



The proteome folding project: Proteome-scale prediction of structure and function

Kevin Drew, Patrick Winters, Glenn L. Butterfoss, et al.

Genome Res. published online August 8, 2011

Access the most recent version at doi:[10.1101/gr.121475.111](https://doi.org/10.1101/gr.121475.111)

P<P	Published online August 8, 2011 in advance of the print journal.
Accepted Manuscript	Peer-reviewed and accepted for publication but not copyedited or typeset; accepted manuscript is likely to differ from the final, published version.
Open Access	Freely available online through the <i>Genome Research</i> Open Access option.
License	This manuscript is Open Access.
Email Alerting Service	Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or click here .



To subscribe to *Genome Research* go to:
<https://genome.cshlp.org/subscriptions>

Copyright © 2011, Cold Spring Harbor Laboratory Press

The proteome folding project: proteome-scale prediction of structure and function

Kevin Drew¹, Patrick Winters¹, Glenn L. Butterfoss¹, Viktors Berstis², Keith Uplinger², Jonathan Armstrong², Michael Riffle³, Erik Schweighofer⁴, Bill Bovermann², David R. Goodlett⁵, Trisha N. Davis³, Dennis Shasha⁶, Lars Malmström⁷,
Richard Bonneau^{1,4,6*}

¹ Center for Genomics and Systems Biology, Department of Biology, New York University, New York, NY 10003, USA

² IBM, Austin, TX, 78758, USA

³ Department of Biochemistry, Department of Genome Sciences, University of Washington, Seattle, WA 98195, USA

⁴ Institute for Systems Biology, Seattle, WA 98103, USA

⁵ Medicinal Chemistry Department, University of Washington, Seattle, WA 98195, USA

⁶ Department of Computer Science, Courant Institute of Mathematical Sciences, New York University, New York, New York, 10003, USA

⁷ Institute of Molecular Systems Biology, ETH Zurich, Zurich CH 8093, Switzerland

* To whom correspondence should be addressed.

Richard Bonneau, PhD.
New York University
Center for Genomics and Systems Biology
100 Washington Sq. E.
1009 Main Building
New York, NY 10003
Voice: 646 460 3026
Email: bonneau@nyu.edu

Running title: proteome folding

Keywords: proteome, *de novo* folding, function prediction, Rosetta, grid computing

Abstract:

The incompleteness of proteome structure and function annotation is a critical problem for biologists and, in particular, severely limits interpretation of high-throughput and next-generation experiments. We have developed a proteome annotation pipeline based on structure prediction, where function and structure annotations are generated using an integration of sequence comparison, fold recognition and grid-computing enabled *de novo* structure prediction. We predict protein domain boundaries and 3D structures for protein domains from 94 genomes (including Human, Arabidopsis, Rice, Mouse, Fly, Yeast, *E. coli* and Worm). *De novo* structure predictions were distributed on a grid of over 1.5 million CPUs worldwide (World Community Grid). We generate significant numbers of new confident fold annotations (9% of domains that are otherwise unannotated in these genomes). We demonstrate that predicted structures can be combined with annotations from the Gene Ontology database to predict new and more specific molecular functions. We provide multiple interfaces to this database at: <http://www.yeastrc.org/pdr/>.

1 Introduction

Annotation of protein structure and function is a fundamental challenge in biology. Accurate structure annotations give researchers a three dimensional view of protein function at the molecular level which enables specific point mutation analysis or the design of custom inhibitors to disrupt function. Accurate function annotations give researchers specific testable hypotheses about the role a protein plays in the cell and also allow biologists to better interpret the role of uncharacterized genes from high-throughput experiments (e.g. Mass Spectrometry co-IP, yeast two-hybrid screens, microarray, RNAi screens or forward-genetic screens). Unfortunately, experimental annotation efforts fail to cover large portions of all proteomes and are often focused on model organisms. Current computational function prediction methods can extend coverage to any sequenced genome (i.e. non-model organisms) but can only annotate proteins that have high sequence similarity to well characterized proteins. Motivated by the observation that structure is more conserved than sequence (Chothia and Lesk 1986), our method extends function annotation coverage to many unannotated protein domains by comparing computationally predicted structures to well characterized proteins with known structures.

Current computational protein annotation methods can be broadly grouped into four categories. (1) Primary sequence-feature annotation methods predict general features of a protein such as disorder content (DISOPRED (Jones and Ward 2003)), secondary structure (Pisipred (Jones 1999)), transmembrane helices (Krogh et al. 2001), coiled coils (COILS (LUPAS et al. 1991)) or signal peptides (SignalP (Bendtsen et al. 2004)) and can be efficiently applied to full proteomes but have limited ability to describe the specific function of individual proteins in a cellular context. (2) Sequence comparison methods, (e.g. Blast (Altschul et al. 1997)) including methods and databases that organize proteins into families (CATH (Orengo et al. 1997), SCOP (Murzin et al. 1995), PFAM (Finn et al. 2008)) and fold recognition methods (FFAS (Jaroszewski et al. 2005)), can generate putative function annotations (see (Lee et al. 2007) for a complete review) but are limited in their application to the set of proteins with sequence detectable homologues or significant sequence matches. (3) Several groups have used machine learning techniques to integrate high throughput experimental data, such as gene expression and protein-protein interactions, to predict protein function ((Lee et al. 2004) (Marcotte et al. 1999; Pena-Castillo et al. 2008) (Troyanskaya et al. 2003) (Bader and Hogue 2002; Hazbun et al. 2003)) but these datasets are not generally available for all organisms and therefore limit these methods' coverage. (4) Several studies have shown that protein structures derived from either large scale protein experimental determination (Dessailly et al. 2009; Matthews 2007) or large scale pro-

tein prediction (homology modeling, fold recognition and *de novo*) significantly increased the annotation coverage of proteomes and often provided site specific information such as probable functional sites and surfaces (Pieper et al. 2009) (Ginalski et al. 2004) (Bonneau et al. 2004; Zhang et al. 2009) (Malmström et al. 2007). Homology modeling is increasingly productive as more structures are solved, however, many proteins and protein domains lack detectable homologous proteins in the Protein Databank (PDB) (Marsden et al. 2007). *De novo* structure prediction methods do not require sequence homology to known structures and can, in principle, provide annotation coverage to proteins unreachable by homology modeling. Unfortunately *de novo* prediction methods require vast computational resources and because of this published pipelines that include *de novo* structure prediction are incapable of keeping pace with incoming genomic data.

This work focuses on this 4th type of annotation method via structure prediction. Here we describe the Proteome Folding Pipeline (PFP), a protein domain-level fold and function annotation method, which combines *de novo* (Rosetta (Rohl et al. 2004)), fold recognition, and sequence-based structure assignment methods into a single pipeline that significantly extends functional and structural proteome annotation coverage for 94 complete genomes (as well as several new protein families from recent metagenomics studies). The computational cost of performing *de novo* structure prediction was distributed on a grid of over 1.5 million CPUs worldwide (World Community Grid, wagrid.org). *De novo* structures and sequence-based comparisons were used to assign predicted protein domains into the Structural Classification of Proteins (SCOP), a hierarchical classification of protein three dimensional structure (Murzin et al. 1995). We demonstrate our ability to integrate these predicted SCOP classifications with information from the Gene Ontology (GO) database (Ashburner et al. 2000) to predict molecular functions, which increases the completeness, accuracy and specificity of protein domain annotations. We developed confidence levels and error models for *de novo* and SCOP classification methods based on a double-blind benchmark of our own construction containing 875 recently solved proteins and characterized the error and yield associated with each product of our pipeline (domains, structure, and function). We provide three interfaces to our database: a Cytoscape (Shannon et al. 2003) network interface that highlights novel structure and function predictions in the context of protein interaction networks (Avila-Campillo et al. 2007), a web interface that allows users to search for specific genes of interest; <http://www.yeastrc.org/pdr/> (Riffle et al. 2005) and a Blast interface that allows searches using individual sequences (<http://pfp.bio.nyu.edu/blast/index>).

2 Results

2.1 Proteome Folding Pipeline applied to 94 genomes

We have applied our pipeline to over 389,000 proteins from 94 genomes (figure 1 describes the full pipeline using a hypothetical protein from *Lactobacillus prophage*, table S1 details the full set of genomes analyzed). First, the domain prediction protocol Ginzu (Chivian et al. 2003; Chivian et al. 2005) uses primary sequence-based annotation methods to predict secondary structure (PSIPRED(Jones 1999)), disordered regions (DISOPRED(Jones and Ward 2003)), signal sequences (SignalP (Bendtsen et al. 2004)), coiled coils (COILS(LUPAS et al. 1991)) and transmembrane regions (TMHMM(Krogh et al. 2001)). PSI-Blast(Altschul et al. 1997) is then used to identify structures in the PDB with high sequence similarity to regions of the query protein, referred to as PDB-Blast hits. At this stage, over 314,000 domains from the set of input proteins were identified.

Next, we use the fold recognition algorithm, FFAS03(Jaroszewski et al. 2005), to match more evolutionarily distant sequences in the PDB, producing an additional 58,000 domains with significant matches to domains in the PDB. We then use additional methods, including identifying Pfam domains (Finn et al. 2008), an algorithm for predicting domains from multiple sequence alignments (MSA) and a Heuristic-based algorithm for delineating domain boundaries to identify additional putative domains (Chivian et al. 2003; Chivian et al. 2005). These methods combined to identify an additional 325,000 domains. In total, our domain prediction produced nearly 700,000 domains for the 389,000 query proteins, which serve as the basis for our domain centric annotation of these proteins. We then used the Rosetta *de novo* protocol(Rohl et al. 2004) to predict the three dimensional structure of those domains lacking structure annotation (i.e. Pfam, MSA and Heuristic domains) and that are less than 150 residues in length. We predicted *de novo* folds for 57,000 domains on IBM's World Community Grid (requiring over 100,000 years of CPU time, resulting in one of the largest repository of protein structure predictions publicly available). The final step in the pipeline classifies predicted protein domains into structural superfamilies (SCOP). PDB-Blast and FFAS03 domains are assigned the superfamily of the PDB structure to which they matched while a logistic regression model was used to classify *de novo* structures into SCOP superfamilies based on the Rosetta predictions and estimated error associated with these structure predictions (Malmström et al. 2007). In all, our pipeline classified over 250,000 domains into SCOP superfamilies of which nearly 43,000 are considered both confident (based on our benchmarks) and novel (FFAS03 and *de novo*).

2.2 Domain prediction

As described above, we first produced protein domain predictions for each protein in the 94 genomes processed. The domain prediction program, GInzu, allows us to hierarchically organize all structure prediction methods, ensuring that each domain has a structure annotation or prediction derived from the most accurate and most computationally efficient method possible. Figure S1 shows the domain types assigned by our pipeline for several representative organisms. PDB-Blast and FFAS03 annotate an average of 47% and 9%, respectively, of all proteomes (figure S1) and together assign a structure to 51% of eukaryotic domain sequences, agreeing very closely with previous efforts (Marsden et al. 2007). Our predictions for the human proteome have a slightly higher coverage of domain sequences, by these two methods (62%), while many pathogenic eukaryotic organisms are sparsely annotated and are observed to have lower PDB-Blast and FFAS03 coverage, for example *Trypanosoma cruzi* (41%) and *Plasmodium vivax* (37%). Domain coverage across all 94 genomes (figure S1 and table S1) confirms that our analysis provides broad coverage of structure annotations even before *de novo* methods are applied and that our strategy for predicting domain boundaries is extensible to a wide variety of organisms. Additionally, our pipeline produced 29,000 Pfam, 105,000 MSA, and 190,000 Heuristic domains totaling 324,000 predicted protein domains annotated by these three methods (which do not assign structures to domains). It is important to note that in our pipeline, Pfam annotates a smaller than expected number of domains because a large number of domains are removed from consideration by PDB-Blast and FFAS before Pfam is applied. This is consistent with our effort throughout the pipeline to first assign a high confident structure to individual domains before proceeding to lower confidence structure methods (i.e. *de novo*).

2.3 *De novo* Structure Prediction and SCOP superfamily classification of structure predictions

We used the Rosetta *de novo* protocol to predict 3D models for domains not annotated by PDB-Blast and FFAS. Resulting structure predictions were then used to predict SCOP superfamilies for each domain. Rosetta is a *de novo* protein structure prediction algorithm that uses fragments of proteins from the PDB to generate local structure during a Monte Carlo optimization that then produces ensembles of low energy protein conformations. The Rosetta *de novo* protocol does not require homology to a solved structure in the PDB and is thus theoretically applicable to all protein domains. In practice, however, the algorithm does not scale to protein domains larger than 150 residues (roughly half of protein domains are too large for Rosetta). Rosetta *de novo* has been very successful in CASP competitions and often predicts structures for protein domains

lacking detectable homology to known structures to within 5 Å RMSD of the native structure or better (Bradley et al. 2005). Structure fragment libraries are produced by comparing the sequence and secondary structure (predicted using PSIPRED) of 3 and 9 amino acid windows of query sequences to a non-redundant set of PDB entries. Because of the computational cost of running Rosetta on the genome-wide scale, we initiated the Human Proteome Folding Project in collaboration with IBM's World Community Grid (WCGrid). WCGrid makes use of idle CPU time on volunteered personal computers around the world to form a very large virtual supercomputer. As of September 25th, 2010, there are over a half million volunteers and nearly 1.6 million devices capable of running Rosetta on WCGrid. Fragment libraries and query sequences were submitted to WCGrid where the Rosetta *de novo* protocol was executed resulting in structure predictions for 57,000 domains between 40 and 150 residues in length (see methods). The resulting ensemble of structures predicted for each domain was clustered using root mean squared distance (RMSD) and five of the top *de novo* predictions (cluster centers) are publicly provided through our web interface where they can be viewed directly in a web browser or downloaded for further analysis.

We have used our *de novo* structure predictions to predict the SCOP superfamily for each predicted domain generated as described above. The SCOP database describes structural and evolutionary relationships of proteins with known structure (the CATH database would also have functioned as an appropriate fold ontology for this study). SCOP organizes protein structure by a set of hierarchical classes, folds, superfamilies and families. We make predictions at the SCOP superfamily level (e.g. a.4.5, "Winged helix" DNA-binding domain), as proteins within the same superfamily have high structural similarity and often share functional features.

For each protein domain, top ranked Rosetta predictions were compared to a representative set of known structures, which are classified in SCOP. Structure comparisons were made with MAMMOTH (Ortiz et al. 2002), which returns the statistical significance (z-score) of the best gapped alignment between two protein structures. A logistic regression model was then used to estimate the probability (MAMMOTH confidence metric, MCM score (Malmström et al. 2007)) of each protein domain belonging to a SCOP superfamily. The regression model's parameters consist of MAMMOTH z-score, Rosetta convergence score, contact order of the predicted structure and a sequence length ratio between the SCOP superfamily representative and the protein domain sequence. We demonstrate, below and in previous work (Malmström et al. 2007), that our MCM score separates true predictions from incorrect conformations and allows us to classify proteins for which no valid predictions were generated by other methods.

Table 1 shows the number of SCOP classifications that were made for all protein domains processed by our pipeline. Over 57,000 domains were *de novo* folded and classified into a SCOP superfamily, of which 12,500 (21.8%) are considered medium confident (MCM score ≥ 0.8 , ~2 out of 3 correct, see section 2.3.1) and 4,500 (7.9%) are considered high confidence (MCM score ≥ 0.9 , ~4 out of 5 correct, see section 2.3.1). In addition, we also assigned SCOP superfamilies based on sequence similarity for domains identified by PDB-Blast and FFAS03, which were annotated using the alignment to the matching PDB entry's SCOP classification (if classified). Using these methods to assign SCOP superfamilies has been shown to be successful while introducing very few false positives (Rychlewski et al. 2000). Table 1 shows the number of PDB-Blast and FFAS03 domains classified by the PFP with SCOP superfamilies. Of the 314,000 PDB-Blast domains, 207,700 (66%) were annotated with a SCOP superfamily and of the 58,000 FFAS03 domains, 30,000 (52%) were annotated with a SCOP superfamily. In total our pipeline assigned or predicted SCOP superfamilies for over 295,000 domains out of the nearly 700,000 predicted domains.

2.3.1 Validation of *De novo* Superfamily Predictions

In this section we assess the accuracy of our SCOP superfamily classifier using a benchmark of 875 proteins solved after our Rosetta predictions were made (table S5). Previous test sets used for determining accuracy of the MCM score function were composed of predictions made on proteins with structures already in the PDB (Malmström et al. 2007). The use of a test set based only on the PDB is not ideal because it does not reflect the error associated with domain prediction, can overestimate the performance (as the PDB is used to create fragment library), and can only account for folds currently present in the SCOP database. To address these issues we compiled a set of 875 *de novo* structure predictions, from the full set of 57,000 *de novo* predictions, whose structures (or highly similar protein, Blast e-value of 10^{-5} or less, see methods) were experimentally solved after Rosetta fragment library selection; we refer to this set of 875 proteins as the Solved After Predicted set (SAP)(table S5). The top 25 cluster centers for each sequence in the SAP set were first classified into SCOP superfamilies (described above) and then the superfamily with the highest MCM score was compared to the true superfamily of the native structure. We correctly classified 47% (407) of the domain structures out of the 875 in the SAP set (table 2) with our top ranked prediction. Figure 2 (top and middle) shows correct vs incorrect predictions broken down by SCOP class. Figure 2 (bottom left) shows a precision/yield curve for the SAP set where predictions are ordered by their MCM score and plotted according to their precision. An MCM score threshold of 0.9 yields 38% (333) of all predictions made in the SAP

set and is 78% accurate (figure 2 bottom left, table 2). Our overall performance on the SAP set demonstrates our ability to accurately classify structure predictions into superfamilies. Additional analysis on the validation of this benchmark can be found in the supplemental material (section S2.2).

2.3.2 Quality of superfamily predictions and underlying structure predictions

In several cases, Rosetta produced relatively accurate structure models but we failed to assign the correct SCOP superfamily because the protein fold was not yet annotated in SCOP or was sparsely sampled. In this section, we characterize sources of error that affect our ability to accurately classify proteins into SCOP superfamilies. To properly account for a correct superfamily not being classified in the SCOP database, we ran our SAP set on an earlier version of SCOP (v1.67), which was released prior to our *de novo* predictions. Our analysis of incorrectly classified SCOP superfamilies in the SAP set reveals three main distinct sources of error: (1) errors due to the absence of the target fold in SCOP, (2) errors due to inaccuracy of predicted structures and (3) errors due to incorrect classification of an accurate structure.

The first type of error is due to the true fold/superfamily not being represented in our structure comparison set. There are 554 new superfamilies in the current version of SCOP (v1.75) that were absent in the previous version (1.67). Figure 2 (bottom right panel) shows a histogram of the incorrectly classified models in the SAP set. Shaded in yellow is the portion of incorrect classifications where the true superfamily is a new superfamily. New superfamilies are the major source of error in our classification method and accounts for roughly 62% of the total error in the SAP set. We believe this error type will become less significant as the protein fold space is more thoroughly sampled and more superfamily representatives are included in our structure comparison set. It should be noted that this type of error is absent from previously reported benchmarks based solely on the PDB because the test structures are already classified in the SCOP database.

The next source of error stems from error in the *de novo* predicted model. We estimate this by determining if the predicted structure is more similar (by MAMMOTH structure alignment) to the incorrect superfamily than the true superfamily (figure 2 bottom right panel, shaded gray) (see methods). Structures that are more similar to the incorrect superfamily are classified incorrectly due to insufficient quality of 3D Rosetta models and this error is estimated to constitute around 28% of the total error. Further developments to the Rosetta *de novo* protocol and increased sampling may alleviate this error in the future.

The final source of error stems from error in our superfamily classifier method. This error arises when Rosetta produces an accurate model but our classifier method chooses an incorrect superfamily in spite of the fact that a structure with the correct superfamily exists in our database (figure 2 bottom right panel, shaded green) (see methods). Error attributed to the classifier method is minimal compared to other error types, approximately 10% of total error. Interestingly, several of the models in this error class have the true superfamily "alpha-beta sandwich, ACT-like" (d.58.18). A recent study describing the CATH hierarchy of structural classification (a classification of 3D protein structures similar to SCOP) describes this fold (CATH Architecture 3.30) to be in a continuous, densely populated region of protein fold space where several independent folds have similar features (Cuff et al. 2009). Our superfamily classifier depends on structural differences, so distinguishing between two closely related folds is more difficult and, therefore, requires the highest level of structural accuracy and increased sampling of fold space. We believe continuous regions of fold space will be a persistent, albeit small, source of error for our method and similar structure-based methods.

In line with our estimates of the relative importance of these three main sources of error, we find that our fold-prediction yield on the SAP set more than doubles when we update our fold database from SCOP release 1.67 to 1.75 (figure S2).

We next asked if the *de novo* structure models are valuable alone, independent of SCOP superfamily classification. We used MAMMOTH to compare each of the 875 *de novo* predictions in the SAP set to their experimental structure. Figure S3 shows our superfamily classifier scores are well correlated with model similarity via the MAMMOTH *z*-score which accounts for length and RMSD of structure similarity to the true structure. This shows that the superfamily classifier score is predictive of not only superfamily classification accuracy but also model accuracy, allowing *de novo* models to be used in other applications that do not require superfamily classification such as active site localization or prediction of residue surface accessibility (see figure S4).

2.4 Molecular Function Prediction

We next demonstrate that SCOP superfamily predictions can be integrated with additional information annotated to a protein to predict Gene Ontology (GO) molecular functions in an automated way. GO is a controlled vocabulary of molecular function (GO-MF), biological process (GO-P) and cellular component (GO-C) terms suitable for automated transfer among proteins. We focus on the integration of structural information with GO-P and GO-C due to the wide availability of these predictors across many of the genomes analyzed. Many structural superfamilies exhibit a diverse set of compatible GO-MF annotations and additional evidence is needed to

determine a specific function for an individual protein domain. For example, the SCOP immunoglobulin superfamily (sccs: b.1.1) is annotated with several function terms including protein binding (pdbid: 3D2U), transporter activity (pdbid: 2ZJS) and structural molecule activity (pdbid: 1ACY), among others. The function of an immunoglobulin protein can only be understood in the context of its localization and interaction partners in the cell (information that can be found in databases or extracted from high throughput experiments for several organisms and described by GO-P or GO-C). Conversely, structural evidence may allow for the refinement of a general function annotation (e.g. binding) inferred by localization or process evidence to a more specific function annotation (e.g. miRNA binding). To this end, we have developed a naïve Bayes classifier that integrates GO-P and GO-C annotations with predicted SCOP superfamily classifications to predict GO-MF terms.

For each GO-MF term in the GO database, our method first calculates the likelihood the MF term is true (LT) and the likelihood the term is false (LF) for an individual predictor (i.e. GO-P, GO-C, superfamily classifications). A ratio of LT over LF is then calculated and the log of the ratio is taken producing an individual log likelihood ratio (LLR) score. The individual LLR is calculated for each independent predictor classified to the query protein. The individual LLR score for the structure (S) predictor is scaled to reflect the confidence we have in the structure assignment (ex. *de novo* LLR is scaled by MCM score). Individual likelihoods are calculated using frequency counts of training sequences annotated with GO-MF and which have GO-P annotations, GO-C annotations or high quality Blast matches to SCOP superfamily representatives. Due to the conditional dependence of GO-P and GO-C terms within their respective branches of the GO tree, a feature selection method was implemented to restrict GO-P and GO-C to a focused conditionally independent set. We selected one GO-P and one GO-C (when available) with the highest mutual information with each potential GO-MF term for the calculation of an individual LLR score. Finally, all individual LLR scores for the GO-MF term are summed to produce an overall LLR score. See methods section for a complete description. An overall LLR score greater than zero can be interpreted as being more likely to be true than false and we therefore consider an $LLR > 0$ to be confident.

2.4.1 Yield of function predictions using superfamily classifications

We applied our function prediction method to 295,000 domains with PDB-Blast, FFAS03 and *de novo* superfamily classifications. We confidently predict specific GO-MF for 44% (129,000) (table S4) of domains ($LLR > 0$: prediction is more likely true than false and GO-MF is annotated to $< 2\%$ of proteins in our training set). We confidently predict novel GO-MF annotations

for 15% (22,000) of the 147,000 domains that lack any GO-MF (table S4) and therefore the addition of structural evidence derived from our predictions significantly increases the coverage of function annotation to unannotated proteins and protein domains. In addition our function predictions are applied directly to domains, which begins to address the need for site specific annotation in multi-domain proteins.

The value of function predictions is seen not only for unannotated domains but also for "under-annotated" domains, which are domains that have either a general function (but not one specific enough to fully characterize the domain) or domains that have multiple functions (but are currently only annotated with one). We confidently predict specific functions for 87,000 domains for which we either extended a generic annotation or predicted a new function altogether. In all, we provide confident function predictions for 109,000 under- and unannotated domains, which makes up a significant fraction (37%) of domains for which structure predictions were generated.

2.4.2 Accuracy of function prediction method

To determine the accuracy of our function prediction method, we created a benchmark using sequences processed with our pipeline that are currently annotated with one or more specific GO-MF term (see methods). We first examine the accuracy of our function predictions made using PDB-Blast structure evidence. PDB-Blast is the highest confidence structure assignment method and provides an upper bound on the confidence and expected yield of function predictions. Figure 3 (upper left) shows that integrating PDB-Blast structural evidence with GO Process (GO-P) and Component (GO-C) (PCS, green) improves prediction accuracy over only using GO-P and GO-C (PC, red) for a random sampling of 5,000 GO-MF annotated eukaryotic proteins. All combinations of predictors involving structural evidence show improved performance, and the combination of PCS performs best with over 20% recall of known molecular functions at 50% precision. This result confirms that molecular function predictions are more accurate using structural evidence at the level of SCOP superfamily.

Performance of function predictions using FFAS03 structure evidence were also evaluated using a random sampling of 5000 eukaryotic proteins. Figure 3 (upper right) highlights the integrative performance of FFAS03 structural evidence with other GO information. This shows the inclusion of fold recognition structural evidence with GO-P and GO-C (green) mostly outperforms GO-P and GO-C alone (red). The decrease in performance of fold recognition evidence in comparison to PDB-Blast (7.5% recall at 50% precision) can be attributed to the error associated with the longer evolutionary distance between query and matched proteins, causing greater functional divergence and a result of incomplete annotation of benchmark domains. The increase in

accuracy when integrating fold recognition structure evidence in comparison to GO-P and GO-C evidence alone, shows structure evidence at multiple levels of confidence can be used to predict molecular function.

Creating a proper benchmark to address the accuracy of function predictions based on *de novo* structure evidence is generally a more difficult task than creating a benchmark to judge structural accuracy of predictions. Because our pipeline annotates domains with the highest confidence method first, domains with *de novo* structure predictions are less likely to be annotated with GO functions shrinking the pool of domains from which we can draw our benchmark. Therefore benchmarking *de novo* function predictions is complicated by the limited ability to distinguish between a false prediction and a true prediction that is not currently annotated. In this work *de novo* function predictions from a subset of the SAP set with known GO-MF annotations were used as a blind benchmark, and show *de novo* structure evidence can be used in combination with GO-P and GO-C to predict molecular function (figure S7). It is important to keep in mind, however, the SAP set is an imperfect benchmark because of its relatively small size with respect to fold recognition based function predictions. Nevertheless, we are encouraged by this result. We anticipate that as additional function and structure annotations are deposited in GO and the PDB, respectively, we will be able to expand our SAP set and better calibrate our error estimates for *de novo*-based function prediction.

2.4.3 Structural information produces unique function predictions and increases the specificity of function predictions derived from other methods.

An examination of the molecular function predictions made by our pipeline shows that structural evidence not only improves our predictive accuracy (see section above) but also allows us to make predictions for a novel set of function terms and a novel set of proteins which are not made using the GO-P and GO-C evidence alone. We first made predictions for a random sampling of FFAS03 annotated eukaryotic domains using two evidence sets, one with structure information, PCS (GO-P, GO-C and SCOP superfamily), and one without structure information, PC (GO-P, GO-C). We then determined all of the unique MF terms predicted by both evidence sets and divided them into sets of functions that were predicted 1) only by PCS, 2) only by PC and 3) by both which resulted in 636 unique function terms (precision > 50%) (figure 3 lower left). Of these 636 function terms, the PCS evidence set (includes structure) predicted 306 unique function terms. This is in comparison to only 29 function terms that were uniquely predicted by the PC evidence set (the remaining 301 unique function terms were predicted by both), which sug-

gests the addition of structural information as a function predictor allows for the prediction of a greater number of function terms.

We also asked what set of proteins required structure evidence for confident function annotation. Figure 3 (lower right) shows 21% (359) of the sample set of proteins are confidently annotated by PCS compared to 10% (183) only annotated by PC (a remaining 1201 proteins were annotated by both evidence sets). These results demonstrate that function predictions incorporating structural evidence aid in annotating both an expanded set of function terms and protein domains. This, taken with the above results, shows that our function predictions accurately annotate substantial portions of proteomes that are unreachable by other methods.

2.5 Example predictions

2.5.1 Expansion of the transglutaminase fold family in *D. radiodurans* represents a key adaptation to ionizing radiation.

D. radiodurans can withstand extremely large doses of ionizing irradiation (Cox and Battista 2005). A previous study of the *D. radiodurans* genome has shown enrichment of specific protein families related to stress response and damage control based on sequence comparisons to other bacterial organisms (Makarova et al. 2001). We performed fold superfamily enrichment analysis based on our proteome wide predicted superfamilies for *D. radiodurans* (4,864 protein domains from the PFP pipeline) (table 3a). Our analysis recovered many of the previous study's enriched protein families including PR-1-like, Subtilisin-like, Nudix hydrolases and DinB/YfiT-like folds. Our work expands structure prediction coverage and reveals several enriched protein folds not reported in these prior comparisons, in particular, we predict the *D. radiodurans* proteome contains 10 or more protein domains with the transglutaminase fold. The transglutaminase fold has been shown to participate in the nucleotide excision repair (NER) pathway (the Yeast protein RAD4 is homologous to the several members of the transglutaminase family) (Anantharaman et al. 2001). The uncharacterized *D. radiodurans* gene, DR1901, predicted by this study to have a transglutaminase fold, exhibits a 10 fold induction in gene expression in response to radiation (Liu et al. 2003) supporting the hypothesis that several of the unannotated proteins we predict to have a transglutaminase fold are directly involved in *D. radiodurans* response to ionizing radiation, likely by participating in the NER pathway in a manner similar to RAD4.

2.5.2 Novel structure predictions reveal an enrichment of several new signaling and virulence related fold families in the *P. vivax* proteome.

P. vivax is a human malaria parasite, which has recently shown resistance to common drug treatments (Rieckmann et al. 1989). We used our structure predictions to uncover specific folds that are significantly expanded in this pathogen (table 3b). Two protein families commonly found in *Plasmodium* genomes, and already known to play a key role in pathogenesis, were recovered: major surface antigen and Duffy binding domain-like. Newly detected enriched folds include the FKBP12-rapamycin-binding (5 instances of this fold were predicted by the Rosetta *de novo* protocol). Rapamycin is an immunosuppressant known to interact with several proteins including mTOR and FKBP12 (Zhou et al.); these domains may be involved in the parasite's interactions with the host immune system during initial stages of infection. Other enriched folds include virulence factors such as the Adhesin YadA fold and the protease cathepsin, which make potential drug targets for further study. Our predictions represent interesting candidate members of the already complex host-pathogen interaction network currently characterized for *P. vivax*.

2.5.4 CC_3056 (*Caulobacter crescentus*) is a truncated hemoglobin fold

The truncated hemoglobin family is a 2/2 α helical fold present in bacteria with functions including NO dioxygenation, oxidation/reduction and respiration (Vinogradov and Moens 2008). The *Caulobacter crescentus* hypothetical protein, CC_3056, was folded using Rosetta (figure 4a) and confidently classified in the "Globin-like" superfamily (scs:a.1.1, MCM = .98); our prediction matched a truncated hemoglobin (1DLWa), producing a structure-structure alignment with 4.97Å RMSD (figure 4b). CC_3056 and the truncated hemoglobin, 1DLWa, have only 23% sequence identity but share a semi-conserved heme binding pocket (figure 4c) where 16 out of 26 ligand binding residues are similar (Blossum62 values > 0), including nine identical residues key to the function of 1DLWa (Lopez et al. 2007). Finally, a close homologue of CC_3056, *C. jejuni* trHbP, was recently crystallized (after our initial prediction was made, 2IG3) confirming the truncated hemoglobin fold prediction (Nardini et al. 2006).

(<http://www.yeastrc.org/pdr/viewProtein.do?id=2823436>, the PDB-Blast hit to 2IG3 is now reported in the updated database)

2.5.5 GOS_6366033 (a new and unannotated protein family from ocean metagenomics) is predicted to be a chorismate mutase

The recent metagenomics Global Ocean Sampling (GOS) expedition (Yooseph et al. 2007) provided more than 1700 new protein families with no detectable homology to known protein families. We folded 2 members for each these novel families containing fewer than 150 residues with Rosetta (877 families). The predicted fold of GOS_6366033 (figure 4d) confidently

matched a chorismate mutase structure (1ECM, scs:a.130.1 Chorismate mutase II, MCM=.93) (LEE et al. 1995) (figure 4e). Chorismate mutase catalyzes chorismate to prephenate in the bacterial biosynthesis pathway of tyrosine and phenylalanine. 1ECM and GOS_6366033 share only 11% sequence identity (based on structural alignment, as sequence-based methods detect no alignment) but a structure-structure alignment of our model to 1ECM predicts that five GOS_6366033 residues are conserved in the chorismate binding pocket of 1ECM (figure 4f). Two of the conserved residues, Arg51 (Arg44 in GOS_6366033) and Glu52 (Glu45), are identical and participate in hydrogen bonds with the transition state analog of chorismate in 1ECM, which possibly stabilize the ligand. Two other conserved residues Leu55 (Met48) and Ile81 (Leu74) are thought to aid binding through hydrophobic contacts with the ligand. A fifth conserved residue, Arg47 (Lys39), stabilizes the Glu52 sidechain through electrostatic interactions in the 1ECM structure leading us to believe the Lys39 in GOS_6366033 protein serves a similar role.

2.5.6 Rumi (*Drosophila melanogaster*)

Rumi is an endoplasmic reticulum protein and an important regulator of the Notch-signaling pathway, which in turn regulates cell-fate decisions. Our pipeline predicted Rumi to be a member of the glycosyltransferase SCOP superfamily (scs: c.87.1) based on fold recognition analysis (FFAS03). Integrating this superfamily prediction with the GO-P term "carbohydrate metabolic process" (GO:0005975) we predict the GO-MF term "transferase activity, transferring glycosyl groups" (GO:0016757). Subsequent to this prediction Rumi was shown to display O-glycosyltransferase activity by Acar, M., et al., who observed lower amounts of O-glycosylated peptides in Rumi knockdown samples compared to controls (Acar et al. 2008) confirming our function prediction for Rumi.

(<http://www.yeastrc.org/pdr/viewPSPOverview.do?id=673166>)

2.5.7 OOep/MOEP19 (*Mus musculus*)

Transfer of genetic material from maternal cells to precise locations in oocytes is important for proper mouse embryogenesis and development. The previously unannotated mouse gene, MOEP19, was predicted by FFAS03 to be a member of the Eukaryotic type KH-domain SCOP superfamily (scs: d.51.1). The KH-domain is a diverse RNA-binding domain superfamily. Our function prediction method using the assigned superfamily, GO-P "cellular macromolecule metabolic process" (GO:0034960) and GO-C "intracellular part" (GO:0044424), confidently predicted MOEP19 to have a "nucleic acid binding" GO-MF annotation (GO:0003676). It was later

shown experimentally by Herr et al. that MOEP19 binds ribonucleotide homopolymers using a RNA binding assay (Herr et al. 2008). The study also showed that MOEP19 is a maternal effect gene involved in blastomere polarity suggesting its involvement in patterning RNA in the developing oocyte. (<http://www.yeastrc.org/pdr/viewPSPOverview.do?id=614967>)

2.6 Interfaces to our proteome-wide structure and function predictions

Biologists can obtain structure, superfamily, and molecular function predictions along with other predictions (such as secondary structure, disordered regions, and multiple sequence alignments) via a standard web interface (<http://www.yeastrc.org/pdr/>, figure S5) (Riffle et al. 2005). A Blast interface (<http://pfp.bio.nyu.edu/blast/index>) is available for searching our database using sequences not in the 94 genomes processed (as well as variants of the sequences processed). Additionally, we have integrated our database with the Cytoscape protein-protein interaction plug-in BioNetBuilder (Avila-Campillo et al. 2007) so that our predictions can be explored in a network context. BioNetBuilder automatically builds gene association networks for any organism in the NCBI taxonomy based on 9 publicly available interaction databases (<http://err.bio.nyu.edu/cytoscape/bionetbuilder/>). Nodes in BioNetBuilder display visual cues (e.g. size, shape) indicating the confidence or relevance of our predictions for each gene in the network. Links to the web database are provided as attributes for each node in the Cytoscape network. Raw data is available upon request.

3 Discussion

We now discuss how the PFP pipeline relates to complementary methods, the need for improved gene prediction and domain boundary prediction as well as the importance of domain-centric function annotation.

3.1 Comparison of PFP to complementary methods and complementary experimental efforts

The Robetta server (Chivian et al. 2005) is a domain centric structure prediction pipeline and, like the PFP, attempts to run the most accurate applicable structure prediction method on each predicted domain (PDB-Blast → Fold recognition → Rosetta). The main disadvantages of this server are extremely long wait times (in many cases > 6 months), the inability to upload whole genomes, the lack of integrated function prediction, and the much smaller amount of sampling used for the *de novo* portion of the structure prediction. In the case of genome annotation, the

most critical of these limitations is the long wait times (due to the computational cost of Rosetta modeling), which makes full-proteome analysis impossible, given the current architecture. Our pipeline solves this problem by pre-computing and archiving *de novo* structure predictions run on WCGrid, allowing immediate access to results and full proteome database queries. For example, a list of all putative transcription factors in a genome is a required input into regulatory network inference algorithms. This and previous work (Bonneau et al. 2004; Bonneau et al. 2007) show that proteome-wide structure prediction significantly expands the list of putative regulators and signaling proteins. Other key advantages to the PFP as a full proteome annotation tool include SCOP superfamily classification of structures and GO function predictions based on those structure annotations. Integrating the Robetta server with the PFP database (such that queries to the Robetta server already contained in our database can be resolved without repeating costly structure prediction calculations) is an attractive area for future work.

The Protein Structure Initiative (PSI) is a multi-group effort to increase the coverage of structure annotations across protein sequence-space using both experimental and computational approaches (Dessailly et al. 2009). The computational effort we describe here and the PSI are extremely complimentary in both methodology and coverage of protein sequence-space. The PSI has greatly increased the number of structures in the Protein Data Bank and thus expanded the mapping of sequence-space to structures. This improves the coverage of PFP sequence and fold recognition-based methods (PDB-Blast, FFAS03) and indirectly improves the accuracy and yield of our *de novo* procedure by adding to the set of structures from which to classify *de novo* predictions (SCOP superfamily classification). The PSI has increased the fraction of protein sequences that can be assigned a structure by 2.0% over a 3 year period (based on the UniProt protein set(Dessailly et al. 2009)). In comparison, we report an increase of 1.8% in structural coverage of 94 genomes based on confident *de novo* structure predictions over all domains predicted (table 1). Although computationally predicted structures are less valuable for many tasks, the additional coverage afforded by our pipeline is mostly orthogonal to the PSI coverage. The PSI and PFP also differ in focus; eukaryotic organisms are not the primary focus of the PSI whereas the PFP has been applied to a balanced distribution of eukaryotic and prokaryotic organisms (table S1). Finally, *de novo* structure predictions from the PFP could be used to aid in target selection by the PSI. A current problem in selecting sequences for experimental structure determination by the PSI is discriminating sequences with novel folds from evolutionary divergent sequences of known folds. Our pipeline reports SCOP superfamily classifications for *de novo* structures, which are potential evolutionary divergent members of the matched superfamily. Sequences

without quality SCOP superfamily classifications are possible members of novel folds and thus could be assigned a higher priority for experimental structure determination with respect to increasing structural coverage of protein space.

3.2 Need for improved gene prediction and domain boundary prediction

Our ability to use structure-based methods to improve proteome annotation is sensitive to the accuracy of gene prediction and domain prediction methods. It has been reported that large numbers of genes have incorrectly assigned start sites (Nielsen and Krogh 2005) and splice sites (Guigó et al. 2006). We believe that as gene prediction techniques continue to improve the identification of protein boundaries, the PFP will show increased accuracy in predicting protein structure annotations for a greater number of proteins. An analysis of the performance of GINZU (Kim et al. 2005) showed that errors in domain boundary prediction are major sources of error in all downstream predictions. Other domain parsing programs such as CHOP (Liu and Rost 2004), could also be used in conjunction with GINZU to predict domain boundaries. Encouragingly, there are several promising recent high throughput experimental approaches to determining protein domains, which could feasibly alleviate significant portions of domain prediction error. For example, Boxem et al. used a modular domain view of protein-protein interactions to determine the minimal region of the protein necessary to maintain a protein-protein interaction (Boxem et al. 2008), thus elucidating individual domains. In many cases the minimal interacting domains detected by Boxem et al. corresponded well to known or predicted structural domains. Similar approaches may be developed in the future to probe domain boundaries in a scalable genome-wide fashion.

3.3 Domain-centric annotation

In general, different domains within the same protein carry out different functions. Unfortunately, many of the function annotations in the Gene Ontology database are not mapped to a specific domain in a multi-domain protein. Applications such as genome wide protein evolution and co-evolution studies are hampered by the lack of domain specific annotations. Chothia et al. have used domain structures of proteins to study gene duplication, recombination, and divergence and they quantitate these processes in terms of their evolutionary impact (Chothia et al. 2003). Vogel et al. describe the co-evolution of domain combinations, called supra-domains, that continually reoccur in proteins (Vogel et al. 2004). They note that over one third of structurally characterized proteins contain a supra-domain and therefore are particularly useful for genome evo-

lution and annotation efforts. Such evolutionary studies require domain level annotation and could be expanded if improved domain predictions were available.

We demonstrate the ability to accurately and efficiently predict protein domains and their three-dimensional structures on a proteome scale, classify unknown proteins into structural superfamilies, and predict functions based on this structural information. We have provided annotations for 94 proteomes, including medically relevant genomes, model organisms and the Human proteome all of which will be valuable to the computational biology, biology and clinical communities.

4 Methods

Previously published methods used in this work are described fully in the supplemental methods section, specifically domain prediction, domain selection for *de novo* prediction, *de novo* structure prediction and SCOP classification of structure predictions. Below are new methods specific to this work.

4.1 Genome Selection

Genomes were chosen to be representative of genomes across the tree of life as well as their value to the broader biological community (see table S1 for complete list). Most major model organisms are covered as well as a broad sampling of bacteria and archaea. Protein FASTA formatted files were downloaded from NCBI or the organism's genome web database circa December 2004 (HPF1 protocol) or November 2007/February 2008 (HPF2 protocol). For the GOS dataset, representatives from only novel protein families as described by Yooseph et al. (Yooseph et al. 2007) were analyzed with the PFP.

4.2 World Community Grid

Rosetta was run on IBM's World Community Grid (WCGrid) using the Boinc framework (Anderson 2004). Owners of PC's running Windows, MAC OS/X, or Linux participate in WCGrid by installing a secure grid client program available from <http://www.worldcommunitygrid.org>. The grid client program requests work units from WCGrid servers and runs Rosetta in the background at lowest possible priority. The state of the computation is checkpointed every few minutes to allow for efficient recovery of calculations in the event of machine failure. The Rosetta program was modified for use in the grid environment, which included adding check pointing and modifications to remove security vulnerabilities.

4.3 Enrichment Analysis

Enrichment scores were determined using the following equation:

$$E_{sf} = \log\left(\frac{n/N}{(b+1)/B}\right) \quad (1)$$

where n is the number of domains with superfamily sf , N is the total number of domains in the organism, b is the number of domains with superfamily sf in background organisms and B is the total number of domains in all background organisms. Background organisms used for *D. radiodurans* analysis were *E. coli*, *M. tuberculosis*, *C. crescentus* and *B. subtilis*. Background organisms used for *P. vivax* analysis were *H. sapiens*, *D. melanogaster*, *C. elegans*, *M. musculus*, *S. cerevisiae* and *A. thaliana*.

4.4 Binding Pocket Conservation Analysis

Example predictions CC_3056 (*Caulobacter crescentus*) and GOS_6366033 (Ocean Metagenomics) were produced as described above using the PFP protocol. Known ligand binding and catalytic residues of the matched superfamily were determined from the FireDB database (Lopez et al. 2007). A structural alignment between the predicted structure and the superfamily representative structure was then obtained using Mammoth. This alignment was used to calculate the Blossum62 conservation score for the ligand binding residues.

4.5 Molecular Function Prediction

Molecular function predictions were made using a naïve Bayes method that estimates a log-likelihood ratio (LLR) based on available evidence for each molecular function term, mf , in the GO ontology. Given the naïve Bayes assumption that available predictors, GO-P (p), GO-C (c) or SCOP classification (s) are conditionally independent, an individual LLR score is calculated for each predictor and summed along with a prior to produce an overall LLR score (equation 2).

$$LLR(s, p, c | mf) = LLR(mf) + LLR(s | mf) * P(s) + LLR(p | mf) + LLR(c | mf) \quad (2)$$

Individual log-likelihood ratios are calculated as the ratio of the likelihood the mf term is true over the likelihood the term is false for an individual predictor x (where x is p , c or s) and is formally defined by the equation:

$$LLR(x | mf) = \log\left(\frac{P(x | mf)}{P(x | \neg mf)}\right) \quad (3)$$

The mf prior LLR is calculated as the ratio of the probability the mf term is true over the probability the term is false and is given by the equation:

$$LLR(mf) = \log\left(\frac{P(mf)}{P(\neg mf)}\right) \quad (4)$$

To account for the error in our structure evidence, the LLR of the structure predictor was scaled by the confidence of the classification. Superfamily classifications based on structural evidence ($P(s)$) are fixed for PDB-BLAST and FFAS alignments at 1.0 and 0.9 respectively to reflect the

expected error of these methods. *De novo* based LLR scores are scaled to 0.8 of the MCM score to represent the estimated uncertainty of our predictions:

$$P_{MCM}(s) = MCM(s) * 0.8 \quad (5)$$

The method was applied to predict functions for all protein domains predicted by Ginzu and had a predicted SCOP superfamily classification. All molecular functions predicted above a cutoff of $LLR \geq -3$ are saved but only $LLR > 0$ are considered confident (i.e. more likely correct than incorrect).

4.5.1 Training data and calculation of function probability tables

A set of sequences annotated with GO terms and structure was created for the purpose of building probability tables required for our function prediction method. We first compared sequences with known GO annotations from June 2009 MYGO lite database (Ashburner et al. 2000) to the Astral95 1.75 database of structurally classified domain sequences (Chandonia et al. 2004) using BLAST. BLAST matches with expectation values better than 10^{-8} and a match length greater than 85% of the full length of the Astral sequence were included in the set. Sequences were clustered using CD-HIT with a sequence identity cutoff of 80% and a length difference of 80% to reduce sequence redundancy and sample bias (Li and Godzik 2006). All GO annotations of members within a cluster were assigned to the cluster's representative sequence.

Probability tables for evidence ($x = \{p, c, s\}$) were calculated using the following equation:

$$P_{w/pseudo}(x | mf) = \frac{N(mf \cap x)}{N(MF \cap x)} * \frac{N(mf \cap x) + P(mf) * M}{N(mf \cap x) + M} \quad (6)$$

Where N is the number of sequences with the given annotation, \cap is the intersection and MF is the set of all GO MF terms. Note, in equation 6, M indicates the number of pseudocounts added to the probability distributions. Negligible differences in results were seen when using a range of M values from 4 to 10. The value of M was conservatively chosen to be 10. Priors used in pseudocount calculations are calculated using the following equation:

$$P(mf) = \frac{N(mf)}{N(MF)} \quad (7)$$

4.5.2 Transfer of GO Annotations for features of our function prediction method.

GO annotations were used as features (GO P and C) to perform function prediction. Often, an unannotated protein has a similar protein that is well characterized and it can be inferred that the two share the same GO terms. GO annotations from annotated sequences in MYGO were transferred to sequences in our 94 genomes based on high confidence Blast matches. In addition, true labels (GO MF) were also generated automatically to benchmark our function prediction classifier. Each PFP protein sequence was searched against the MYGO sequence database (Ashburner et al. 2000) and annotations were transferred from Blast matches with an expectation value bet-

ter than 10^{-10} and a match alignment consisting of 85% of the smaller sequence. The MYGO database June 2009 was used for evaluating confidence of function prediction methods. The MYGO database August 2010 was used to generate automated GO annotations for final molecular function predictions.

4.5.3 Selection of relevant features for function prediction

Since the GO ontology is represented as a directed acyclic graph (i.e. specific terms are children of related less-specific terms), multiple terms annotated to a single protein are often conditionally dependent which violates our naïve Bayes assumption of independent predictors. We remedy this by selecting the most relevant GO-P and GO-C terms for predicting each molecular function by way of mutual information. Initially, we observed high mutual information scores for non-specific GO P and C terms with specific MF terms and correct for this by only calculating mutual information for joint probabilities where MF and P or C terms are true. The modified mutual information (MI) of a GO term (x) with a molecular function (mf) is given by the following equation:

$$MI(x;mf) = P(x,mf) * \log\left(\frac{P(x,mf)}{P(x)P(mf)}\right) \quad (8)$$

The P and C terms (x) with the maximum mutual information out of all terms annotated were selected as predictors of function in the naïve Bayes method.

4.6 Evaluating Confidence

4.6.1 *De Novo* Solved After Predicted (SAP)

We built a benchmark of *de novo* structure predictions whose predictions were made prior to an experimental structure being solved. Sequences were compared to Astral 1.75 database using BLAST. A BLAST expectation value of 10^{-5} and alignment length of 80% was used to transfer SCOP superfamily assignments of Astral domains to query sequences. Sequences were then clustered with CD-HIT using 40% identity and 70% length parameters to limit overrepresentation and sequence bias. Sequences of each cluster were then searched in the full PDB using BLAST to identify sequences that were possibly crystallized prior to our structure predictions (2005-01-01). Sequences with Blast matches (<10 eval), which met a threshold of 20% identity and 50% length to PDB sequences deposited prior to 2005, were filtered out. This conservative filtering ensures sequences with *de novo* structure predictions in the SAP set were folded prior to any similar structure solved in the PDB. A list of the sequences used in the SAP set can be found in table S5.

4.6.2 Superfamily Classification Error

To estimate the error of our *de novo* superfamily classifier, the SAP benchmark was classified using an earlier version of SCOP (1.67). We compared these predicted superfamilies to the true superfamilies as determined by sequence comparisons to Astral 1.75 sequences (as described in the previous section). We then only included incorrectly classified entries (i.e. predicted superfamily \neq true superfamily) for further analysis, as these are examples of error in our classification method. We were then able to determine three types of error from this set, specifically, new superfamily error, model error and misclassification error. New superfamily error was calculated as the number of SAP entries in which the true superfamily was not present in SCOP 1.67. To calculate model error and misclassification error we compared predicted structures of a SAP entry to a structure in the incorrectly assigned superfamily and to a structure in the true superfamily using Mammoth z-score. SAP structures that are more similar to the incorrect superfamily than the true superfamily (i.e. higher z-score) are due to insufficient quality of 3D Rosetta models and are therefore considered model error. SAP structures that are closer to the correct superfamily but still incorrectly classified are incorrectly predicted due to misclassification error.

4.6.3 Function Prediction

To analyze the performance of the naïve Bayes function prediction method using PDB-Blast and FFAS03 superfamily classifications as predictors, function predictions were made for 5000 randomly sampled eukaryotic PDB-Blast and FFAS03 proteins with known GO-MF annotations. To analyze performance using *de novo* superfamily classifications as predictors, function predictions were made for 231 domains of the SAP set with known GO-MF. Predictions were compared to the known MF annotation terms (i.e. automated GO annotations based on MYGO June 2009) for accuracy. Only specific molecular functions (annotated to $< 2\%$ of proteins) are included in the set of known annotations. Performance of the function prediction classifier is determined by precision vs recall:

$$\text{Precision} = \frac{TP}{TP + FP} \quad (9)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (10)$$

where TP is true predictions, FP is false predictions and FN is true annotations not predicted. The predictions labeled “random” in precision-recall graphs were produced sampling GO-MF terms from all predictions and ranking by base LLR calculated from priors.

To determine the uniqueness of function prediction between evidence sets (with structure, PCS and without structure, PC), functions were predicted with both evidence sets for a random sampling of 5000 proteins with FFAS03 domain predictions that had SCOP classifications and known GO MF annotations. Predictions were ordered by log likelihood ratio and precision esti-

mates were created based on comparing the predicted MF term to the known (or electronically transferred) MF annotation term.

5 Acknowledgments

We would like to thank the volunteers participating in the Human Proteome Folding Project on IBM's World Community Grid and the FFAS03 development team. We would also like to thank Sasha Levy for comments on the manuscript, Shannon Tsai, John Bako and P. Doug Renfrew for computational support, and Jane Carlton for helpful discussions. In addition, we thank the Rosetta Commons and the Yeast Resource Center. RB, KD and GB were supported by NSF-0820757, NIH-RC4 AI092765-01, and NIH-U54CA143907-01, KD was also supported by NIH T32-GM-088118-01, TD and MR and the development of the YRC were supported by NIH-P41-RR11823.

6 Figure Legends

Figure 1: Flow diagram of the Proteome Folding Pipeline using *Lactobacillus prophage hypothetical protein Ljo_0324* as an example. (A) The Ljo_0324 sequence is first annotated with primary and secondary structure (PSIPRED, DISOPRED and not shown TMHMM, COILS and SignalP). (B) PDB-Blast (PSI-Blast) and fold recognition (FFAS03) are then used to compare the input sequence to the PDB database to identify a similar sequence/structures. (C) Regions of the protein sequence still unannotated are processed by PFAM, multiple sequence alignments (MSA) and a Heuristic method for domain identification. (D) Domains < 150 amino acids are then sent to the computational grid where the three dimensional structure is predicted by Rosetta. (E) Finally, domains with structural annotations (*de novo*, Blast or fold recognition) are classified into SCOP superfamilies. Regions annotated in each level of the figure are outlined in black and regions annotated in previous levels are outlined in dotted lines. Example can be found at <http://www.yeastrc.org/pdr/viewProtein.do?id=2155068>

Figure 2: Distribution of SCOP superfamily classifications. (Top and middle) A set of 875 protein domains (SAP set) folded by Rosetta and whose structure (or protein with strong sequence similarity) has been solved after our prediction was used to determine the accuracy of our method's ability to correctly classify SCOP superfamilies. Plotted in solid lines and dotted lines are the number of correct and incorrect classifications respectively. Classifications are broken down by SCOP class A (α , blue), B (β , yellow), C (β - α - β , red) and D (segregated α and β , green). This graph demonstrates classifications with high MCM scores are the most accurate in classifying SCOP superfamily. (Bottom left) The Precision/Yield plot shows the percentage of protein domains in the SAP set classified using SCOP v1.75 for varying precisions. The line is colored relative to MCM score (right axis). (Bottom right) The histogram of superfamily classification error types represents the total number of incorrectly classified models in the SAP set for different MCM score ranges using a previous version of SCOP (v1.67). New SF error is the error due to a new superfamily (i.e. the true superfamily was not represented in the structure comparison set), 62% of total error. Decoy error is the error due to insufficient *de novo* model quality, 28% of total error. Classifier error is the error due to the superfamily classifier being inaccurate, 10%

of total error.

Figure 3: Function prediction method with structure information out-performs method without structure information. (Upper panel) Precision vs Recall for eukaryotic function predictions separated by the type of structure evidence. The graph shows the precision of function predictions versus recall (fraction of all true annotations) for sequences that were structurally classified by PDB-Blast (left) and FFAS03 (right). The red lines represent the prediction method using GO Process and Component (PC). The green lines represent the prediction method using GO Process, Component and Structure (PCS). The graph shows adding structure information from PDB-Blast and FFAS03 improves precision for function prediction. **(Lower panel)** Function prediction methods using structure information produce novel function predictions for a unique set of proteins. Functions were predicted for FFAS03 domains with SCOP classifications from a random sampling of 5,000 proteins with known molecular function. Black lines represent functions (left) or proteins (right) predicted by both predictor methods. Red lines represent functions (or proteins) predicted only by PC. Blue lines represent functions (or proteins) predicted only by PCS. The predictions made with the integration of structure with process and localization terms annotate a unique range of molecular functions and proteins that would be otherwise unreachable.

Figure 4: *C. crescentus* hypothetical protein CC_3056 is a truncated hemoglobin fold and marine metagenome GOS_6366033 is predicted a chorismate mutase. A. Rosetta model of CC_3056. **B.** Structural alignment of Rosetta model and 1DLWa, a truncated hemoglobin with MAMMOTH alignment z score of 14.39. **C.** Rosetta model viewed with superimposed 1DLW heme ligand. **D.** Rosetta model of GOS_6366033. **E.** Structure alignment of Rosetta model and 1ECMa, a chorismate mutase with structural alignment z-score of 9.71. **F.** Rosetta model viewed with 1ECM transition state analog of chorismate. Ligand contacting residues that are identical and conserved with matched PDB structures are colored in blue and red respectively for **C** and **F**.

7. Tables

SCOP class	PDB-BLAST Total (%)	FFAS03 Total (%)	<i>De novo</i> Total (%)	<i>De novo</i> MedConf (%) *	<i>De novo</i> HighConf (%) **
A	45500 (21.9%)	6148 (20.4%)	35765 (62.3%)	8649 (24.2%)	3432 (9.6%)
B	35117 (16.9%)	4929 (16.3%)	3117 (5.4%)	526 (16.9%)	140 (4.5%)
C	57196 (27.5%)	3999 (13.2%)	3874 (6.8%)	590 (15.2%)	170 (4.4%)
D	38152 (18.4%)	5433 (18.0%)	12463 (21.7%)	2204 (17.7%)	584 (4.7%)
Other	31721 (15.3%)	9674 (32.1%)	2147 (3.7%)	559 (26.0%)	197 (9.2%)
All	207686	30183	57366	12528 (21.8%)	4523 (7.9%)

*MCM score > 0.8

**MCM score > 0.9

Table 1: Superfamily classifications for domains processed by the PFP.

SCOP class	Total (%)	Total correct (%)	MedConf (%)*	MedConf correct (%)*	Yield Med-Conf*	HighConf (%)**	HighConf correct (%)**	Yield HighConf**
A	303 (34.6%)	122 (40.3%)	186 (61.4%)	106 (57.0%)	21.3%	138 (45.5%)	95 (68.8%)	15.8%
B	97 (11.1%)	41 (42.3%)	27 (27.8%)	12 (44.4%)	3.1%	11 (11.3%)	9 (81.8%)	1.3%
C	113 (12.9%)	35 (31.0%)	59 (52.2%)	30 (50.8%)	6.7%	33 (29.2%)	24 (72.7%)	3.8%
D	340 (38.9%)	202 (59.4%)	209 (61.5%)	170 (81.3%)	23.9%	142 (41.8%)	125 (88.0%)	16.2%
Other	22 (2.5%)	7 (31.8%)	10 (45.5%)	6 (60.0%)	1.1%	9 (40.9%)	6 (66.7%)	1.0%
All	875	407 (46.5%)	491 (56.1%)	324 (66.0%)		333 (38.1%)	259 (77.8%)	

*MCM score > 0.8

**MCM score > 0.9

Table 2: Superfamily classifications for SAP structures.

<i>(a) D. radiodurans</i>				
Rank	SCOP id	Enrichment Score	Number of domains	Superfamily name
1	b.1.5	3.774	10	Transglutaminase, two C-terminal domains
2	d.110.7	3.418	7	Roadblock/LC7 domain
3	d.111.1	3.081	5	PR-1-like
4	c.41.1	1.654	6	Subtilisin-like
5	a.3.1	1.523	10	Cytochrome c
6	h.1.5	1.472	6	Tropomyosin
7	d.159.1	1.472	17	Metallo-dependent phosphatases
8	b.1.18	1.317	12	E set domains
9	d.185.1	1.271	9	LuxS/MPP-like metallohydrolase
10	c.58.1	1.220	7	Aminoacid dehydrogenase-like, N-terminal domain
<i>(b) P. vivax</i>				
Rank	SCOP id	Enrichment Score	Number of domains	Superfamily name
1	h.4.2	5.756	120	Clostridium neurotoxins, "coiled-coil" domain
2	b.6.2	5.245	8	Major surface antigen p30, SAG1
3	b.42.4	5.208	131	STI-like
4	a.264.1	5.158	22	Duffy binding domain-like
5	b.61.5	3.348	6	Dipeptidyl peptidase I (cathepsin C), exclusion domain
6	a.24.7	3.166	5	FKBP12-rapamycin-binding domain of FKBP-rapamycin-associated protein (FRAP)
7	a.118.1 1	3.166	11	Cytochrome c oxidase subunit E
8	b.81.3	2.878	6	Adhesin YadA, collagen-binding domain

9	a.56.1	2.809	7	CO dehydrogenase ISP C-domain like
10	a.24.26	2.809	7	YppE-like

Table 3: Top enriched folds of (a) *D. radiodurans* and (b) *P. vivax*

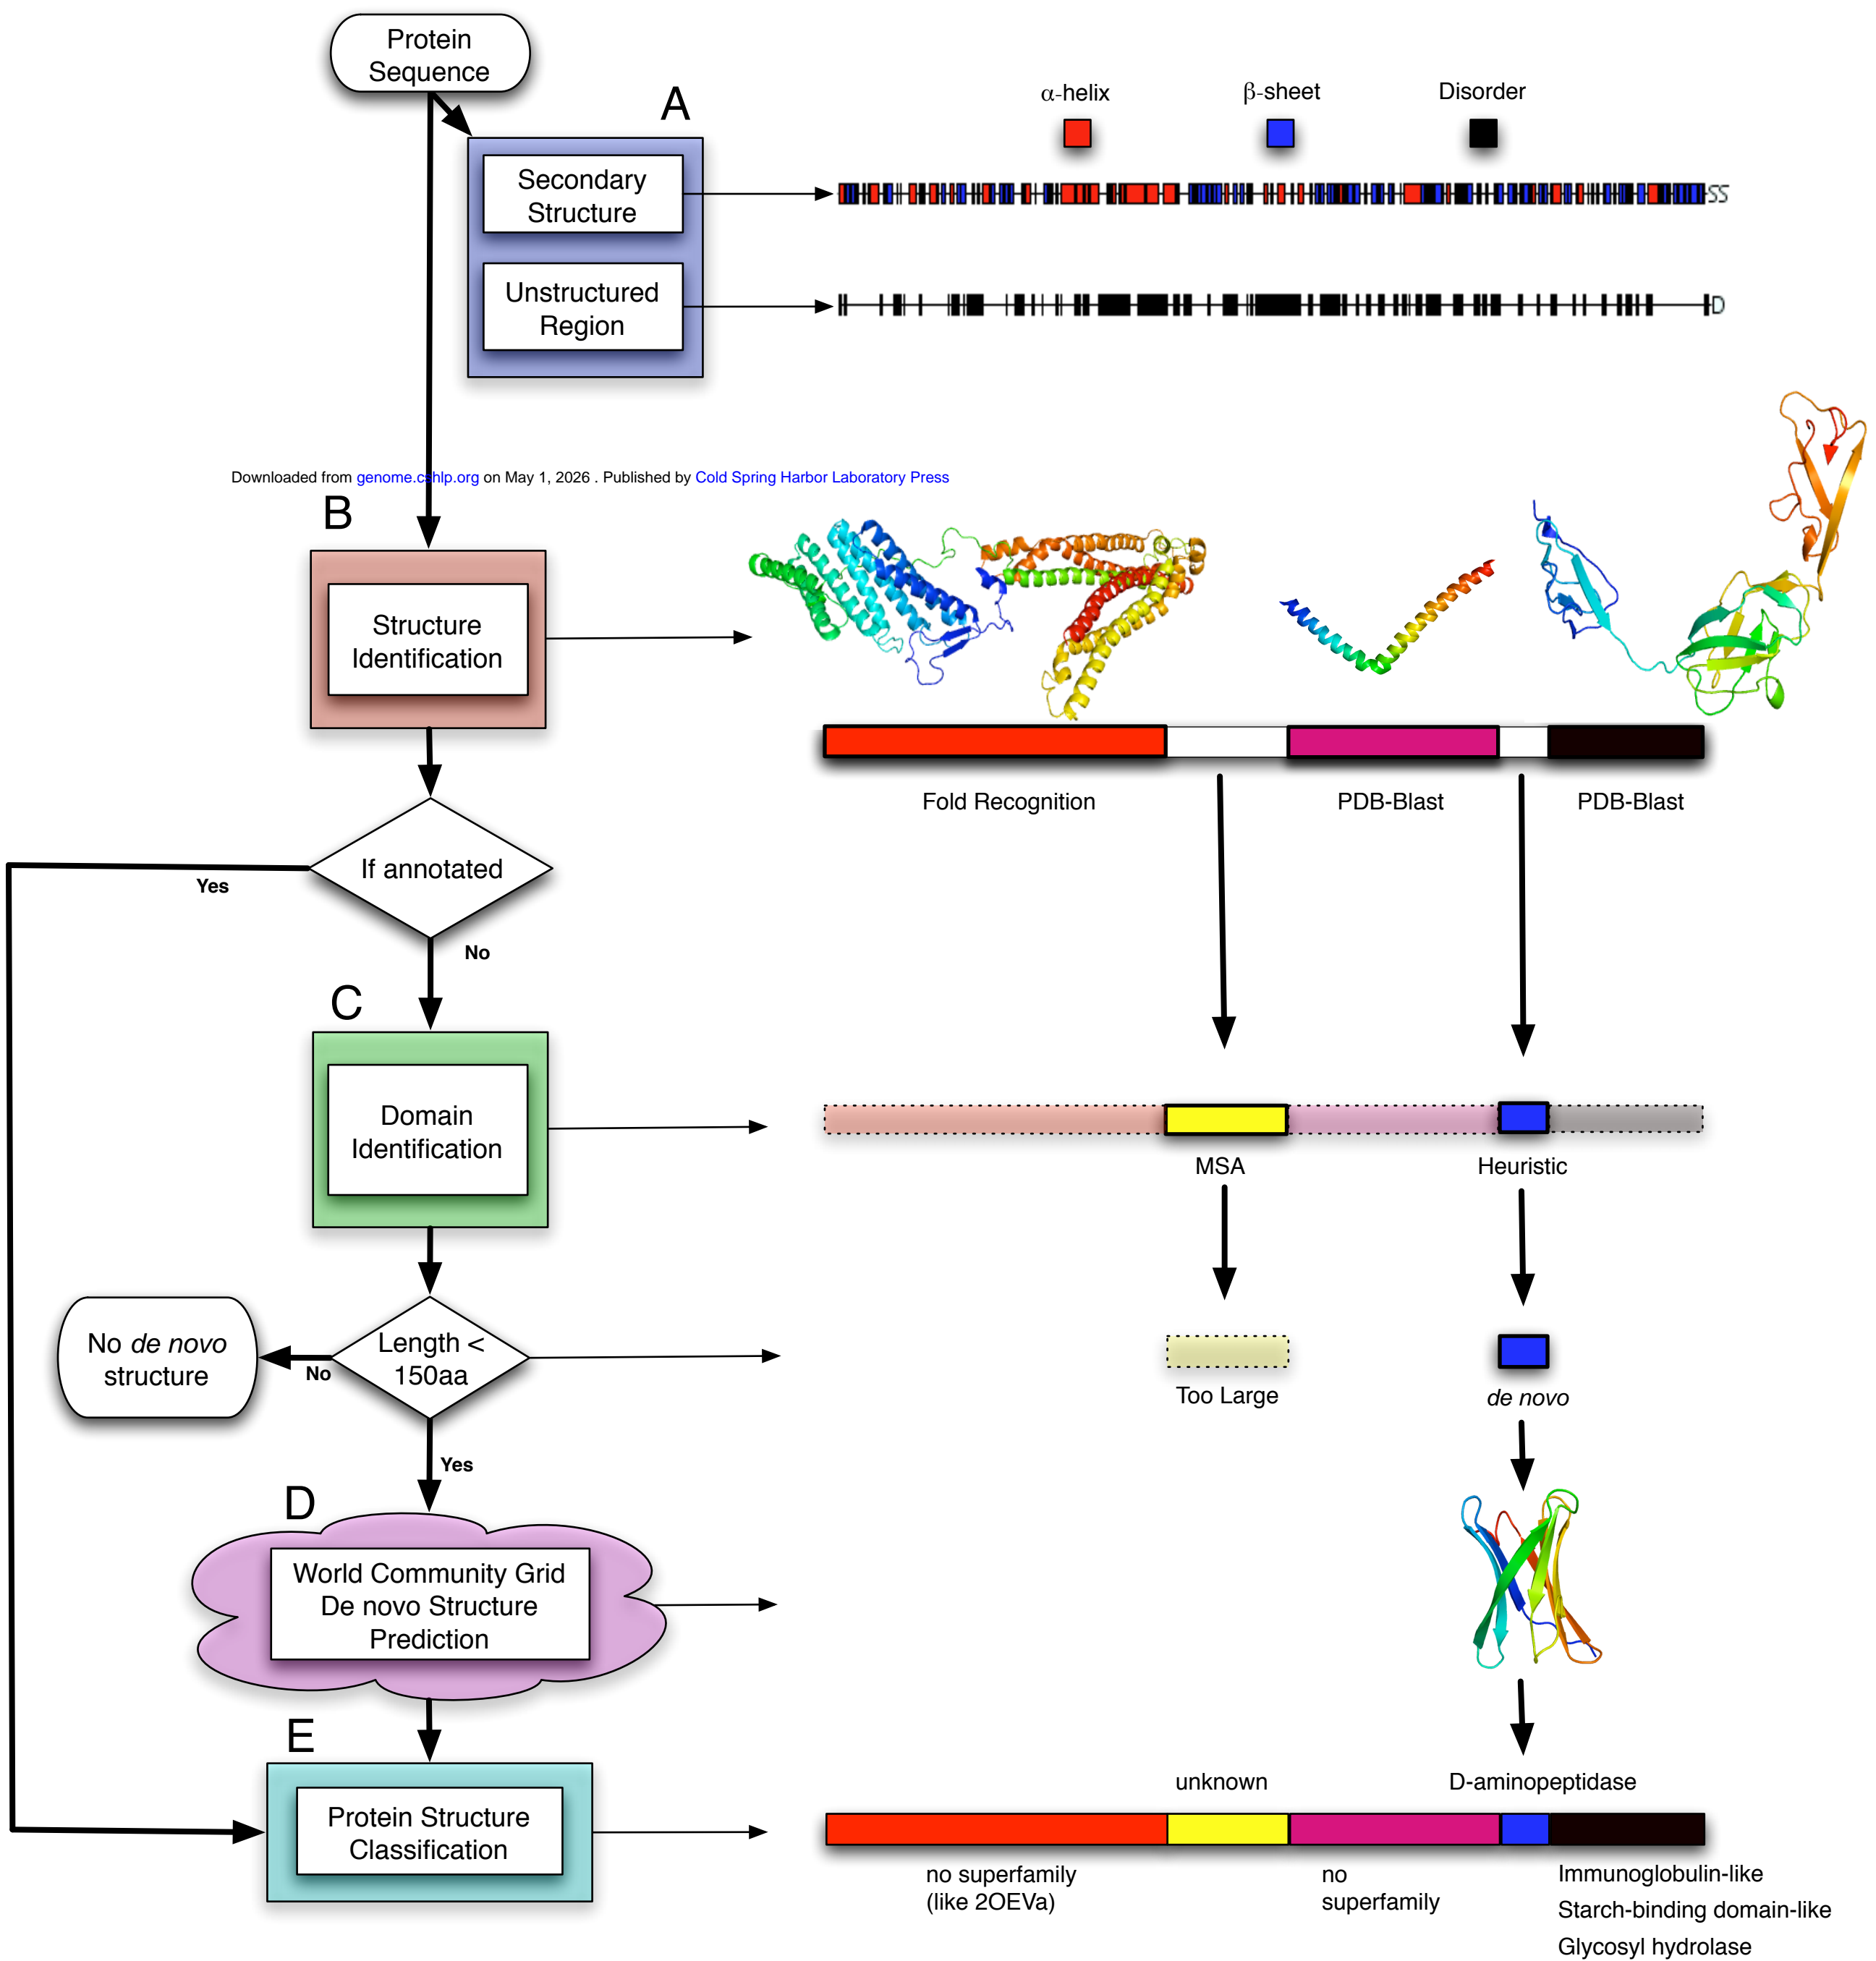
8 References

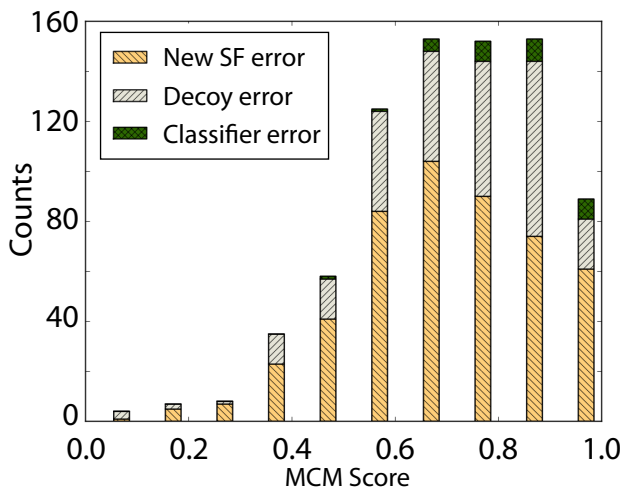

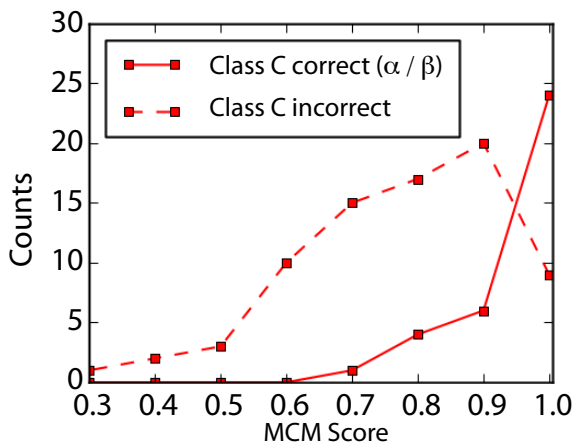
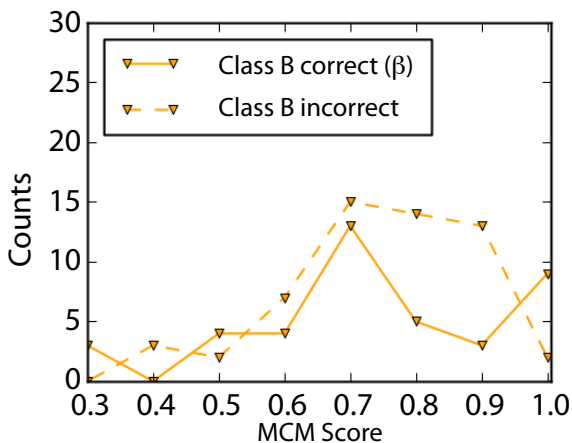
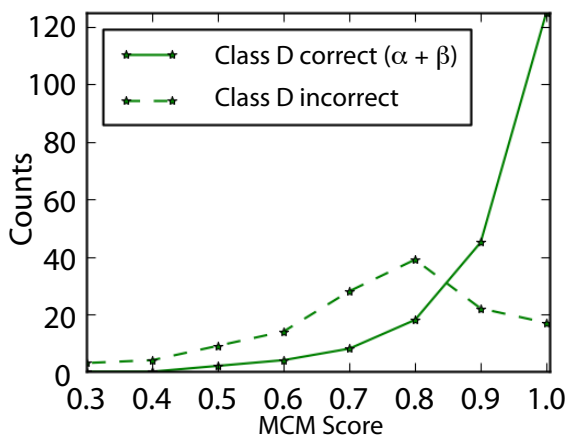
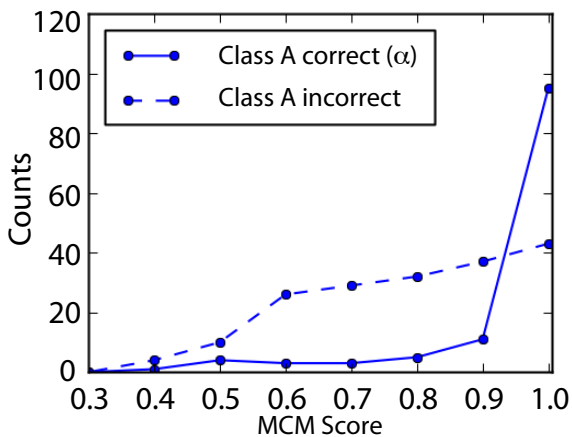
- Acar, M., H. Jafar-Nejad, H. Takeuchi, A. Rajan, D. Ibrani, N.A. Rana, H. Pan, R.S. Haltiwanger, and H.J. Bellen. 2008. Rumi is a CAP10 domain glycosyltransferase that modifies Notch and is required for Notch signaling. *Cell* **132**: 247-258.
- Altschul, S.F., T.L. Madden, A.A. Schäffer, J. Zhang, Z. Zhang, W. Miller, and D.J. Lipman. 1997. Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic Acids Res.* **25**: 3389--3402.
- Anantharaman, V., E.V. Koonin, and L. Aravind. 2001. Peptide-N-glycanases and DNA repair proteins, Xp-C/Rad4, are, respectively, active and inactivated enzymes sharing a common transglutaminase fold. *Hum Mol Genet* **10**: 1627-1630.
- Anderson, D.P. 2004. BOINC: A System for Public-Resource Computing and Storage. In *GRID* (ed. R. Buyya), pp. 4-10. IEEE Computer Society.
- Ashburner, M., C.A. Ball, J.A. Blake, D. Botstein, H. Butler, J.M. Cherry, A.P. Davis, K. Dolinski, S.S. Dwight, J.T. Eppig et al. 2000. Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat Genet* **25**: 25-29.
- Avila-Campillo, I., K. Drew, J. Lin, D.J. Reiss, and R. Bonneau. 2007. BioNetBuilder: automatic integration of biological networks. *Bioinformatics* **23**: 392--393.
- Bader, G.D. and C.W.V. Hogue. 2002. Analyzing yeast protein-protein interaction data obtained from different sources. *Nat Biotechnol* **20**: 991-997.
- Bendtsen, J., H. Nielsen, G. von Heijne, and S. Brunak. 2004. Improved prediction of signal peptides: SignalP 3.0. *Journal of Molecular Biology* **340**: 783-795.
- Bonneau, R., N. Baliga, E. Deutsch, P. Shannon, and L. Hood. 2004. Comprehensive de novo structure prediction in a systems-biology context for the archaea *Halobacterium* sp NRC-1. *Genome Biology* **5**.
- Bonneau, R., M.T. Facciotti, D.J. Reiss, A.K. Schmid, M. Pan, A. Kaur, V. Thorsson, P. Shannon, M.H. Johnson, J.C. Bare et al. 2007. A predictive model for transcriptional control of physiology in a free living cell. *Cell* **131**: 1354-1365.
- Boxem, M., Z. Maliga, N. Klitgord, N. Li, I. Lemmens, M. Mana, L. de Lichtenvelde, J.D. Mul, D. van de Peut, M. Devos et al. 2008. A protein domain-based interactome network for *C. elegans* early embryogenesis. *Cell* **134**: 534-545.
- Bradley, P., L. Malmström, B. Qian, J. Schonbrun, D. Chivian, D.E. Kim, J. Meiler, K.M.S. Misura, and D. Baker. 2005. Free modeling with Rosetta in CASP6. *Proteins* **61 Suppl 7**: 128-134.
- Chandonia, J.M., G. Hon, N.S. Walker, L. Lo Conte, P. Koehl, M. Levitt, and S.E. Brenner. 2004. The ASTRAL Compendium in 2004. *Nucleic Acids Res.* **32**: D189--192.
- Chivian, D., D.E. Kim, L. Malmstrom, P. Bradley, T. Robertson, P. Murphy, C.E. Strauss, R. Bonneau, C.A. Rohl, and D. Baker. 2003. Automated prediction of CASP-5 structures using the Robetta server. *Proteins* **53 Suppl 6**: 524-533.
- Chivian, D., D.E. Kim, L. Malmström, J. Schonbrun, C.A. Rohl, and D. Baker. 2005. Prediction of CASP6 structures using automated Robetta protocols. *Proteins* **61 Suppl 7**: 157-166.
- Chothia, C., J. Gough, C. Vogel, and S.A. Teichmann. 2003. Evolution of the protein repertoire. *Science* **300**: 1701-1703.
- Chothia, C. and A.M. Lesk. 1986. The relation between the divergence of sequence and structure in proteins. *EMBO J* **5**: 823-826.

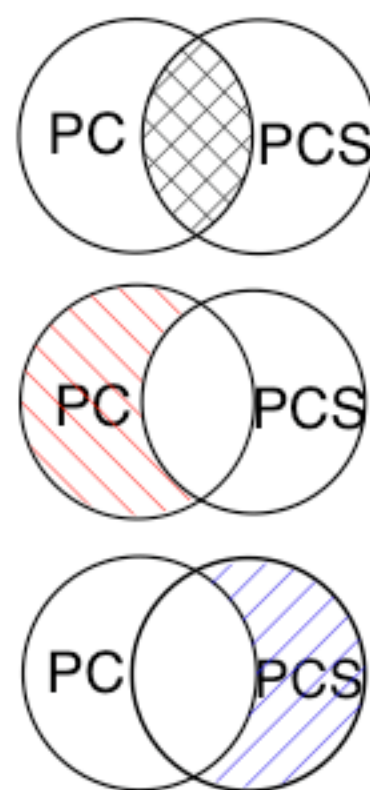



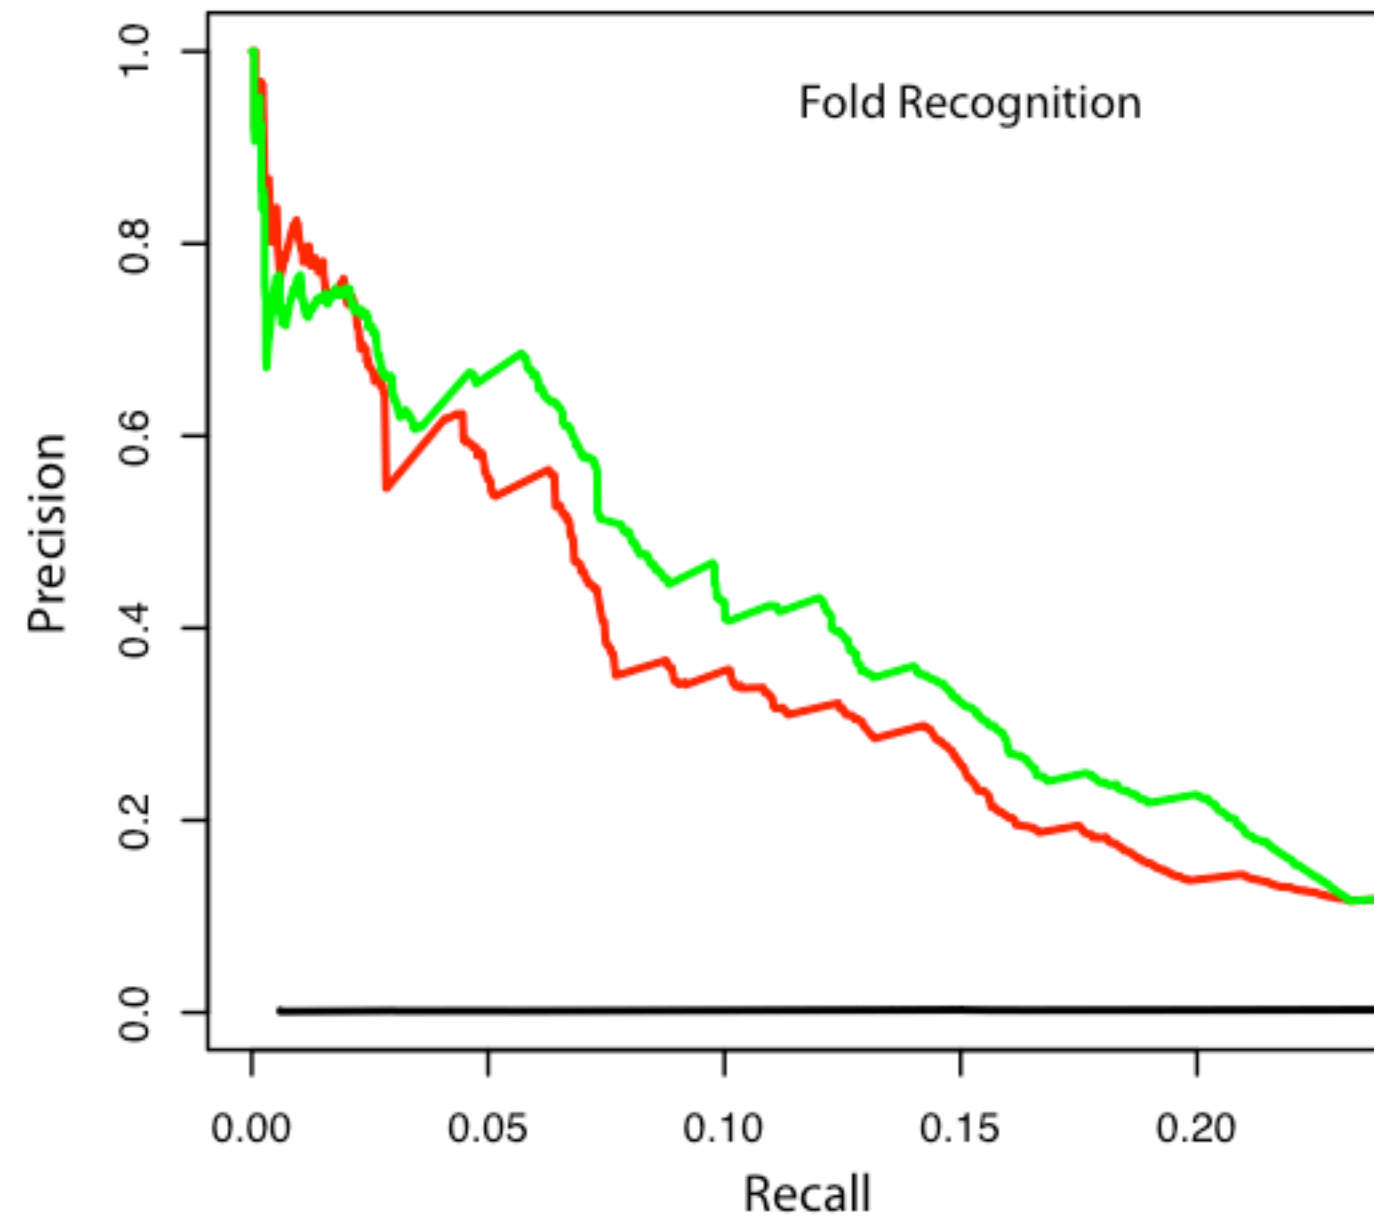
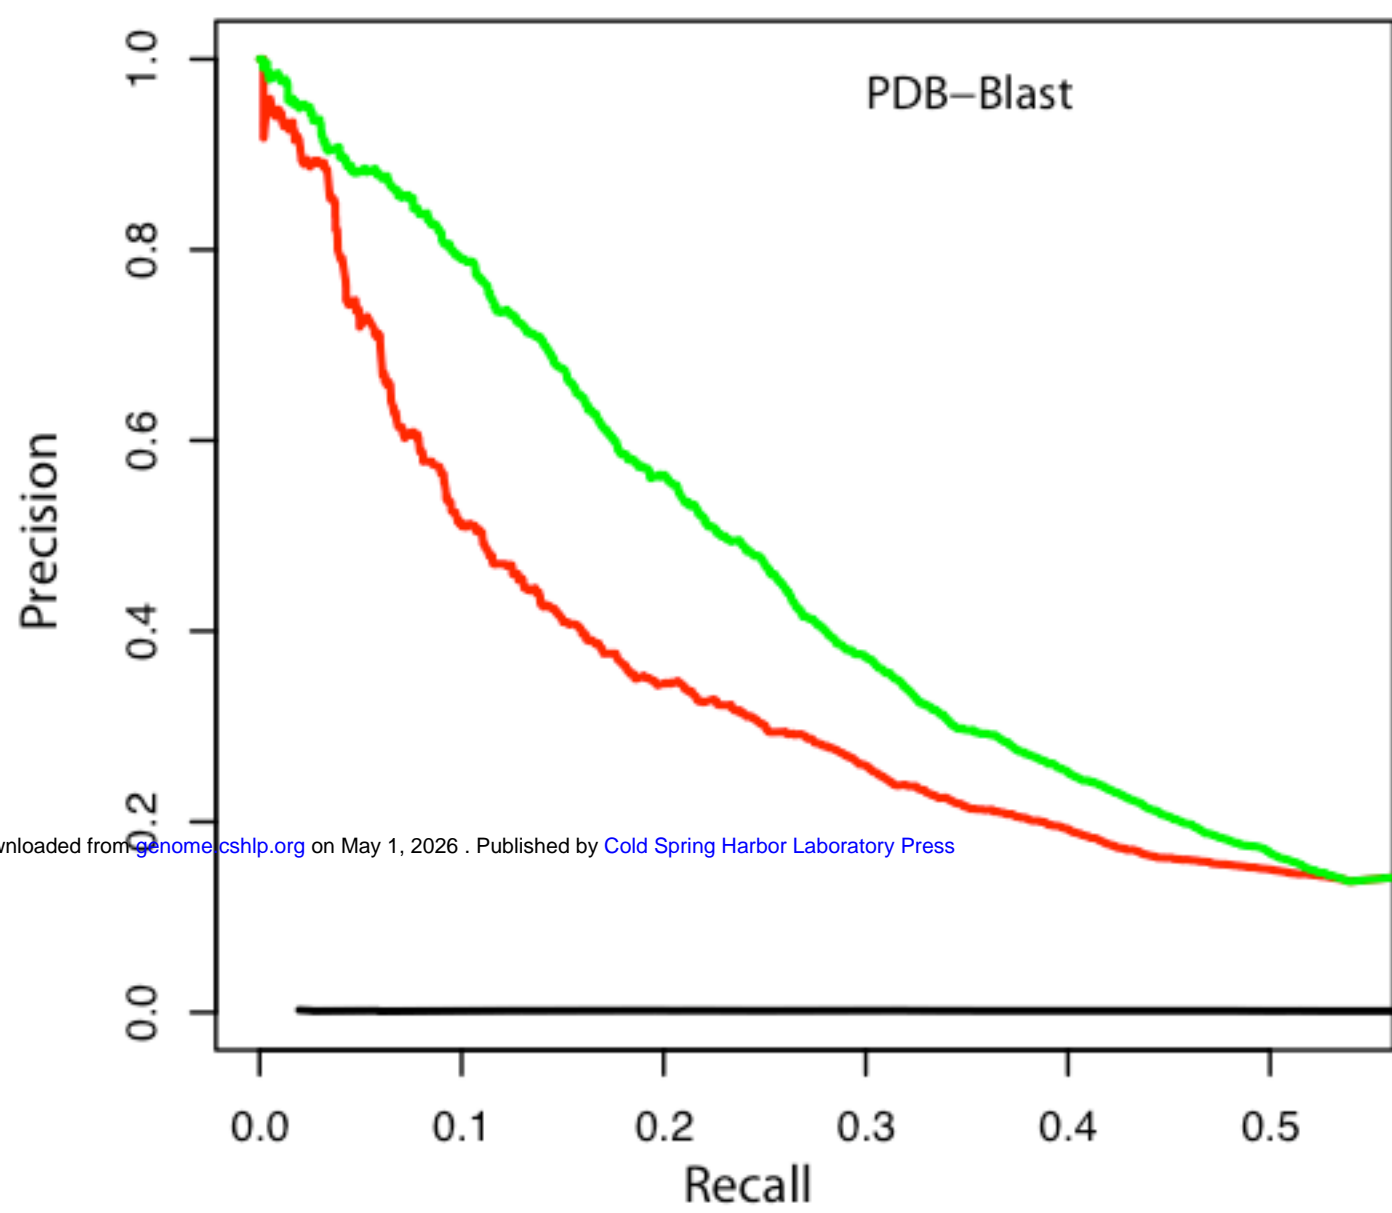
- Cox, M.M. and J.R. Battista. 2005. Deinococcus radiodurans - the consummate survivor. *Nat Rev Microbiol* **3**: 882-892.
- Cuff, A., O.C. Redfern, L. Greene, I. Sillitoe, T. Lewis, M. Dibley, A. Reid, F. Pearl, T. Dallman, A. Todd et al. 2009. The CATH Hierarchy Revisited-Structural Divergence in Domain Superfamilies and the Continuity of Fold Space. *Structure* **17**: 1051-1062.
- Dessailly, B.H., R. Nair, L. Jaroszewski, J.E. Fajardo, A. Kouranov, D. Lee, A. Fiser, A. Godzik, B. Rost, and C. Orengo. 2009. PSI-2: structural genomics to cover protein domain family space. *Structure* **17**: 869-881.
- Finn, R.D., J. Tate, J. Mistry, P.C. Coggill, S.J. Sammut, H.R. Hotz, G. Ceric, K. Forslund, S.R. Eddy, E.L. Sonnhammer et al. 2008. The Pfam protein families database. *Nucleic Acids Res.* **36**: D281--288.
- Ginalski, K., L. Rychlewski, D. Baker, and N.V. Grishin. 2004. Protein structure prediction for the male-specific region of the human Y chromosome. *Proc Natl Acad Sci U S A* **101**: 2305-2310.
- Guigó, R., P. Flicek, J.F. Abril, A. Reymond, J. Lagarde, F. Denoeud, S. Antonarakis, M. Ashburner, V.B. Bajic, E. Birney et al. 2006. EGASP: the human ENCODE Genome Annotation Assessment Project. *Genome Biol* **7 Suppl 1**: S2.1-31.
- Hazbun, T.R., L. Malmström, S. Anderson, B.J. Graczyk, B. Fox, M. Riffle, B.A. Sundin, J.D. Aranda, W.H. McDonald, C.-H. Chiu et al. 2003. Assigning function to yeast proteins by integration of technologies. *Mol Cell* **12**: 1353-1365.
- Herr, J.C., O. Chertihin, L. Digilio, K.N. Jha, S. Vemuganti, and C.J. Flickinger. 2008. Distribution of RNA binding protein MOEP19 in the oocyte cortex and early embryo indicates pre-patterning related to blastomere polarity and trophectoderm specification. *Dev Biol* **314**: 300-316.
- Jaroszewski, L., L. Rychlewski, Z. Li, W. Li, and A. Godzik. 2005. FFAS03: a server for profile-profile sequence alignments. *Nucleic Acids Res.* **33**: W284--288.
- Jones, D. 1999. Protein secondary structure prediction based on position-specific scoring matrices. *Journal of Molecular Biology* **292**: 195-202.
- Jones, D. and J. Ward. 2003. Prediction of disordered regions in proteins from position specific score matrices. *Proteins-Structure Function and Bioinformatics* **53**: 573-578.
- Kim, D.E., D. Chivian, L. Malmström, and D. Baker. 2005. Automated prediction of domain boundaries in CASP6 targets using GinzU and RosettaDOM. *Proteins* **61 Suppl 7**: 193-200.
- Krogh, A., B. Larsson, G. von Heijne, and E. Sonnhammer. 2001. Predicting transmembrane protein topology with a hidden Markov model: Application to complete genomes. *Journal of Molecular Biology* **305**: 567-580.
- LEE, A., P. KARPLUS, B. GANEM, and J. CLARDY. 1995. ATOMIC-STRUCTURE OF THE BURIED CATALYTIC POCKET OF ESCHERICHIA-COLI CHORISMATE MUTASE. *Journal of the American Chemical Society* **117**: 3627-3628.
- Lee, D., O. Redfern, and C. Orengo. 2007. Predicting protein function from sequence and structure. *Nat Rev Mol Cell Biol* **8**: 995-1005.
- Lee, I., S.V. Date, A.T. Adai, and E.M. Marcotte. 2004. A probabilistic functional network of yeast genes. *Science* **306**: 1555-1558.
- Li, W. and A. Godzik. 2006. Cd-hit: a fast program for clustering and comparing large sets of protein or nucleotide sequences. *Bioinformatics* **22**: 1658-1659.


- Liu, J. and B. Rost. 2004. CHOP: parsing proteins into structural domains. *Nucleic Acids Res* **32**: W569-571.
- Liu, Y., J. Zhou, M.V. Omelchenko, A.S. Beliaev, A. Venkateswaran, J. Stair, L. Wu, D.K. Thompson, D. Xu, I.B. Rogozin et al. 2003. Transcriptome dynamics of *Deinococcus radiodurans* recovering from ionizing radiation. *Proc Natl Acad Sci U S A* **100**: 4191-4196.
- Lopez, G., A. Valencia, and M. Tress. 2007. FireDB--a database of functionally important residues from proteins of known structure. *Nucleic Acids Res* **35**: D219-223.
- LUPAS, A., M. VANDYKE, and J. STOCK. 1991. PREDICTING COILED COILS FROM PROTEIN SEQUENCES. *Science* **252**: 1162-1164.
- Makarova, K.S., L. Aravind, Y.I. Wolf, R.L. Tatusov, K.W. Minton, E.V. Koonin, and M.J. Daly. 2001. Genome of the extremely radiation-resistant bacterium *Deinococcus radiodurans* viewed from the perspective of comparative genomics. *Microbiol Mol Biol Rev* **65**: 44-79.
- Malmström, L., M. Riffle, C.E.M. Strauss, D. Chivian, T.N. Davis, R. Bonneau, and D. Baker. 2007. Superfamily Assignments for the Yeast Proteome through Integration of Structure Prediction with the Gene Ontology. *PLoS Biol* **5**: e76.
- Marcotte, E.M., M. Pellegrini, H.L. Ng, D.W. Rice, T.O. Yeates, and D. Eisenberg. 1999. Detecting protein function and protein-protein interactions from genome sequences. *Science* **285**: 751-753.
- Marsden, R.L., T.A. Lewis, and C.A. Orengo. 2007. Towards a comprehensive structural coverage of completed genomes: a structural genomics viewpoint. *BMC Bioinformatics* **8**: 86.
- Matthews, B.W. 2007. Protein Structure Initiative: getting into gear. *Nat Struct Mol Biol* **14**: 459-460.
- Murzin, A.G., S.E. Brenner, T. Hubbard, and C. Chothia. 1995. SCOP: a structural classification of proteins database for the investigation of sequences and structures. *J. Mol. Biol.* **247**: 536--540.
- Nardini, M., A. Pesce, M. Labarre, C. Richard, A. Bolli, P. Ascenzi, M. Guertin, and M. Bolognesi. 2006. Structural determinants in the group III truncated hemoglobin from *Campylobacter jejuni*. *J Biol Chem* **281**: 37803-37812.
- Nielsen, P. and A. Krogh. 2005. Large-scale prokaryotic gene prediction and comparison to genome annotation. *Bioinformatics* **21**: 4322-4329.
- Orengo, C.A., A.D. Michie, S. Jones, D.T. Jones, M.B. Swindells, and J.M. Thornton. 1997. CATH--a hierarchic classification of protein domain structures. *Structure* **5**: 1093-1108.
- Ortiz, A.R., C.E. Strauss, and O. Olmea. 2002. MAMMOTH (matching molecular models obtained from theory): an automated method for model comparison. *Protein Sci.* **11**: 2606--2621.
- Pena-Castillo, L., M. Tasan, C. Myers, H. Lee, T. Joshi, C. Zhang, Y. Guan, M. Leone, A. Pagnani, W. Kim et al. 2008. A critical assessment of *Mus musculus* gene function prediction using integrated genomic evidence. *Genome Biology* **9**: S2.
- Pieper, U., N. Eswar, B.M. Webb, D. Eramian, L. Kelly, D.T. Barkan, H. Carter, P. Mankoo, R. Karchin, M.A. Marti-Renom et al. 2009. MODBASE, a database of annotated comparative protein structure models and associated resources. *Nucleic Acids Res* **37**: D347-354.

- Rieckmann, K.H., D.R. Davis, and D.C. Hutton. 1989. Plasmodium vivax resistance to chloroquine? *Lancet* **2**: 1183-1184.
- Riffle, M., L. Malmström, and T.N. Davis. 2005. The Yeast Resource Center Public Data Repository. *Nucleic Acids Res.* **33**: D378--382.
- Rohl, C.A., C.E.M. Strauss, K.M.S. Misura, and D. Baker. 2004. Protein structure prediction using Rosetta. *Methods Enzymol* **383**: 66-93.
- Rychlewski, L., L. Jaroszewski, W. Li, and A. Godzik. 2000. Comparison of sequence profiles. Strategies for structural predictions using sequence information. *Protein Sci* **9**: 232-241.
- Shannon, P., A. Markiel, O. Ozier, N.S. Baliga, J.T. Wang, D. Ramage, N. Amin, B. Schwikowski, and T. Ideker. 2003. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res* **13**: 2498-2504.
- Troyanskaya, O.G., K. Dolinski, A.B. Owen, R.B. Altman, and D. Botstein. 2003. A Bayesian framework for combining heterogeneous data sources for gene function prediction (in *Saccharomyces cerevisiae*). *Proc Natl Acad Sci U S A* **100**: 8348-8353.
- Vinogradov, S.N. and L. Moens. 2008. Diversity of globin function: enzymatic, transport, storage, and sensing. *J Biol Chem* **283**: 8773-8777.
- Vogel, C., C. Berzuini, M. Bashton, J. Gough, and S.A. Teichmann. 2004. Supra-domains: evolutionary units larger than single protein domains. *J Mol Biol* **336**: 809-823.
- Yooseph, S., G. Sutton, D.B. Rusch, A.L. Halpern, S.J. Williamson, K. Remington, J.A. Eisen, K.B. Heidelberg, G. Manning, W. Li et al. 2007. The Sorcerer II Global Ocean Sampling expedition: expanding the universe of protein families. *PLoS Biol* **5**: e16.
- Zhang, Y., I. Thiele, D. Weekes, Z. Li, L. Jaroszewski, K. Ginalski, A.M. Deacon, J. Wooley, S.A. Lesley, I.A. Wilson et al. 2009. Three-dimensional structural view of the central metabolic network of *Thermotoga maritima*. *Science* **325**: 1544-1549.
- Zhou, H., Y. Luo, and S. Huang. Updates of mTOR inhibitors. *Anticancer Agents Med Chem.*






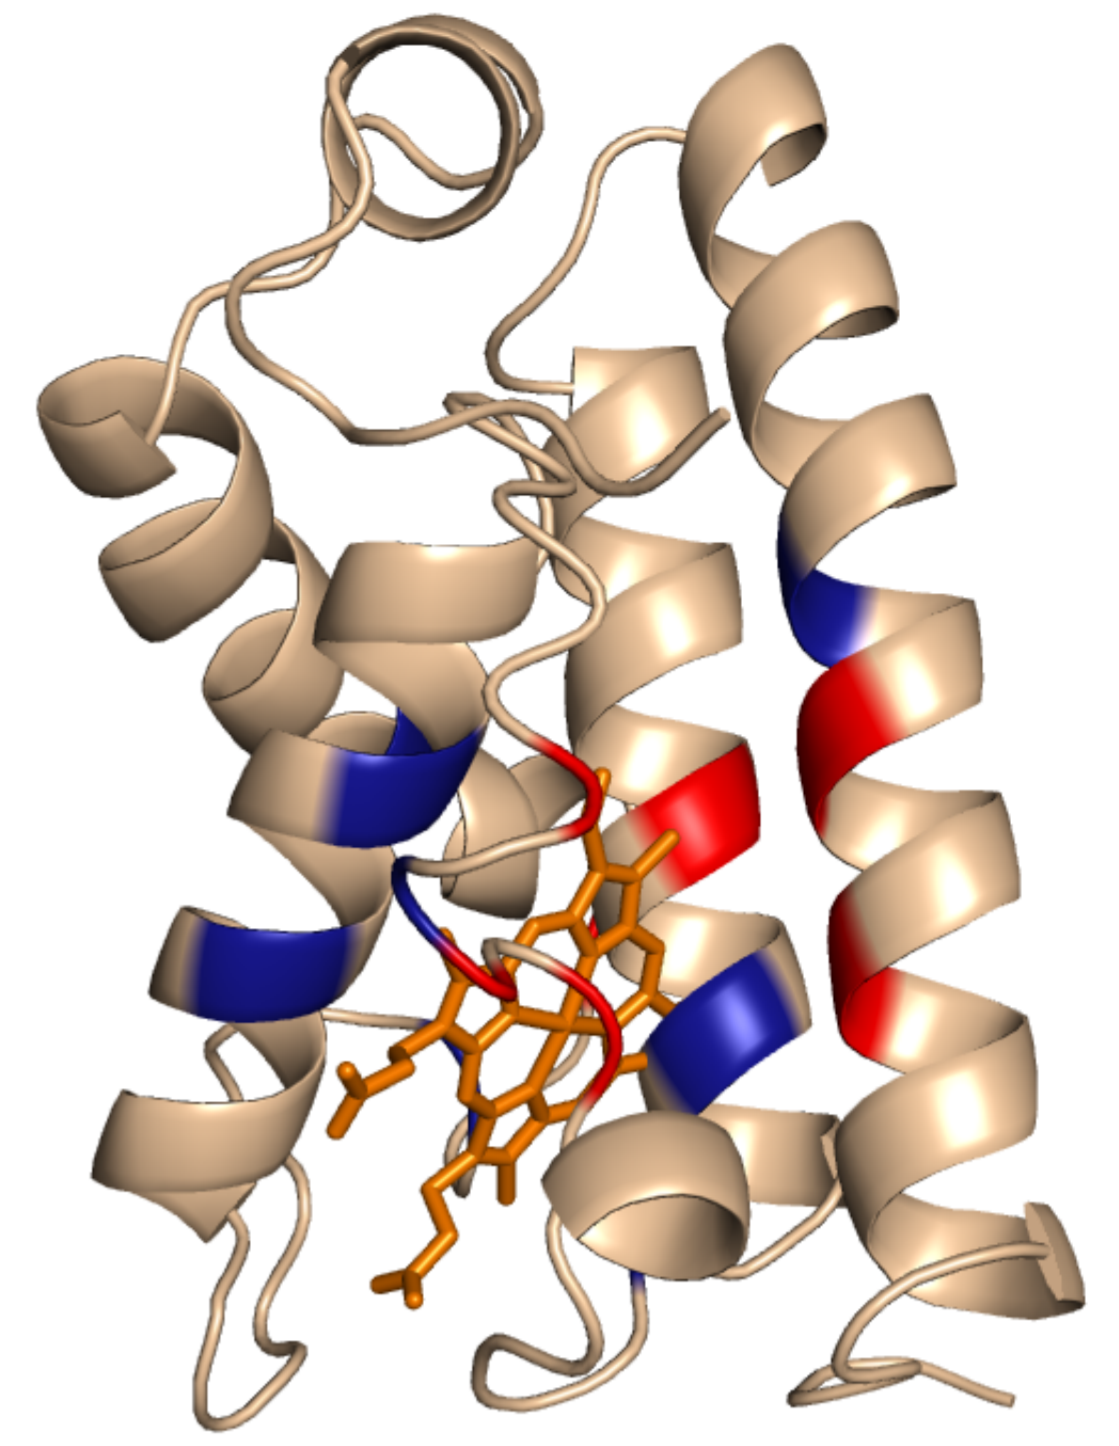




A. Rosetta model of CC_3056




B. Structural alignment w/ 1DLWa




C. Rosetta Model and 1DLWa heme


Downloaded from genome.cshlp.org on May 1, 2026 . Published by Cold Spring Harbor Laboratory Press



D. Rosetta model of GOS_6366033



E. Structural alignment w/ 1ECMa



F. Rosetta Model and 1ECMa TSA