Ovchinnikov S, Kamisetty H, Baker D (2014) Robust and accurate prediction of residue-residue interactions across protein interfaces using evolutionary information. *eLife* pp. 1–21.
25. Hopf TA, et al. (2017) Mutation effects predicted from sequence co-variation. *Nature Biotechnology* 35(2):128–135.
26. Cheng J, Baldi P (2007) Improved residue contact prediction using support vector machines and a large feature set. *BMC Bioinformatics* 8(1).
27. Jones DT, Singh T, Kosciolek T, Tetchner S (2015) MetaPSICOV: Combining coevolution methods for accurate prediction of contacts and long range hydrogen bonding in proteins. *Bioinformatics* 31(7):999–1006.
28. Michel M, Skwark MJ, Hurtado DM, Ekeberg M, Elofsson A (2017) Predicting accurate contacts in thousands of Pfam domain families using PconsC3. *Bioinformatics* 33(18):2859–2866.
29. Tegge AN, Wang Z, Eickholt J, Cheng J (2009) NNcon: Improved protein contact map prediction using 2D-recursive neural networks. *Nucleic Acids Research* 37:515–518.
30. Lena PD, Nagata K, Baldi P (2012) Deep architectures for protein contact map prediction. *Bioinformatics* 28(19):2449–2457.
31. Xiong D, Zeng J, Gong H (2017) A deep learning framework for improving long-range residue-residue contact prediction using a hierarchical strategy. *Bioinformatics* 33(17):2675–2683.
32. Stahl K, Schneider M, Brock O (2017) EPSILON-CP: Using deep learning to combine information from multiple sources for protein contact prediction. *BMC Bioinformatics* 18(1):1–11.
33. He B, Mortuza SM, Wang Y, Shen HB, Zhang Y (2017) NeBcon: Protein contact map prediction using neural network training coupled with naïve Bayes classifiers. *Bioinformatics* 33(15):2296–2306.
34. Wang S, Sun S, Li Z, Zhang R, Xu J (2017) *Accurate De Novo Prediction of Protein Contact Map by Ultra-Deep Learning Model.* Vol. 13.
35. Riesselman AJ, Ingraham JB, Marks DS (2018) Deep generative models of genetic variation capture mutation effects. *Nature Methods.*
36. Liu Y, Palmedo P, Ye Q, Berger B, Peng J (2018) Enhancing Evolutionary Couplings with Deep Convolutional Neural Networks. *Cell Systems* 6(1):65–74.e3.
37. Wozniak PP, Pelc J, Skrzypecki M, Vriend G, Kotulska M (2018) Bio-knowledge based filters improve residue- residue contact prediction accuracy. *Bioinformatics* (May):1–9.
38. Jones DT, Kandathil SM (2018) High precision in protein contact prediction using fully convolutional neural networks and minimal sequence features. *Bioinformatics* (April):1–8.
39. Adhikari B, Hou J, Cheng J (2018) DNCON2: Improved protein contact prediction using two-level deep convolutional neural networks. *Bioinformatics* 34(9):1466–1472.
40. Hanson J, Paliwal K, Litfin T, Yang Y, Zhou Y (2018) Accurate Prediction of Protein Contact Maps by Coupling Residual Two-Dimensional Bidirectional Long Short-Term Memory with Convolutional Neural Networks. *Bioinformatics* (June):bty481–bty481.
41. Figliuzzi M, Barrat-Charlaix P, Weigt M (2018) How pairwise coevolutionary models capture the collective residue variability in proteins? *Molecular Biology and Evolution* 35(4):1018–1027.
42. Schaarschmidt J, Monastyrskyy B, Kryshtafovych A, Bonvin AM (2018) Assessment of contact predictions in CASP12: Co-evolution and deep learning coming of age. *Proteins: Structure, Function and Bioinformatics* 86(October 2017):51–66.
43. Wang S, Sun S, Xu J (2018) Analysis of deep learning methods for blind protein contact prediction in CASP12. *Proteins: Structure, Function and Bioinformatics* 86(August 2017):67–77.
44. Seeliger D, de Groot BL (2007) Atomic contacts in protein structures. A detailed analysis of atomic radii, packing, and overlaps. *Proteins* 68:595–601.
45. Sathyapriya R, Duarte JM, Stehr H, Filippis I, Lappe M (2009) Defining an essence of structure determining residue contacts in proteins. *PLoS Computational Biology* 5(12).
46. Duarte JM, Sathyapriya R, Stehr H, Filippis I, Lappe M (2010) Optimal contact definition for reconstruction of Contact Maps. *BMC Bioinformatics* 11.
47. Yuan C, Chen H, Kihara D (2012) Effective inter-residue contact definitions for accurate protein fold recognition. *BMC Bioinformatics* 13(1).
48. Lin CP, et al. (2008) Deriving protein dynamical properties from weighted protein contact number. *Proteins* 72(3):929–935.
49. Marcos ML, Echave J (2015) Too packed to change: side-chain packing and site-specific substitution rates in protein evolution. *PeerJ* 3:e911.
50. Shahmoradi A, Wilke CO (2016) Dissecting the roles of local packing density and longer-range effects in protein sequence evolution. *Proteins* 84(6):841–854.
51. Jones DT, Buchan DW, Cozzetto D, Pontil M (2012) PSICOV: Precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. *Bioinformatics* 28(2):184–190.
52. Chaudhury S, Lyskov S, Gray JJ (2010) PyRosetta: A script-based interface for implementing molecular modeling algorithms using Rosetta. *Bioinformatics* 26(5):689–691.
53. Leaver-Fay A, et al. (2011) Rosetta3: An object-oriented software suite for the simulation and design of macromolecules. *Methods in Enzymology* 487(Chapter 19):545–574.
54. Kellogg EH, Leaver-Fay A, Baker D (2011) Role of conformational sampling in computing mutation-induced changes in protein structure and stability. *Proteins* 79(3):830–838.
55. Teufel AI, Wilke CO (2017) Accelerated simulation of evolutionary trajectories in origin-fixation models. *Journal of the Royal Society Interface* 14(127).
56. Jiang Q, Teufel AI, Jackson EL, Wilke CO (2018) Beyond Thermodynamic Constraints: Evolutionary Sampling Generates Realistic Protein Sequence Variation. *Genetics* 208:1387–1395.
57. Seemayer S, Gruber M, Söding J (2014) CCMpred - Fast and precise prediction of protein residue-residue contacts from correlated mutations. *Bioinformatics* 30(21):3128–3130.
58. Anishchenko I, Ovchinnikov S, Kamisetty H, Baker D (2017) Origins of coevolution between residues distant in protein 3D structures. *Proceedings of the National Academy of Sciences* p. 201702664.
59. Abdel-Azeim S, Chermak E, Vangone A, Oliva R, Cavallo L (2014) MDcons : Intermolecular contact maps as a tool to analyze the interface of protein complexes from molecular dynamics trajectories. *BMC Bioinformatics* 15:1–11.
60. Doshi U, Holliday MJ, Eisenmesser EZ, Hamelberg D (2016) Dynamical network of residue–residue contacts reveals coupled allosteric effects in recognition, catalysis, and mutation. *Proc Natl Acad Sci USA* 113(17):4735–4740.
61. Mercadante D, Gräter F, Daday C (2018) CONAN: A tool to decode dynamical information from molecular interaction maps. *Biophysical Journal* 114:1267–1273.
62. Ekeberg M, Lövkvist C, Lan Y, Weigt M, Aurell E (2013) Improved contact prediction in proteins: Using pseudolikelihoods to infer Potts models. *Physical Review E* 87(1):1–16.
63. Vorberg S, Seemayer S, Söding J (2018) Synthetic protein alignments by CCMgen quantify noise in residue-residue contact prediction. *PLoS Computational Biology* 14(11):e1006526.
64. Potter SC, et al. (2018) HMMER web server: 2018 update. *Nucleic Acids Research* 46(W1):W200–W204.
65. Chen C, et al. (2011) Representative Proteomes: A Stable, Scalable and Unbiased proteome set for sequence analysis and functional annotation. *PLoS ONE* 6(4):e18910.
66. Katoh K, Standley DM (2013) MAFFT multiple sequence alignment software version 7: Improvements in performance and usability. *Molecular Biology and Evolution* 30(4):772–780.
67. Hopf TA, et al. (2018) The EVcouplings Python framework for coevolutionary sequence analysis Thomas. *Bioinformatics* p. bty862.

DRAFT

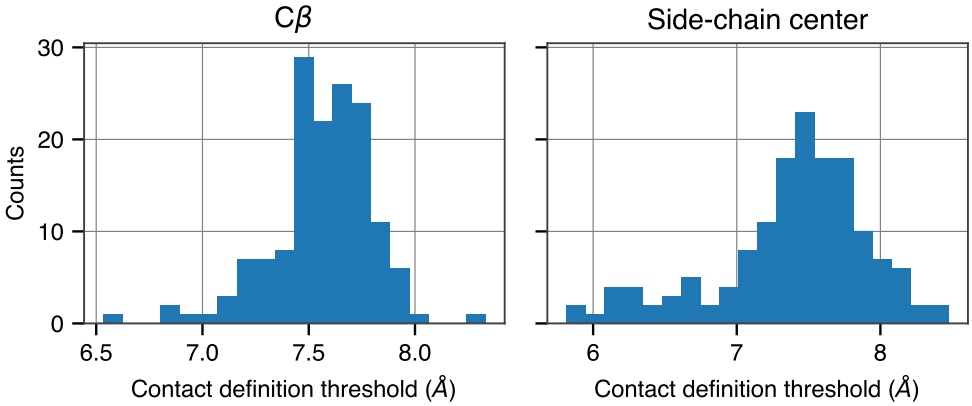

SI Fig. S1. To directly compare contact identification metrics, contacts for a given protein were first defined according to Cα atoms with an 8Å distance threshold. For Cβ (left) and side-chain center (right), an equal number of putative contacts were then identified regardless of their distances. Histograms show the relevant distance thresholds that ensure an equal number of putative contacts for each metric for the set of 150 proteins.

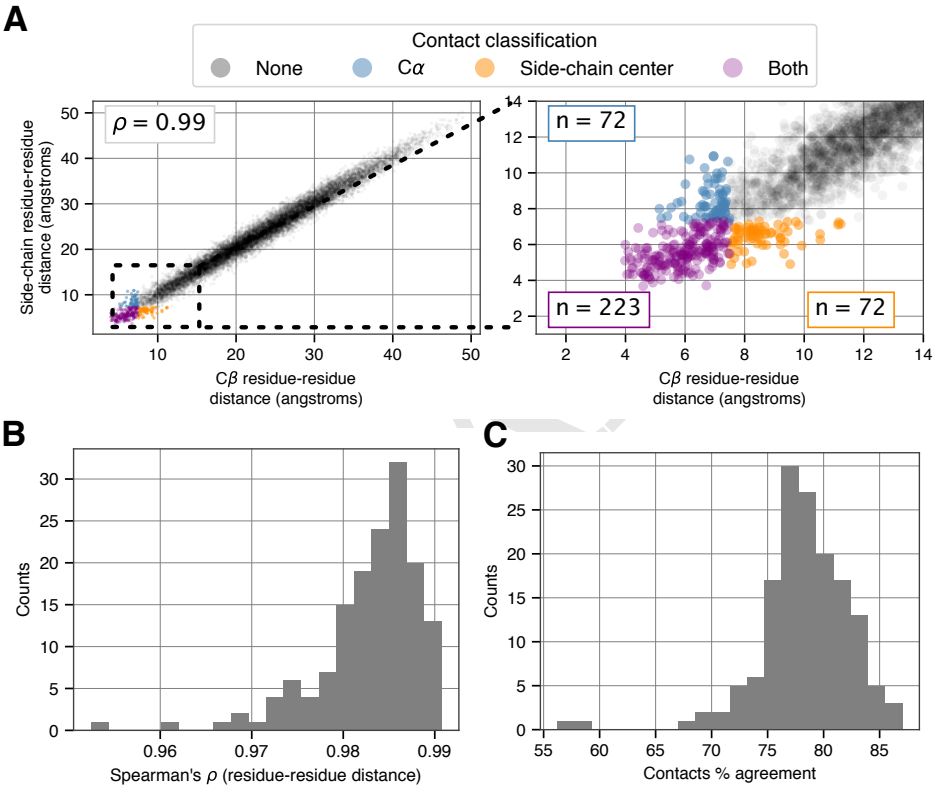

SI Fig. S2. Similar to Fig. 2 of the main text; here, we compare residue–residue distances defined via Cβ and side-chain centers. (*A*) For PDB:1AOE, Cβ and side-chain center contacts were defined according to 7.5Å and 7.33Å distance thresholds, respectively. (*B*) Distribution of Spearman's correlation coefficient values (ρ) between residue–residue distances for 150 different proteins. (*C*) Distribution of the percent agreement for contact definitions for the same set of proteins.

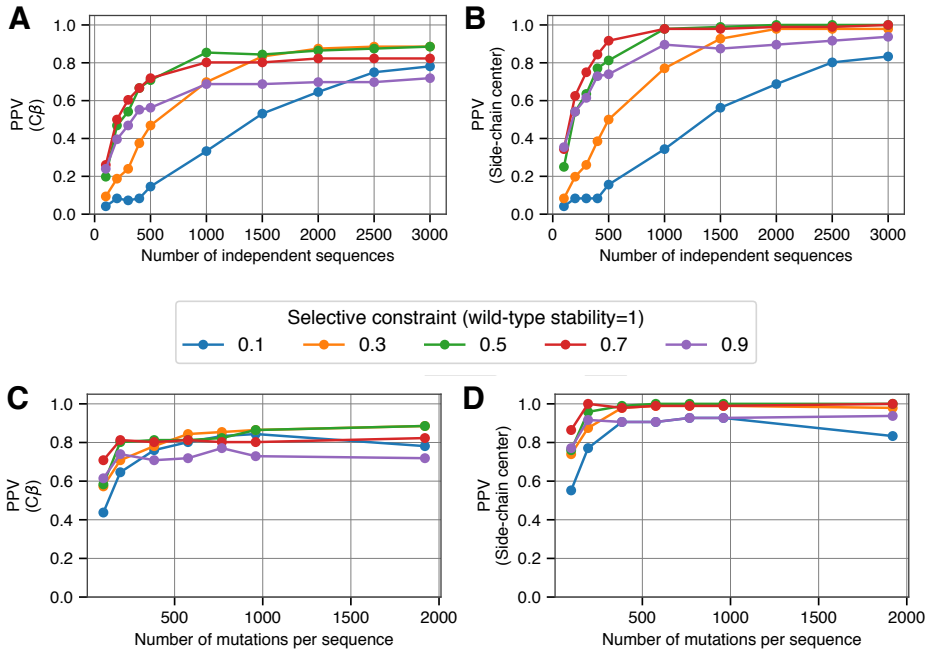

SI Fig. S3. Similar to Fig. 3 of the main text; here, we compare Cβ and side-chain center contact definitions. (*A*) For each of 5 separate selection strengths (colored lines), we ran simulations until a number of mutations totaling 10 times the length of the protein were accumulated per replicate. We varied the number of independent replicate sequences (x-axis) that were used as input for evolutionary coupling analysis, and found that couplings fail to fully recover Cα defined contacts for PDB:1AOE. (*B*) By contrast, contacts defined via side-chain centers are near-perfectly recovered across a range of simulation parameters. (*C*) and (*D*) Similar to parts (A) and (B), but along the x-axis we now show results from simulations where a different number of accepted mutations were accumulated per sequence. We fixed the number of replicate sequences that were simulated—and used for evolutionary coupling analysis—at 3,000 for each of these data points. (note: side-chain center plots—(B) and (D)—are the same as in Fig. 3 and are included here only for comparison.)

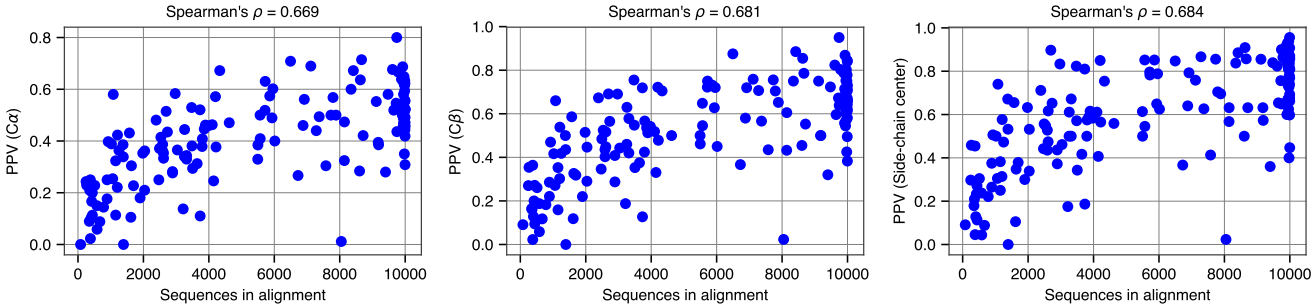

SI Fig. S4. PPVs for all three different contact identification methods (y-axes) are strongly correlated with the number of sequences in the empirical sequence alignments used as input for calculating evolutionary couplings. PPV data here is the same data from Fig. 4.

DRAFT

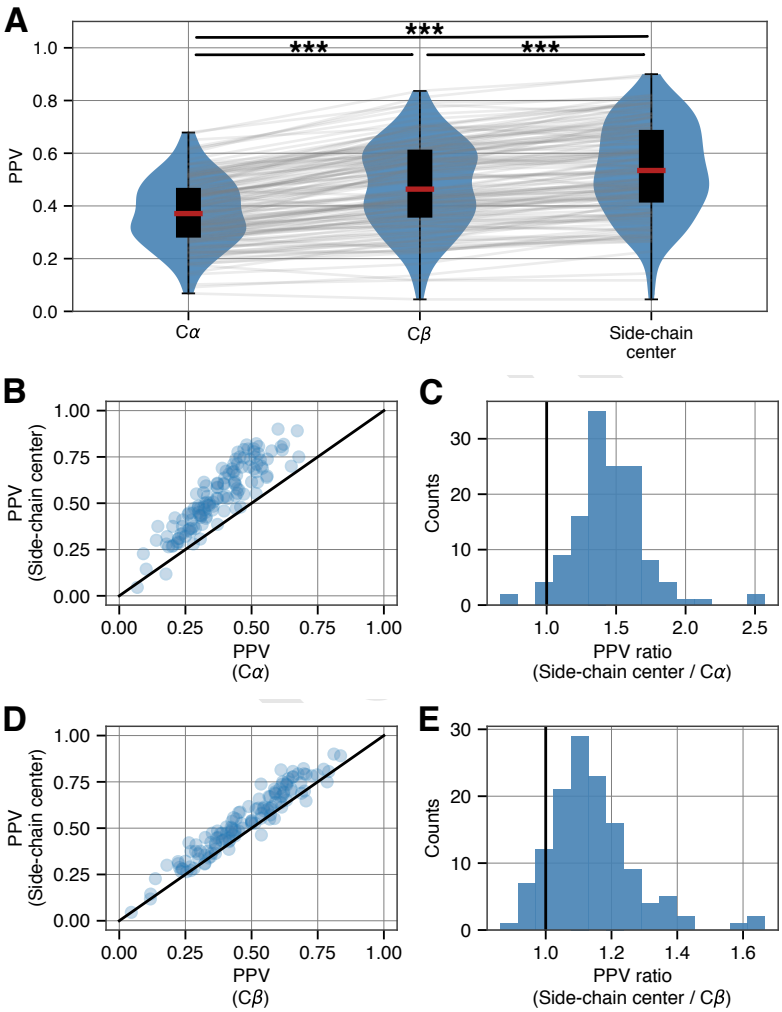

SI Fig. S5. Similar to Fig. 4 of the main text; here, we use the PSICOV method to calculate evolutionary couplings. (*A*) For a diverse set of 150 proteins, the PPV of the top $L/2$ evolutionary coupling scores—derived from empirical sequence alignments—is progressively higher when intramolecular contacts are defined according to Cα atoms, Cβ atoms, and side-chain centers. (*** indicates $p < 10^{-20}$, Wilcoxon signed-rank test) (*B*) Scatter plot of PPVs for each protein according to Cα and side-chain center contact identification methods. (*C*) Histogram of the ratios from the data in (B) to estimate the effect size. (*D*) and (*E*) As in (B) and (C), comparing Cβ and side-chain center contact identification methods.

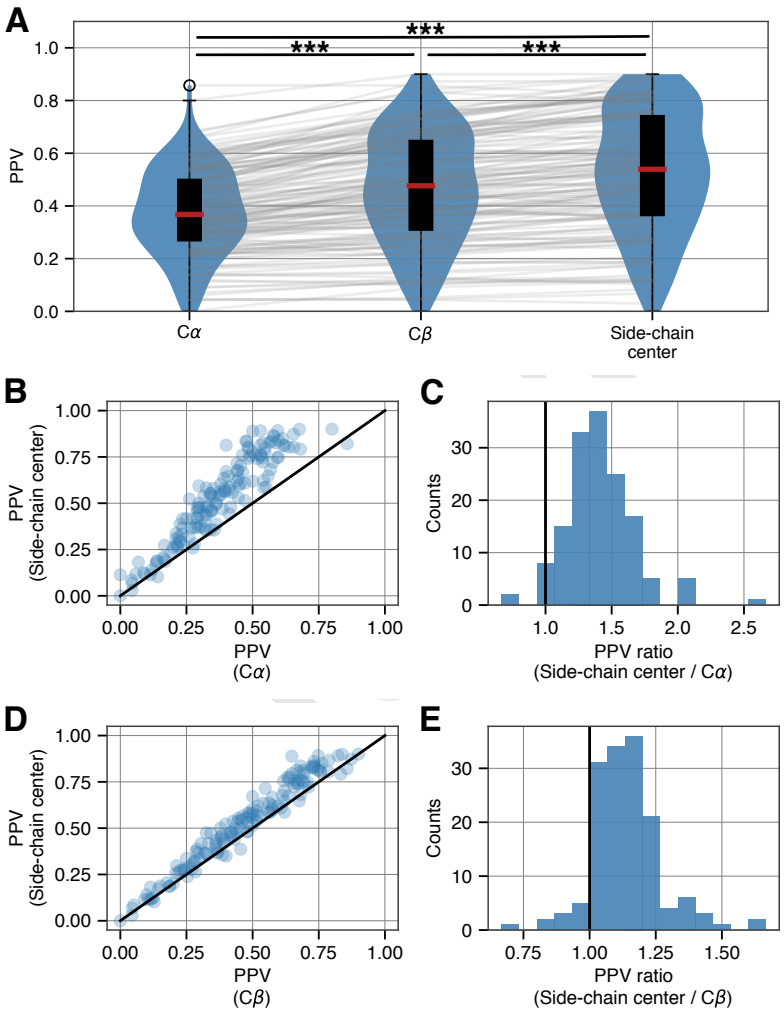

SI Fig. S6. Similar to Fig. 4 of the main text; here, we use the PLMC method to calculate evolutionary couplings. (*A*) For a diverse set of 150 proteins, the PPV of the top $L/2$ evolutionary coupling scores—derived from empirical sequence alignments—is progressively higher when intramolecular contacts are defined according to Cα atoms, Cβ atoms, and side-chain centers. (*** indicates $p < 10^{-20}$, Wilcoxon signed-rank test) (*B*) Scatter plot of PPVs for each protein according to Cα and side-chain center contact identification methods. (*C*) Histogram of the ratios from the data in (B) to estimate the effect size. (*D*) and (*E*) As in (B) and (C), comparing Cβ and side-chain center contact identification methods.

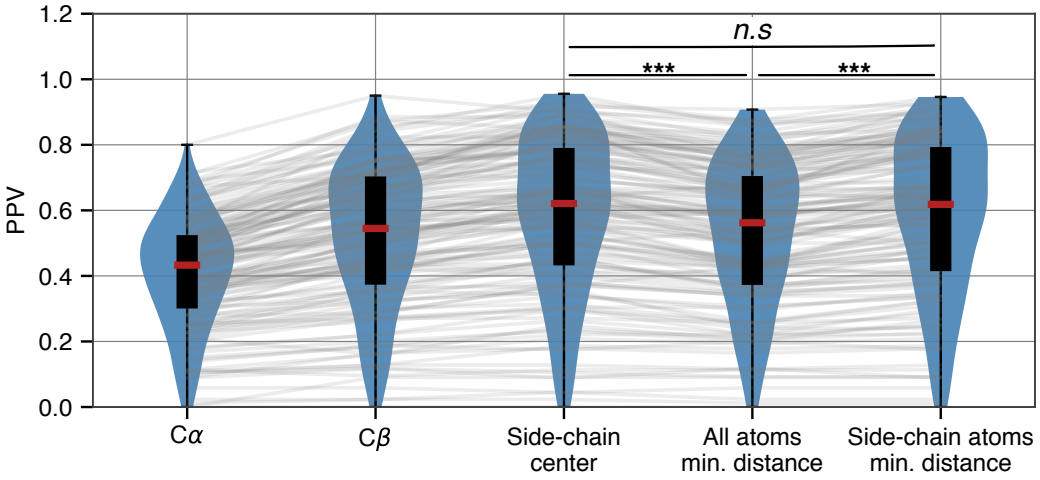

SI Fig. S7. Similar to Fig. 4A of the main text; here, we use the CCMpred method (as in the main text) to calculate evolutionary couplings but include two new methods for calculating residue–residue contacts. Contacts are defined based on the minimum distance between any two heavy atoms ('All atoms min. distance') or based on the minimum distance between any two heavy atoms that are part of the side chain ('Side-chain atoms min. distance'). Cα, Cβ and side-chain centers are included for comparison only and are the exact same data as displayed in Fig. 4A (we have ommited significance comparisons between these categories for clarity). (*** indicates $p < 10^{-20}$, Wilcoxon signed-rank test)