Published in final edited form as:

J Proteome Res. 2012 September 7; 11(9): 4476–4487. doi:10.1021/pr300227y.

Computational Framework for Analysis of Prey–Prey Associations in Interaction Proteomics Identifies Novel Human Protein–Protein Interactions and Networks

Sudipto Saha[†], Jean-Eudes Dazard[†], Hua Xu[†], and Rob M. Ewing^{†,‡,*}

[†]Center for Proteomics and Bioinformatics, Case Western Reserve University School of Medicine, Cleveland, Ohio 44106, United States

[‡]Department of Genetics and Genome Sciences, Case Western Reserve University School of Medicine, Cleveland, Ohio 44106, United States

Abstract

Large-scale protein–protein interaction data sets have been generated for several species including yeast and human and have enabled the identification, quantification, and prediction of cellular molecular networks. Affinity purification-mass spectrometry (AP-MS) is the preeminent methodology for large-scale analysis of protein complexes, performed by immunopurifying a specific “bait” protein and its associated “prey” proteins. The analysis and interpretation of AP-MS data sets is, however, not straightforward. In addition, although yeast AP-MS data sets are relatively comprehensive, current human AP-MS data sets only sparsely cover the human interactome. Here we develop a framework for analysis of AP-MS data sets that addresses the issues of noise, missing data, and sparsity of coverage in the context of a current, real world human AP-MS data set. Our goal is to extend and increase the density of the known human interactome by integrating bait–prey and cocomplexed preys (prey–prey associations) into networks. Our framework incorporates a score for each identified protein, as well as elements of signal processing to improve the confidence of identified protein–protein interactions. We identify many protein networks enriched in known biological processes and functions. In addition, we show that integrated bait–prey and prey–prey interactions can be used to refine network topology and extend known protein networks.

Keywords

protein–protein interaction network; affinity purification-mass spectrometry; interactome

© 2012 American Chemical Society

*Corresponding Author: Rob M. Ewing, Ph.D., Center for Proteomics and Bioinformatics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106. Tel: (216) 368-4380. Fax: (216) 368-6846. rob.ewing@case.edu.

The authors declare no competing financial interest.

INTRODUCTION

Interactions between proteins and the protein complexes and networks that these interactions form are fundamental units of biological organization that mediate most cellular processes. Understanding the topology of protein interaction networks is therefore a primary goal of systems biology. The principal techniques for large-scale analyses of protein interactions are yeast two-hybrid and affinity purification-mass spectrometry (AP-MS). These two approaches provide complementary views of the protein interactome;¹ the yeast-two-hybrid assay identifies binary interactions between pairs of proteins, while AP-MS identifies protein complexes associated with a given “bait” protein. AP-MS experiments have been conducted by purifying protein complexes using native antibodies or on a larger-scale by epitope tagging bait proteins and recovering associated protein complexes with antibodies directed against the epitope tag. Despite the power of AP-MS to map protein complexes, however, several technical hurdles exist, including the presence of nonspecific interacting proteins in these large-scale data sets as well as lack of reproducibility and sparse coverage of the underlying networks.

The yeast protein interactome is by far the most well covered eukaryotic interactome with multiple studies and techniques contributing large-scale data sets.^{2–5} With its larger size and greater complexity, the human protein interactome has been mapped more selectively, with studies focusing on specific complexes or classes of protein.^{6–9} Diverse computational methods have been developed for analysis of large-scale interaction proteomics data sets. For AP-MS data sets, the focus of these analysis methods has been the assignment of confidence scores to observed interactions and to distinguish specific from nonspecific interactions.^{6,7,10–14} In yeast, the socio-affinity index was developed to score a large-scale yeast interactome data set with high coverage and reciprocal AP-MS experiments.¹⁰

In human data sets, with less dense coverage of the underlying networks, computational methods and scores have focused on distinguishing specific from nonspecific interactions. For example, the normalized spectral abundance factor (NSAF) was used as a measure of abundance for each protein and the ratio of the vectors of counts for each protein between “control” and “bait” experiments used to distinguish specific from nonspecific interactors.⁶ The D-score, a metric combining total spectral count, reproducibility and overall frequency of prey proteins in AP-MS experiments into a single confidence score for bait–prey associations, was developed for the analysis of human AP-MS data.⁷ The D-score was applied to a large-scale study of human deubiquitinating enzymes and shown to out-perform the socio-affinity index, the NSAF method and Z-score on the data in hand.⁷ An alternative approach, using mixture modeling and Bayesian statistics named SAINT (Significance Analysis of Interactome) was developed¹³ and applied to the analysis of phosphatase interaction networks.¹⁵ The Decontaminator uses a Bayesian approach to model false-positive protein–protein interactions (PPIs) by comparing the score of a putative prey protein in induced vs control experiments.¹⁴

Two broad models for interpreting AP-MS data in a network context have been proposed. The “spoke” model holds that the bait protein interacts with each of the identified “prey” proteins (the bait representing the center of a wheel with spokes connecting to each of the prey proteins) whereas the “matrix” model assumes that each of the identified proteins (bait and prey) interacts with each of the others in a given AP-MS experiment.¹⁶ Although the matrix model will capture a higher proportion of the underlying protein associations, this comes with the price of increased false positives. Therefore, although both models have merits, it is likely that a combination of spoke and matrix models represent underlying biological reality in most cases. Several scoring metrics for AP-MS data explicitly implement these concepts. The socio-affinity index is a summary score including spoke and

matrix terms, as well as accounting for reciprocal bait-prey AP-MS experiments and computes the log ratio of protein co-occurrence given their observed frequencies.¹⁰ The hypergeometric distribution has been used to calculate expected frequencies of co-occurrence for proteins using the matrix model in the large-scale yeast AP-MS data sets and found to be an effective means of identifying protein-protein interactions.¹⁷ Other schemes that make use of co-occurrence frequencies have been proposed for resolving protein complexes using large-scale yeast AP-MS data sets^{18,19} but with added refinements such as consideration of the variation in bait affinity when computing the results.¹⁸

Probabilistic approaches using small to medium scale human AP-MS data sets have also been developed.^{6,20} These studies analyze data sets in which the coverage of given protein complexes by bait proteins is relatively high. Notably, these studies^{6,20} make use of the quantitative features of mass-spectrometry data, rather than using binary co-occurrence frequencies. Measures of abundance such as spectral counts and other mass-spectrometry based confidence measures are very informative features of AP-MS data, since protein-protein affinities are likely to vary considerably. An emerging consensus of these diverse computational methods is that both bait-prey and prey-prey associations in AP-MS data sets can be used to identify protein-protein interactions. Utilization of prey-prey associations is most obvious in AP-MS data sets with dense bait coverage, so that co-occurrence profiles are well-defined. However, it is less clear whether this concept may be applied to less dense AP-MS data sets such as larger-scale human AP-MS data sets.^{7,12} Our goal in this work is to test the utility of mining prey-prey associations from large-scale, but less dense, human AP-MS data.

We previously generated a large-scale human interaction proteomics data set focused on 338 human bait proteins, many of which are linked to human diseases.¹² This data set is the largest human AP-MS data set to date and mass-spectrometry experiments were performed in a uniform manner, making it a useful resource for development of data analysis techniques. In contrast to other human AP-MS data sets that focus on individual complexes or processes, our data set represents a broad survey of human proteins and their interactions. In addition, the data set represents several complexes with relatively high coverage (in terms of number of baits), such as the proteasome, and Eukaryotic Initiation Factors, and many other human protein complexes that are sparsely covered. Thus, our motivation is to develop a method that can identify probable protein complexes from a heterogeneous (in terms of coverage or sampling of protein complexes) data set. In our initial analysis of the data set, we focused on bait-prey interactions (spoke model) by using partial least-squares to predict the reproducibility of each prey protein observation based upon a training set of highly reproduced AP-MS experiments.¹² In the current work, we mine the data set more comprehensively by identifying prey-prey associations based upon co-occurrence profiles. We show that these prey-prey interactions are a valuable source of protein interactions and that integration of bait-prey and prey-prey associations yields improved network models of protein complexes. Importantly, we show that prey-prey associations may be identified from less well covered interactomes, such as the current human interactome. Since comprehensive efforts to map the human protein interactome are still in their early states, methods to mine protein-protein interactions from incomplete AP-MS data sets will be important tools for the foreseeable future.

Our framework exploits the quantitative features of AP-MS data to assign confidence scores to each bait-prey observation. To improve signal-to-noise ratio, the sparse matrix of scores are then transformed to eigenvalues using singular value decomposition. A similarity measure between the profiles of each pair of prey proteins is then to detect potential prey-prey associations, which are the starting point for the clustering and reconstruction of protein complexes, which we validate by reference to known complexes and annotations.

Finally, we integrate bait–prey and prey–prey associations and show that these integrated networks have improved biological coherence. Importantly, we are able to directly compare bait–prey and prey–prey interactions in terms of their biological coherence and show that prey–prey interactions are a rich source of protein–protein associations. We illustrate our approach using selected specific protein networks and show how novel interactions can be identified.

The principal contribution of our work is to show how large-scale AP-MS data sets may be data-mined for identification of protein–protein associations. Our analysis pipeline is generically applicable to large-scale AP-MS data sets, and we anticipate that as increasing volumes of AP-MS data are available it can be applied and used to discover novel protein interactions and elucidate the topology of protein interaction networks.

METHODS

Data Set

Data used to develop our approach is principally derived from a previously described human AP-MS data set, in which 832 single-step anti-FLAG immunoprecipitation experiments representing 384 human bait proteins (~50% of baits were replicated) identified 5269 distinct prey proteins.¹² The human bait proteins were selected based on association with diseases such as cancer and obesity. The data were generated as previously described,¹² except that all spectra were researched against an IPI human protein sequence database (version 3.31) (92012 sequences) using the Mascot mass-spectrometry search engine (version 2.1, Matrix Science; fixed modification: Carbamidomethyl (C), variable modification: Oxidation (M); peptide mass tolerance, 2 Da; fragment mass tolerance, 0.4 Da; missed cleavages, 2). Peptide and protein identifications were processed through Peptide and Protein Prophet⁴³ and imported to LabKey server⁴⁴ (version 10.3) for data management. For subsequent steps of the analysis, corresponding gene symbols, where available, were used to represent baits and preys.

Protein Identification False Discovery Rates

To assess the protein identification false discovery rate, we used decoy database searches. Using the same Mascot search parameters as above, we searched all data against a concatenated decoy human IPI protein sequence database (version 3.31) and computed the mean false discovery rate across all searches (3.63%). We also searched the complete data sets using another search-engine, MassMatrix⁴² (version 2.4.2, <http://www.massmatrix.net>) that provides the false discovery rate in terms of the % decoy hits when searched against concatenated decoy databases. The database and search parameters were the same as for the Mascot analysis, except for the following additional options: peptide length: 6–40 amino acid residues and score thresholds of 5.3 and 1.3 for the pp and pp_{tag} scores respectively. Proteins identified using MassMatrix were cross-referenced to those identified using Mascot. Of the set of 34383 bait–prey associations from the original Mascot searches, and for which a D-score was computed (see below), 82% were also identified with MassMatrix. The mean false discovery rate across this set is 4.52%.

Bait–Prey Confidence Score

The D-score (eq 1) is a previously described confidence score for AP-MS data that combines measures of abundance (spectral counts), specificity and reproducibility into a single score for each bait–prey observation.⁷

$$D_{ij} = \sqrt{\left(\frac{k}{\sum_{i=1}^p f_{ij}}\right)^p} X_{ij} \text{ for all } i, j \quad (1)$$

where k = total number of unique bait proteins; X_{ij} = total spectral counts for prey i from bait j ; $f_{ij} = \{1, 0\}$; p = number of replicates runs in which the prey protein is present. This generated a complete bait–prey D-score matrix (**D**) of dimensions (5269 proteins $\times$ 384 baits).

Contaminant Identification

False-positive proteins may occur in AP-MS experiments for different reasons. First, there are many proteins that occur frequently in AP-MS experiments as a result of nonspecific affinity. In the AP-MS data set used in this study, a set of 200 control AP-MS experiments (using HEK293 cells with vector alone, i.e., no bait) provide a data set for identifying “control” proteins. Since the D-score accounts for protein frequency of occurrence, these proteins typically have low D-scores. D-scores for the top 100 highly frequent control proteins in this study have median value equal to 3 (Supplementary Figure S1, Supporting Information). These were removed along with any protein with D-score < 3 from the initial data set (13688 bait–prey interactions with D-score < 3 were removed). A second type of contaminant results from cross-contamination at the experimental level between samples. Since the experimental protocol used to generate data used here¹² used gel-based separation of proteins prior to mass-spectrometry, we identified proteins with D-score > 10 co-occurring on the same gels or samples that were processed on the same date with different bait proteins as potential cross-contaminants and removed them from the data set (1550 bait–prey interactions were removed using this criterion).

Latent Semantic Analysis (LSA)

The D-score matrix (**D**) is necessarily sparse (1.25% nonzero values), since the frequency of most proteins is low (identified with few or unique baits). To identify relationships between proteins in this matrix, we made use of the principles of latent semantic analysis (LSA) which uses Singular Value Decomposition (SVD) to reduce matrix sparseness.²¹ SVD was applied to the rectangular bait–prey matrix (eq 2), and a dimensional reduction of the singular values was performed after identifying the optimum k -value which determines the degree of reduction.

$$\mathbf{D} = \mathbf{U} \sum \mathbf{V}^T \approx \mathbf{D}_k = \mathbf{U}_k \sum_k \mathbf{V}_k^T \quad (2)$$

D denotes the D-score bait–prey matrix ($m \times n$). **U** ($m \times m$ matrix) and **V** ($n \times n$ matrix) are orthogonal matrices with unit-length columns (i.e., $\mathbf{U}\mathbf{U}^T = \mathbf{I}$ and $\mathbf{V}^T\mathbf{V} = \mathbf{I}$) and Σ is a diagonal matrix containing the ordered singular values of rank $r = \min(m, n)$. **D_k** denotes the approximated latent semantic representation of the bait–prey matrix **D** where k denotes the selected number of singular values used in the approximation. This generated a final bait–prey D-score matrix (**D_k**) of dimensions (2242 proteins $\times$ 384 baits). LSA computations were performed using the Perl Data Language (PDL, <http://pdl.perl.org>).

Prey–Prey Similarity Score

For each bait protein, the bait vector is the vector of D-scores for all preys. Similarly, for each prey protein, the prey vector is the vector of all D-scores for all baits. LSA projects the vectors of original bait and prey vectors into lower dimensional semantic space. Transformed bait and prey vectors are the vectors of bait or prey eigenvalues. Cosine

similarity (eq 3) was used to compute pairwise similarities between prey–prey vectors from the bait–prey matrix. The similarity score for a pair of prey proteins a, and b is then the cosine similarity of their associated D-score vectors (0 indicates minimal similarity, 1 indicates maximal similarity). Hereafter the cosine similarity is referred to as the prey–prey score (PPS).

$$\begin{aligned} \text{sim}(a, b) &= \cos \alpha(a, b) \\ &= \frac{a \times b}{(|a| \times |b|)} \\ &= \frac{\sum_{i=1}^n a_i b_i}{\left(\sqrt{\sum_{i=1}^n a_i^2} \sqrt{\sum_{i=1}^n b_i^2} \right)} \end{aligned} \quad (3)$$

Analysis of Known Interactions

Bait–prey and prey–prey interactions were compared to CORUM version 2.0²² and BioGRID version 3.1.78²³ data sets. The complete set of CORUM mammalian protein complexes was used (2577 protein complexes; 4304 distinct gene Ids). Human protein interactions were extracted from BioGRID (402127 total interactions; 37910 human interactions). Data originating from our previous publication¹² was carefully removed from BioGRID prior to the current analysis. For each human protein interaction in BioGRID, a 2-hop neighborhood network was computed. Thus, for two pairs of interacting proteins A–B and B–C, the 2-hop network for protein A would include C. Bait–prey and prey–prey interactions in the current study were compared to the 2-hop network for corresponding proteins in BioGRID and the protein complexes in CORUM.

Protein–Protein Interaction False Discovery Rates

The mean False Discovery Rate (FDR) estimate was computed by generating a null distribution of prey–prey cosine similarities from a reshuffled bait–prey D-score matrix (random permutation of rows and columns independently). This was repeated 100 times to estimate the standard error of the mean FDR. The number of prey–prey cosine similarity values greater than a given Prey–Prey Similarity (PPS) threshold was used to estimate the number of false prey–prey interactions. The FDR estimate is then the ratio of the number of false prey–prey interactions to the number of prey–prey interactions discovered. This strategy is similar to the ones used to estimate the FDR of peptide identifications from Peptide Spectrum Match with the help of decoy databases (reviewed in ref 47), or to the resampling based methods used by Lavalley-Adam et al.⁴⁸ and Dazard et al.⁴⁹ to assess the FDR of Protein–protein Interactions (PPI) from Affinity Purification Mass Spectrometry data with the help of a control set of experiments. A controlled FDR of $7.58 \pm 0.02\%$ was achieved for the chosen PPS threshold of PPS = 0.75 (see Results section). Comparison plots of distributions of prey–prey cosine similarities in the original matrix and in the null matrix under the resampling scheme are shown in Supplementary Figure S5 (Supporting Information). We also looked at the FDR profile within a range of PPS, and found that the FDR was quite stable around 6–8% within the PPS range [0.5, 0.75] and decreasing (as expected) within the PPS range [0.75, 0.99].

Semantic Similarity

The semantic similarity is a metric for computing the similarity of Gene Ontology (GO) terms, their ancestors, or the descendants for two genes.²⁴ Semantic similarity uses information content (IC) as a measure of the specificity of a term and is quantified as the negative log likelihood,

$$IC = -\log p(c) \quad (4)$$

where $p(c)$ is the probability of occurrence of c in the GO structure. We used Resnik's max measure of similarity between two terms as the IC of their most informative common ancestor (MICA):

$$sim_{Res}(c_1, c_2) = IC(c_{MICA}) \quad (5)$$

This measure is effective in determining the information shared by the two terms.²⁵ Analysis was performed using the *GOsim* R package.²⁶

Protein Clustering and Visualization

Assembly of protein clusters from the prey–prey similarity matrix was performed using Pearson correlation and hierarchical clustering (using TIBCO Spotfire software). Enrichment analyses of clustered sets of proteins were performed using FuncAssociate 2.²⁷ FuncAssociate performs a Fisher's Exact Test analysis to identify GO terms and the results are corrected for multiple hypotheses via empirical resampling, and adjusted p -values computed for significance. Cytoscape (version 2.8) was used for biological network visualization.²⁸

RESULTS AND DISCUSSION

Initial Data Processing and Analysis

An overview of the data analysis workflow is shown in Figure 1. The workflow incorporates steps to address key issues in AP-MS data sets including the occurrence of nonspecific prey proteins and missing data due to under-sampling of protein complexes. Scores for both bait–prey and prey–prey protein associations are computed and used to construct protein networks.

Abundance and specificity of prey protein peptides in AP-MS experiments are important features that allow the confidence of prey proteins to be assessed. The D-score⁷ was used as the primary metric for each bait–prey observation, since it combines measures of abundance and specificity into a single score and gives weight to well-replicated and less frequent prey proteins.

A principal challenge in interpreting AP-MS data is the presence of nonspecific prey proteins. Typically, these are identified by their frequency across the data set or as proteins present in control experiments. Our previous study¹² analyzed ~200 empty vector control AP-MS experiments that were then used to define nonspecific prey proteins. We set a defined frequency threshold and used this threshold to filter out frequently occurring prey proteins.¹² In this analysis, we first observed that D-score values are typically very low for highly frequent proteins (since the D-score negatively weights frequency). Highly frequent proteins in the initial data set such as PRMT5 (present in 94% of the 995 AP-MS experiments in the data set) have low median D-scores (the median D-score for PRMT5 is 1.8 and 95% of baits for which PRMT5 is identified have D-score <10). The top 100 most frequent prey proteins in the data set have a median D-score of ~3.0 (Supplementary Figure S1, Supporting Information). We used this value as a threshold to remove any bait–prey interaction with D-score < 3. Thus, 13688 bait–prey interactions with D-score < 3 were removed from the data set. The final matrix of D-scores used for subsequent analyses has dimensions of 384 baits × 5269 prey proteins, with ~1.25% (25344) nonzero values (Supplementary Table S1, Supporting Information).

To better identify relations between associated proteins in the sparse bait–prey matrix (Figure 1A), we used singular value decomposition (SVD), to project the bait–prey matrix into a lower dimensional space, and assign eigenvalues to bait–prey pairs (Figure 1B). Application of SVD to the bait–prey matrix increases the overall similarity of prey protein vectors within the matrix as shown in Supplementary Figure S2 (Supporting Information). SVD in the form of latent semantic analysis was previously applied to detect analytical trends in the HUPO Plasma Proteome Project (HUPO PPP).²⁹ Our goal in applying SVD to the bait–prey matrix is to increase the probability of detecting biologically relevant associations between proteins in the matrix. The extent to which the matrix dimension is reduced is determined by the k -value in SVD (Figure 1B). A high k -value corresponds to a small reduction of matrix dimensions with possible retention of too much noise, whereas a small k -value may retain too little information from the original matrix. We estimated an appropriate k -value by plotting the singular value at the k^{th} rank vs k -value (Supplementary Figure S3, Supporting Information). A k -value of 150 was selected at the plateau in the graph so that the majority of significant information in the matrix is retained (singular values with rank $< k$). To identify relationships between prey proteins in the bait–prey matrix, we then generated a prey–prey similarity matrix (Figure 1C) by computing the similarity between each pair of preys in the matrix (cosine similarity). Although a similar methodology could also be applied to the bait–bait matrix, we focused on the prey–prey matrix (5269×5269 proteins), because it is much larger than the bait–bait matrix (384×384 proteins), and we reasoned that it would represent a richer resource for discovery of novel interactions and complexes. To identify similarities between prey proteins, we computed the cosine distance between each pair of prey proteins for all prey proteins occurring with 2 or more baits (2242 preys) (Figure 1C).

Benchmarking Bait–Prey and Prey–Prey Interactions

To gauge the biological relevance of the bait–prey and prey–prey interactions and their associated scores, we compared protein–protein pairs in our data set to known data sets and measured the degree to which protein pairs shared functional annotations. Since the bait–prey and prey–prey interactions identified using our framework represent putative physical associations between proteins, we used two complementary sources of protein interaction data as our benchmark sources. CORUM²² is a map of mammalian protein complexes curated from individual studies, whereas BioGRID²³ is a repository of physical protein interactions including data from high-throughput interaction studies. CORUM groups proteins according to protein complexes, whereas BioGRID represents protein interactions as protein–protein pairs. To benchmark versus BioGRID, we first computed the 2-hop neighborhood of each protein–protein pair in BioGRID. Since our data is derived from AP-MS experiments, representing complexes of physically associated proteins that may or may not directly interact, we compared our data to these neighborhoods of associated proteins rather than binary protein–protein pairs in BioGRID to ensure a representative comparison. For each bait–prey or prey–prey protein pair, we determined whether they co-occurred in a CORUM complex or within a BioGRID 2-hop network.

To compare bait–prey and prey–prey interactions to known interactions, we calculated log likelihood scores for the relative enrichment of BioGRID and CORUM known interactions in our data, as previously formulated.³⁰ Figure 2 shows the log likelihood of known interactions for bait–prey D-scores (Figure 2A) and prey–prey scores (Figure 2B). Although the absolute numbers of interactions overlapping between our data and BioGRID are higher than the overlap between our data and CORUM, CORUM interactions represented in our bait–prey or prey–prey data sets have significantly higher D-score or prey–prey scores, respectively. Ranked by prey–prey score, the 25th percentile CORUM prey–prey score is 0.44 whereas the 25th percentile BioGRID prey–prey score is 0.24. This is not unexpected,

since CORUM is derived from manually curated protein complexes whereas BioGRID includes high-throughput protein interaction studies. This also demonstrates the value of comparisons of protein interaction data to multiple sources; CORUM provides higher specificity with low sensitivity whereas BioGRID improves the sensitivity at the cost of lower specificity. Figure 2 also shows that the relative enrichment of BioGRID or CORUM interactions decreases as bait–prey D-score or prey–prey score decrease, showing that both scores provide some discrimination of true positive interactions from false. In the case of the bait–prey interactions, BioGRID and CORUM interactions decrease very sharply below D-score ~20 (corresponding to ~95th percentile of all of the D-scores, as determined by Sowa et al.⁷). In the case of the prey–prey interactions, BioGRID and CORUM interactions decrease more evenly across the range of prey–prey scores. The large predicted size of the human protein interactome,^{30,31} the incompleteness of “known” human protein interactions, as well as noise and context-sensitivity mean that intersections between experimental protein interaction data sets and known interactions tend to be small. For example, although BioGRID is a comprehensive source of available protein interaction data (~400000 total interactions; ~38000 human interactions),²³ 35% of the bait proteins used in our original AP-MS study¹² have 1 or fewer interacting protein in BioGRID (27% have no interactions at all). The overlap between bait–prey or prey–prey interactions and known interactions in the BioGRID set or in CORUM was 26.9% of bait–prey interactions (D-score >20) and 4.6% of prey–prey interactions (prey–prey score >0.75).

Biological Coherence of Bait–Prey and Prey–Prey Associations

To ascertain whether bait–prey and prey–prey protein pairs represent associations between proteins with related functions, and to benchmark the bait–prey and prey–prey scores, we analyzed functional annotations of associated proteins using the Gene Ontology (GO).³² We first observed that computing coannotation of GO terms for protein–protein pairs has low sensitivity, since many proteins, although biologically related may not be assigned the same term. We therefore used semantic similarity (SS) of GO terms which has proven to be robust measures of biological similarity for pairs or sets of genes.²⁴ Although there are multiple implementations of semantic similarity, here we use the Resnik max method,²⁵ since it was previously shown to perform best in a comparative analysis of semantic similarity metrics using large-scale protein–protein interactions.³³ Semantic similarity vastly increases the sensitivity of analyzing coannotations oversimple analysis of coannotated proteins pairs (Supplementary Tables S2 and S3, Supporting Information).

For both bait–prey and prey–prey associations, we reasoned that gene-ontology semantic similarity scores should be higher for bait–prey or prey–prey pairs with higher D-scores and prey–prey scores respectively. GO semantic similarity scores were computed for each bait–prey or prey–prey protein pair (Supplementary Tables S2 and S3, Supporting Information) and analyzed as follows. Bait-prey protein pairs were binned according to D-score: high (D-score >100), medium (100.00 < D-score > 20.00), low (20.0 < D-Score > 5.00), very low (D-score < 5.00) and a randomly selected set ($n = 1000$) of protein–protein pairs. Distributions of semantic similarity scores for bait–prey pairs are shown in Figure 3A for the molecular function GO ontology. The high and medium bins were found to be significantly higher than the random set (Wilcoxon test p -values: 1.5×10^{-8} , 3.3×10^{-6} , 0.59 for high, medium and low respectively). Similar trends were observed with the biological process and cellular compartment GO ontologies (Supplementary Figure S4, Supporting Information). We used these analyses along with the data shown in Figure 2A to calibrate the D-score threshold and therefore focused subsequent analyses on bait–prey pairs with D-score ≥ 20 .

Semantic similarity scores for all prey–prey associations were also computed and binned according to prey–prey scores (Supplementary Table S3, Supporting Information). Although not strictly monotonic, the log likelihood of known interaction enrichment with respect to

prey–prey score broadly increases as prey–prey score increases (Figure 2B). We therefore grouped the prey–preys into four bins spanning the prey–prey score (PPS) range high (PPS ≥ 0.75), medium ($0.75 > \text{PPS} \geq 0.50$), low ($0.50 > \text{PPS} \geq 0.25$) and very low ($\text{PPS} < 0.25$) and compared with semantic similarity of GO molecular function as shown in Figure 3B. A random set of 10000 prey–prey scores were generated to compute the significance test. Notably, the high scoring bin (PPS ≥ 0.75) enriches for interactions with higher semantic similarity in all 3 GO ontologies (Figure S3, Supporting Information), and the difference between the high bin and the random set was statistically significant (Wilcoxon Test p -values: 2.1×10^{-12} , 0.54 for high and medium respectively). These results show that both the D-score and prey–prey scores can be used to define sets of bait–prey or prey–prey pairs that are enriched for interactions with higher biological coherence, based upon their GO annotations. These analyses also provide a guide for selected subsets of interactions for further analysis, and as such we used bait–preys with D-score ≥ 20 and prey–preys with score ≥ 0.75 for building network models of protein complexes. Approximately 5.6% (1900) of bait–prey interactions and 7.7% (23000) of prey–prey interactions meet these criteria and were used in subsequent analyses.

Identification of Protein Complexes

Since AP-MS experiments identify cocomplexed proteins, rather than binary pairs of interacting proteins, we organized the data into more meaningful biological groupings by clustering sets of proteins with significant prey–prey scores. In addition, since the sets of high-scoring bait–prey and prey–prey interactions are large, and likely contain significant numbers of false positives, clustering provides a means of focusing on higher-likelihood associations of proteins. The prey–prey similarity matrix was hierarchically clustered as shown in Figure 4. We identified 107 clusters comprised of a total of 754 proteins (each cluster was required to have at least 3 proteins and all prey–prey associations > 0.9) as shown on the diagonal of the prey–prey matrix in Figure 4. For each of the 107 protein clusters, we identified significantly enriched GO categories, and for 43 of the 107 protein clusters, one or more significant ($p < 0.05$) GO annotation terms were identified (Supplementary Table S4, Supporting Information).

Network Models Incorporating Bait–Prey and Prey–Prey Interactions

As demonstrated in the previous sections, the global biological coherence of prey–prey interactions is similar to that of bait–prey interactions. Since matrix models of AP-MS data, in which all pairwise interactions are assumed to occur, are prone to false positives,¹⁶ we sought to selectively combine high-scoring bait–prey and prey–prey interactions into integrated network models. Protein clusters identified in Figure 4 were used as network seeds and extended by addition of selected high scoring bait–prey (D-score > 20) and prey–prey (PPS > 0.75) interactions. Four constructed network models, corresponding to Eukaryotic Initiation Factor (EIF) complexes, G-protein signaling and regulation, chromatin assembly factor complex and nucleosome regulation, and the proteasome are shown in Figure 5A–D, respectively, and were selected to illustrate the potential of combining bait–prey and prey–prey interactions for delineation of network topology and identification of new protein complex components and interactions.

Figure 5A shows a network model constructed by integrating bait–prey and prey–prey interactions corresponding to Eukaryotic Initiation Factor (EIF) complexes. Six bait proteins: EIF1B (also known as GC20), EIF2B1, EIF3S10, EIF4A2, EIF4A1, EIF4EBP1 and their associated prey proteins were integrated. Of particular interest is the separation of the cliques corresponding to EIF1/2/3 components (red shaded nodes) and EIF4 components (green nodes). A large number of prey proteins (D-score > 20) were found for EIF1B and EIF2B1 baits, including many ribosomal proteins, in line with known associations between

these components and ribosomes.³⁴ In contrast, EIF4A2 and EIF4A1 baits had relatively few high-scoring prey proteins as shown in Figure 5A. Of particular note, the LSM14A protein was identified uniquely with EIF4A2 bait and with high prey-prey interactions with EIF4G1 and PDCD4. LSM14A is a component of P-bodies, cytoplasmic foci that govern mRNA storage and degradation.³⁵ In line with previous findings, EIF4 components are present in P-bodies while other EIF components are conspicuously absent.³⁵ PDCD4 is uniquely associated with EIF4A2 and EIF4A1 baits in our data and is known to interact directly with EIF4A2, EIF4A1 and EIF4G1.³⁶

While EIF complexes are well represented in our data set, we also analyzed protein complexes with sparser coverage such as the networks in Figure 5B and C. Figure 5B shows a network of proteins involved in G-protein signaling. Members of the Rho family of GTPases (RHOB, RHOC) and the Ras superfamily of small GTP-binding proteins (RAC1, RAC2) are associated through their common bait, ARHGDIA. All of these proteins function in vesicular transport and endosomal signaling. An additional high-scoring interaction was observed with LYST, the lysosomal trafficking regulator. Endosomes may ultimately fuse with lysosomes for degradation of constituent molecules. In addition, LYST is associated with a rare lysosomal disorder, Chediak-Higashi syndrome, and thus the association with endosome signaling may potentially shed additional light on the disease mechanisms.

We also observed a cluster of prey-prey interactions corresponding to the Chromatin Assembly Factor (CAF-1) complex, and expanded this cluster of proteins into the network shown in Figure 5C. Several protein complexes with known functions were identified. The CAF-1 complex and Polycomb Recessive Complex 2 (PRC2) have related functions in chromatin metabolism and share components such as RBBP4. The recently identified MMS22L-TONSL complex³⁷ that mediates recombination mediated repair of replication forks and is comprised of MMS22L and TONSL (a.k.a. NFKBIL2) proteins was also identified. Other significant interactions between these proteins and the Tausled-like kinases (TLK1, TLK2) were also observed. TLK1 and TLK2 heterodimerize and are also involved in chromatin assembly.³⁸ Within this rich, overlapping network of complexes with related functions, we searched for other high-scoring prey-prey interactions that might be novel components. Two other proteins with significant prey-prey interactions with proteins in the chromatin modification network were observed. SMG6, a protein functioning in telomere regulation and nonsense-mediated mRNA decay was observed with high-scoring prey-prey interactions with several proteins (TLK1, 0.98; JARID2, 0.99). SMG6 is conserved (across eukaryotes) and has been found to physically interact with telomerase.³⁹ A second protein, UBN2, was recently identified as an ortholog of a yeast protein involved found in senescence associated chromatin foci.⁴⁰ The yeast ortholog of UBN2 interacts with the yeast ortholog of ASF1A, and both proteins function to create transcriptionally silent heterochromatin.⁴¹ Thus, integration of bait-prey and prey-prey interactions serves to identify proteins linked to known complexes and functions.

Finally, we constructed a network based on coverage of the proteasome complex in our data (Figure 5D). This network was constructed from four bait proteins and 16 prey proteins, highly enriched for components of the eukaryotic proteasome. Three of the baits used (PSMD6, PSMD10 AND PSMD13) function as 26S proteasome non-ATPase regulatory subunits of the proteasome and whereas one bait (PSMC4) is a 26S protease regulatory subunit. Although the structure of the eukaryotic proteasome is comparatively well understood, we note that the network model delineates sub components of the proteasome. For example, PSMC1, PSMC2 and PSMC5 cluster while PSMD subunits (PSMD8, PSMD12 and PSMD14) cluster separately. Although the coverage of protein-protein interactions is relatively sparse, we note that the alpha-type proteasome subunits, PSMA5 and PSMA6 cluster exclusively with PSMD subunits and not PSMC subunits, possibly

indicating intermediate forms of the proteasome comprised of specific sets of components. In summary, integrated network models as shown by these examples have the potential to yield novel components of protein complexes as well as further delineating the topology and substructure of networks.

We next compared the biological coherence of integrated (bait–prey and prey–prey) and matrix models of equivalent complexes. Matrix models for selected baits consisted of all pairwise protein interactions (bait–prey and prey–prey), whereas integrated models consisted of all bait–prey interactions with D-score >20 and all prey–prey interactions with score >0.75. Semantic similarity distributions were used to compare the integrated and matrix models of several protein complexes as shown in Figure 6. In addition to the Eukaryotic Initiation Factor (EIF) and Proteasomal complexes, we also analyzed data corresponding to four single baits (CTNNBIP1, VHL, REA and WDR8), the four most well replicated baits in the original study. For all four single baits and in the EIF complex, semantic similarity scores were significantly higher for integrated network models than for the matrix models. In only 1 case, the proteasome, the integrated and matrix model showed no significant difference between the semantic similarity distributions, suggesting that the biological coherence of integrated and matrix models of our proteasomal data is similar. This may be explained by the fact that most of the prey proteins identified by proteasome baits in our study are already known components of the proteasome, and so selection of prey–prey interactions with high PPS score for incorporation in the network model does not improve biological coherence over assuming that all preys interact with all other preys. In summary, these data show that selective integration of high-scoring bait–prey and prey–prey interactions can be used to generate protein network models with high biological coherence that can reveal new connections between proteins as well as new components of protein complexes.

CONCLUSIONS

We present a data-driven framework for analysis of large-scale interaction proteomics data that uses integrated computational techniques to derive additional value from these data sets in terms of novel protein–protein interactions. We used this framework to analyze a unique, systematically generated human AP-MS data set, in which complexes were determined for 384 disease-relevant bait proteins.¹² Specifically, our analysis integrates associations between the baits and their identified proteins (bait–prey), and associations that we detect among prey proteins (prey–prey), by analysis of the complete data matrix. Analyzed globally, high-scoring bait–prey and prey–prey are enriched for known interactions and interactions between proteins with related biological function. In addition, by integrating prey–prey and bait–prey interactions into network models, we increase the biological coherence of those networks. We also show that integrated networks of bait–prey and prey–prey interactions provide a basis for delineation of network topology and identification of new protein complex members.

A major motivation for our work is to develop methods that enable novel protein–protein associations to be derived from large-scale data sets. Since mapping the complete human protein interactome experimentally is such an enormous undertaking, approaches that can be used to identify novel interactions from existing data will continue to play a role in extending and refining the known human protein interactome. In addition, as previously observed,²² despite the rapid accumulation of large volumes of proteomics data, there has been surprisingly little reanalysis and evaluation of most proteomics data sets. The exceptions to this include unique data sets such as the yeast interaction proteomics data sets^{2–5} that have fueled much of the development of scoring algorithms as well as analysis of interaction networks.

Other studies with similar intention to the current work have primarily focused on the relatively complete yeast AP-MS data sets, or on AP-MS data sets that although not complete, are focused on specific complexes or sets of complexes.^{6,20} In addition, most methods that have been proposed for AP-MS data analysis are not appropriate to all data sets, for reasons of data size or completeness. The heterogeneity of current AP-MS data sets in terms of biological and analytical methodology (epitope tag, expression construct, organism, cell-type. etc.) along with the challenges associated with the data itself (missing data, noise, contaminant proteins) have hindered the development of widely applicable analysis tools. Exceptions to this include the SAINT algorithm where the explicit goal is to develop a widely applicable AP-MS analysis method.¹³

With these challenges in mind, we have constructed a computational framework for analysis of large-scale human AP-MS data. Noteworthy features of our framework include representation of protein observations using the D-score that takes into account spectral counts, overall frequency, and replication for each protein observation. Although other studies with similar intent have made use of present/absent calls, particularly in analysis of the yeast AP-MS data sets, quantitative values derived via label-free analysis provide an additional dimension for ranking and discrimination of true positive interactions.²⁰ In the latter study, Choi et al. used nested biclustering to group sets of baits or preys with similar quantitative spectral count levels. In our study, the data set is more heterogeneous both in terms of the actual baits used as well as the number of replicate AP-MS experiments per bait. For this reason, rather than use the spectral counts directly for each protein, we used the D-score as the starting point for the analysis, so that frequency and number of replicates as well as spectral counts are taken into account. Second, we address the issue of sparseness of the primary bait–prey matrix by first transforming the matrix of bait–prey scores through singular value decomposition. Singular value decomposition in the form of latent semantic analysis has previously been used to large-scale expression proteomics albeit with a different goal.²⁹ In the latter study, SVD was used to analyze heterogeneous plasma proteomics data sets acquired using different technologies in different laboratories. The method used by Klie et al represented protein observations as binary measures, presumably so that data from different instruments and different laboratories could be appropriately integrated. Regardless of whether quantitative values are used however, tools such as SVD, that address the sparseness and missing data challenges of proteomics data sets are essential. SVD was also previously applied to the assembly of protein interaction networks from more focused human AP-MS data, although in this case, SVD was applied to the problem of identifying clusters of proteins.⁶

To mitigate the problems of false positives in studying the large volume of prey–prey associations, we use semantic similarity measures of annotation to first calibrate our protein–protein scores (bait–prey and prey–prey) and then to select small high-scoring subsets of protein–protein interactions to study. A principal challenge of data-driven inference of human protein–protein interactions remains the lack of “gold-standard” protein interactions and complexes. Although curated annotations of mammalian protein complexes exist,²² these represent only a fraction of the total interactome. In the absence of gold-standard data sets, the semantic similarity measures provide a means of benchmarking protein–protein scores as shown here. Our method increases the amount of information in terms of protein–protein associations that may be gleaned from large-scale AP-MS studies. However, increasing the density of coverage (in terms of bait proteins used) will be the method of choice for truly defining protein complexes in the cell. Studies that iterate through a network, testing all proteins as baits in AP-MS experiments may provide the most detailed representations of protein complexes.⁸ The complexity of mapping the protein interactome, in terms of distinguishing complexes that share components was highlighted in an AP-MS study of chromatin remodeling complexes.⁶ This study reiterates the point that although

computational and statistical analyses may allow for the prediction of protein–protein interactions and network topology, a detailed map of overlapping protein complexes may ultimately only be achieved through high density experimental analyses.

Integration of other data may further help refine of protein network topology. Although protein quantifications are only loosely correlated with protein–protein interactions, integration of protein abundance values may be used to refine prediction of protein interaction. The recently developed PaxDb⁴⁵ is a cross-species database of protein quantifications, thus providing an independent source of proteome-wide abundance. In PaxDb, protein–protein interactions are used as a metric of consistency for quantitative proteomics studies, with the rationale that interacting proteins tend to have more similar levels of expression than randomly selected or noninteracting proteins.⁴⁵ Proteome-wide protein abundance values may also prove useful for normalizing the abundance values of proteins identified in AP-MS experiments. For example, a recent study used PaxDb values to account for the differing cellular abundance of proteins identified in AP-MS analysis of chromatin remodeling complexes.⁴⁶

Future work might therefore utilize the approach that we have described in conjunction with other proteome-wide information for further refinement of protein networks and protein interaction discovery.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Joseph Abraham and Parminder Kaur for productive scientific discussions. R.M.E. acknowledges funds from the Cleveland Foundation used in part to fund this study.

References

1. Saha S, Kaur P, Ewing RM. The bait compatibility index: computational bait selection for interaction proteomics experiments. *J Proteome Res.* 2010; 9(10):4972–81. [PubMed: 20731387]
2. Ho Y, Gruhler A, Heilbut A, Bader GD, Moore L, Adams SL, Millar A, Taylor P, Bennett K, Boutilier K, Yang L, Wolting C, Donaldson I, Schandorff S, Shewnarane J, Vo M, Taggart J, Goudreau M, Muskut B, Alfarano C, Dewar D, Lin Z, Michalickova K, Willems AR, Sassi H, Nielsen PA, Rasmussen KJ, Andersen JR, Johansen LE, Hansen LH, Jespersen H, Podtelejnikov A, Nielsen E, Crawford J, Poulsen V, Sørensen BD, Matthiesen J, Hendrickson RC, Gleeson F, Pawson T, Moran MF, Durocher D, Mann M, Hogue CW, Figgeys D, Tyers M. Systematic identification of protein complexes in *Saccharomyces cerevisiae* by mass spectrometry. *Nature.* 2002; 415(6868):180–3. [PubMed: 11805837]
3. Gavin AC, Bösch M, Krause R, Grandi P, Marzioch M, Bauer A, Schultz J, Rick JM, Michon AM, Cruciat CM, Remor M, Höfert C, Schelder M, Brajenovic M, Ruffner H, Merino A, Klein K, Hudak M, Dickson D, Rudi T, Gnau V, Bauch A, Bastuck S, Huhse B, Leutwein C, Heurtier MA, Copley RR, Edelmann A, Querfurth E, Rybin V, Drewes G, Raida M, Bouwmeester T, Bork P, Seraphin B, Kuster B, Neubauer G, Superti-Furga G. Functional organization of the yeast proteome by systematic analysis of protein complexes. *Nature.* 2002; 415(6868):141–7. [PubMed: 11805826]
4. Krogan NJ, Cagney G, Yu H, Zhong G, Guo X, Ignatchenko A, Li J, Pu S, Datta N, Tikuisis AP, Punna T, Peregrín-Alvarez JM, Shales M, Zhang X, Davey M, Robinson MD, Paccanaro A, Bray JE, Sheung A, Beattie B, Richards DP, Canadien V, Lalev A, Mena F, Wong P, Starostine A, Canete MM, Vlasblom J, Wu S, Orsi C, Collins SR, Chandran S, Haw R, Rilstone JJ, Gandhi K, Thompson NJ, Musso G, St Onge P, Ghanny S, Lam MH, Butland G, Altaf-Ul AM, Kanaya S, Shilatifard A, O'Shea E, Weissman JS, Ingles CJ, Hughes TR, Parkinson J, Gerstein M, Wodak SJ,

- Emili A, Greenblatt JF. Global landscape of protein complexes in the yeast *Saccharomyces cerevisiae*. *Nature*. 2006; 440(7084):637–43. [PubMed: 16554755]
5. Yu H, Braun P, Yildirim MA, Lemmens I, Venkatesan K, Sahalie J, Hirozane-Kishikawa T, Gebreab F, Li N, Simonis N, Hao T, Rual JF, Dricot A, Vazquez A, Murray RR, Simon C, Tardivo L, Tam S, Svrikapa N, Fan C, de Smet AS, Motyl A, Hudson ME, Park J, Xin X, Cusick ME, Moore T, Boone C, Snyder M, Roth FP, Barabási AL, Tavernier J, Hill DE, Vidal M. High-quality binary protein interaction map of the yeast interactome network. *Science*. 2008; 322(5898):104–10. [PubMed: 18719252]
 6. Sardi ME, Cai Y, Jin J, Swanson SK, Conaway RC, Conaway JW, Florens L, Washburn MP. Probabilistic assembly of human protein interaction networks from label-free quantitative proteomics. *Proc Natl Acad Sci USA*. 2008; 105(5):1454–9. [PubMed: 18218781]
 7. Sowa ME, Bennett EJ, Gygi SP, Harper JW. Defining the human deubiquitinating enzyme interaction landscape. *Cell*. 2009; 138(2):389–403. [PubMed: 19615732]
 8. Goudreaux M, D'Ambrosio LM, Kean MJ, Mullin MJ, Larsen BG, Sanchez A, Chaudhry S, Chen GI, Sicheri F, Nesvizhskii AI, Aebersold R, Raught B, Gingras AC. A PP2A phosphatase high density interaction network identifies a novel striatin-interacting phosphatase and kinase complex linked to the cerebral cavernous malformation 3 (CCM3) protein. *Mol Cell Proteomics*. 2009; 8(1): 157–71. [PubMed: 18782753]
 9. Glatter T, Wepf A, Aebersold R, Gstaiger M. An integrated workflow for charting the human interaction proteome: insights into the PP2A system. *Mol Syst Biol*. 2009; 5:237. [PubMed: 19156129]
 10. Gavin AC, Aloy P, Grandi P, Krause R, Boesche M, Marzioch M, Rau C, Jensen LJ, Bastuck S, Dümpelfeld B, Edelmann A, Heurtier MA, Hoffman V, Hoefert C, Klein K, Hudak M, Michon AM, Schelder M, Schirle M, Remor M, Rudi T, Hooper S, Bauer A, Bouwmeester T, Casari G, Drewes G, Neubauer G, Rick JM, Kuster B, Bork P, Russell RB, Superti-Furga G. Proteome survey reveals modularity of the yeast cell machinery. *Nature*. 2006; 440(7084):631–6. [PubMed: 16429126]
 11. Collins SR, Kemmeren P, Zhao XC, Greenblatt JF, Spencer F, Holstege FC, Weissman JS, Krogan NJ. Toward a comprehensive atlas of the physical interactome of *Saccharomyces cerevisiae*. *Mol Cell Proteomics*. 2007; 6(3):439–50. [PubMed: 17200106]
 12. Ewing RM, Chu P, Elisma F, Li H, Taylor P, Climie S, McBroom-Cerajewski L, Robinson MD, O'Connor L, Li M, Taylor R, Dharsee M, Ho Y, Heilbut A, Moore L, Zhang S, Ornatsky O, Bukhman YV, Ethier M, Sheng Y, Vasilescu J, Abu-Farha M, Lambert JP, Duewel HS, Stewart II, Kuehl B, Hogue K, Colwill K, Gladwish K, Muskat B, Kinach R, Adams SL, Moran MF, Morin GB, Topaloglou T, Figeys D. Large-scale mapping of human protein-protein interactions by mass spectrometry. *Mol Syst Biol*. 2007; 3:89. [PubMed: 17353931]
 13. Choi H, Larsen B, Lin ZY, Breitkreutz A, Mellacheruvu D, Fermin D, Qin ZS, Tyers M, Gingras AC, Nesvizhskii AI. SAINT: probabilistic scoring of affinity purification-mass spectrometry data. *Nat Methods*. 2011; 8(1):70–3. [PubMed: 21131968]
 14. Lavallée-Adam M, Cloutier P, Coulombe B, Blanchette M. Modeling contaminants in AP-MS/MS experiments. *J Proteome Res*. 2011; 10(2):886–95. [PubMed: 21117706]
 15. Skarra DV, Goudreaux M, Choi H, Mullin M, Nesvizhskii AI, Gingras AC, Honkanen RE. Label-free quantitative proteomics and SAINT analysis enable interactome mapping for the human Ser/Thr protein phosphatase 5. *Proteomics*. 2011; 11(8):1508–16. [PubMed: 21360678]
 16. Bader GD, Hogue CW. Analyzing yeast protein-protein interaction data obtained from different sources. *Nat Biotechnol*. 2002; 20(10):991–7. [PubMed: 12355115]
 17. Hart GT, Lee I, Marcotte ER. A high-accuracy consensus map of yeast protein complexes reveals modular nature of gene essentiality. *BMC Bioinform*. 2007; 8:236.
 18. Yu X, Ivanic J, Wallqvist A, Reifman J. A novel scoring approach for protein co-purification data reveals high interaction specificity. *PLoS Comput Biol*. 2009; 5(9):e1000515. [PubMed: 19779545]
 19. Geva G, Sharan R. Identification of protein complexes from co-immunoprecipitation data. *Bioinformatics*. 2011; 27(1):111–7. [PubMed: 21115439]

20. Choi H, Kim S, Gingras AC, Nesvizhskii AI. Analysis of protein complexes through model-based biclustering of label-free quantitative AP-MS data. *Mol Syst Biol.* 2010; 6:385. [PubMed: 20571534]
21. Deerwester S, Dumais St, Furnas GW, Landauer TK, Harshman R. Indexing by Latent Semantic Analysis. *J Am Soc Inf Sci.* 1990; 41(6):391–407.
22. Ruepp A, Waegele B, Lechner M, Brauner B, Dunger-Kaltenbach I, Fobo G, Frishman G, Montrone C, Mewes HW. CORUM: the comprehensive resource of mammalian protein complexes—2009. *Nucleic Acids Res.* 2010; 38(Database issue):D497–501. [PubMed: 19884131]
23. Stark C, Breitkreutz BJ, Chatr-Aryamontri A, Boucher L, Oughtred R, Livstone MS, Nixon J, Van Auken K, Wang X, Shi X, Regulj T, Rust JM, Winter A, Dolinski K, Tyers M. The BioGRID Interaction Database: 2011 update. *Nucleic Acids Res.* 2011; 39(Database issue):D698–704. [PubMed: 21071413]
24. Pesquita C, Faria D, Falcão AO, Lord P, Couto FM. Semantic similarity in biomedical ontologies. *PLoS Comput Biol.* 2009; 5(7):e1000443. [PubMed: 19649320]
25. Resnik, P. Using information content to evaluate semantic similarity in a taxonomy; *Proc 14th Int Joint Conf Artif Intell*; 1995. p. 448–453.
26. Fröhlich H, Speer N, Poustka A, Beissbarth T. GOSim—an R-package for computation of information theoretic GO similarities between terms and gene products. *BMC Bioinform.* 2007; 8:166.
27. Berriz GF, Beaver JE, Cenik C, Tasan M, Roth FP. Next generation software for functional trend analysis. *Bioinformatics.* 2009; 25(22):3043–4. [PubMed: 19717575]
28. Smoot ME, Ono K, Ruscheinski J, Wang PL, Ideker T. Cytoscape 2.8: new features for data integration and network visualization. *Bioinformatics.* 2011; 27(3):431–2. [PubMed: 21149340]
29. Klie S, Martens L, Vizcaíno JA, Côté R, Jones P, Apweiler R, Hinneburg A, Hermjakob H. Analyzing large-scale proteomics projects with latent semantic indexing. *J Proteome Res.* 2008; 7(1):182–91. [PubMed: 18047271]
30. Ramani AK, Bunescu RC, Mooney RJ, Marcotte EM. Consolidating the set of known human protein-protein interactions in preparation for large-scale mapping of the human interactome. *Genome Biol.* 2005; 6(5):R40. [PubMed: 15892868]
31. Stumpf MP, Thorne T, de Silva E, Stewart R, An HJ, Lappe M, Wiuf C. Estimating the size of the human interactome. *Proc Natl Acad Sci USA.* 2008; 105(19):6959–64. [PubMed: 18474861]
32. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, Davis AP, Dolinski K, Dwight SS, Eppig JT, Harris MA, Hill DP, Issel-Tarver L, Kasarskis A, Lewis S, Matese JC, Richardson JE, Ringwald M, Rubin GM, Sherlock G. The Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nat Genet.* 2000; 25(1):25–9. [PubMed: 10802651]
33. Xu T, Du L, Zhou Y. Evaluation of GO-based functional similarity measures using *S. cerevisiae* protein interaction and expression profile data. *BMC Bioinform.* 2008; 9:472.
34. Pestova TV, Kolupaeva VG, Lomakin IB, Pilipenko EV, Shatsky IN, Agol VI, Hellen CU. Molecular mechanisms of translation initiation in eukaryotes. *Proc Natl Acad Sci USA.* 2001; 98(13):7029–36. [PubMed: 11416183]
35. Eulalio A, Behm-Ansmant I, Schweizer D, Izaurralde E. P-body formation is a consequence, not the cause, of RNA-mediated gene silencing. *Mol Cell Biol.* 2007; 27(11):3970–81. [PubMed: 17403906]
36. Yang HS, Jansen AP, Komar AA, Zheng X, Merrick WC, Costes S, Lockett SJ, Sonenberg N, Colburn NH. The transformation suppressor Pdc4 is a novel eukaryotic translation initiation factor 4A binding protein that inhibits translation. *Mol Cell Biol.* 2003; 23(1):26–37. [PubMed: 12482958]
37. Piwko W, Olma MH, Held M, Bianco JN, Pedrioli PG, Hofmann K, Pasero P, Gerlich DW, Peter M. RNAi-based screening identifies the Mms22L-Nfkbil2 complex as a novel regulator of DNA replication in human cells. *EMBO J.* 2010; 29(24):4210–22. [PubMed: 21113133]
38. Groth A, Lukas J, Nigg EA, Silljé HH, Wernstedt C, Bartek J, Hansen K. Human Tousled like kinases are targeted by an ATM- and Chk1-dependent DNA damage checkpoint. *EMBO J.* 2003; 22(7):1676–87. [PubMed: 12660173]

39. Redon S, Reichenbach P, Lingner J. Protein RNA and protein protein interactions mediate association of human EST1A/SMG6 with telomerase. *Nucleic Acids Res.* 2007; 35(20):7011–22. [PubMed: 17940095]
40. Banumathy G, Somaiah N, Zhang R, Tang Y, Hoffmann J, Andrade M, Ceulemans H, Schultz D, Marmorstein R, Adams PD. Human UBN1 is an ortholog of yeast Hpc2p and has an essential role in the HIRA/ASF1a chromatin-remodeling pathway in senescent cells. *Mol Cell Biol.* 2009; 29(3): 758–70. [PubMed: 19029251]
41. Zhang R, Poustovoitov MV, Ye X, Santos HA, Chen W, Daganzo SM, Erzberger JP, Serebriiskii IG, Canutescu AA, Dunbrack RL, Pehrson JR, Berger JM, Kaufman PD, Adams PD. Formation of MacroH2A-containing senescence-associated heterochromatin foci and senescence driven by ASF1a and HIRA. *Dev Cell.* 2005; 8(1):19–30. [PubMed: 15621527]
42. Xu H, Freitas MA. MassMatrix: A database search program for rapid characterization of proteins and peptides from tandem mass spectrometry data. *Proteomics.* 2009; 9(6):1548–1555. [PubMed: 19235167]
43. Nesvizhskii AI, Keller A, Kolker E, Aebersold R. A statistical model for identifying proteins by tandem mass spectrometry. *Anal Chem.* 2003; 75:4646–58. [PubMed: 14632076]
44. Nelson EK, Piehler B, Eckels J, Rauch A, Bellew M, Hussey P, Ramsay S, Nathe C, Lum K, Krouse K, Stearns D, Connolly B, Skillman T, Igra M. LabKey Server: an open source platform for scientific data integration, analysis and collaboration. *BMC Bioinform.* 2011; 12:71.
45. Wang M, Weiss M, Simonovic M, Haertinger G, Schrimpf SP, Hengartner MO, von Mering C. PaxDb, a Database of Protein Abundance Averages Across All Three Domains of Life. *Mol Cell Proteomics.* 2012; 10.1074/mcp.O111.014704
46. Tsai Y-C, Greco TM, Boonmee A, Miteva Y, Cristea IM. Functional proteomics establishes the interaction of SIRT7 with chromatin remodeling complexes and expands its role in regulation of RNA polymerase I transcription. *Mol Cell Proteomics.* 2012; 11:M111.015156.
47. Nesvizhskii AI. A survey of computational methods and error rate estimation procedures for peptide and protein identification in shotgun proteomics. *J Proteomics.* 2010; 73(11):2092–123. [PubMed: 20816881]
48. Lavalley-Adam M, Cloutier P, Coulombe B, Blanchette M. Modeling contaminants in AP-MS/MS experiments. *J Proteome Res.* 2011; 10(2):886–95. [PubMed: 21117706]
49. Dazard JE, Saha S, Ewing RM. ROCS: A reproducibility index and confidence score for interaction proteomics. *BMC Bioinform.* 2012; 13(1):128.

Figure 1.

Data analysis work-flow. (A) D-score⁷ matrix with score for each pair of bait–prey proteins. (B) D-score matrix approximation using singular value decomposition. (C) Pairwise cosine similarity computed for each vector of prey scores. (D) Topological overlap matrix created using hierarchical clustering to group prey proteins with similar prey–prey profiles. (E) Spoke, matrix and integrated models for a hypothetical bait (B1) and 4 prey proteins (P1–P4). Integrated model incorporates selected edges from spoke and matrix models (solid lines represent bait–prey interactions, dotted lines represent prey–prey interactions). Figure 1B adapted with permission from Klie et al. (*J. Proteome Res.* **2008**, *7*, 182–119). Copyright 2008 American Chemical Society

Figure 2.

Comparison with known interactions. (A) Bait–prey interactions compared to known interactions. Bait–prey protein pairs within shared 2-hop BioGRID network or present in same CORUM complex counted as “known”. Log likelihood of relative enrichment of known vs unknown interactions computed for each D-score threshold. (B) As in A, with log likelihood computed for each prey–prey protein pair and corresponding prey–prey interaction score threshold.

Figure 3.

GO semantic similarity distributions and protein interaction scores. (A) Box plots represent distributions of semantic similarity scores for bait-prey protein pairs binned according to D-scores (4 bins) for molecular function gene ontologies. Bins 1–4 represent D-score bins of D-score > 100, 100 > D-score > 20, 20 > D-score > 5 and D-score < 5, respectively. Bins 1 and 2 have statistically significantly (*) higher semantic similarity than random set (R) (Wilcoxon Test p -values are 1.5×10^{-8} and 3.3×10^{-6}). (B) Box plots represent distributions of semantic similarity scores for prey-prey protein pairs binned according to prey-prey scores (4 bins) for molecular function (MF) gene ontologies. Bins 1–4 represent prey-prey scores (PPS) of 0.75, 0.75 > PPS > 0.5, 0.5 > PPS > 0.25 and PPS < 0.25, respectively. Bin 1 has significantly (*) higher semantic similarity than random set (R) of prey-prey similarity scores (Wilcoxon Test P value is 2.1×10^{-12}). Median of each distribution is represented by a horizontal bar in each box plot.

Figure 4.

Topological overlap matrix of prey-prey interactions. The prey-prey similarity matrix (2242×2242) was used to create a topological overlap matrix using hierarchical clustering to group preys with similar prey-prey vectors. Red areas indicate prey-prey vector similarity and diagonal shows 107 protein modules. Selected modules are labeled with significant GO terms.

Figure 5.

Selected network models integrating high-scoring bait-prey and prey-prey edges. (A) Eukaryotic Initiation Factor (EIF) bait proteins (EIF1B, EIF3S10, EIF2B1, EIF4EBP1, EIF4A1 and EIF4A2) and their associated prey proteins (D-score >20) were combined with associated high-scoring prey-prey interactions. EIF1/2/3, EIF4 and EIF2 components are represented in red, green and gray nodes respectively. (B) G-protein signaling complex. ARHGDI, RAC1 and RAC2; GTPase, RHOB and RHOC are identified preys with extensive interconnectivity (high prey-prey scores). (C) Chromatin metabolism/CAF1 complex, PRC2 complex and MMS22L-TONSL complex are shown in green, red and pink color respectively. (D) Proteasome complex, four bait proteins corresponding to proteasome, PSMD6, PSMD10, PSMD13 and PSMC4 and 16 prey proteins forming highly enriched eukaryotic proteasome complex. The gray nodes correspond to proteasome core complex subunit (PSMA5/6); green nodes to 26S proteasome non-ATPase regulatory subunit (PSMD*); red nodes represents 26S protease regulatory subunits (PSMC*). Solid edges indicate bait-prey interactions and dashed edges indicate prey-prey interactions, edge thickness indicates bait-prey or prey-prey scores appropriately (thicker edge/higher score).

Figure 6.

Biological coherence of integrated and matrix network models of protein complexes. Semantic similarity (GO biological process) distributions of protein-protein interactions in Integrated (I) and Matrix (M) models. Networks were constructed for data from 4 single baits (CTNNBIP1, VHL, REA and WDR8) and two models incorporating multiple baits, Proteasome complex (P) and Eukaryotic Initiation Factor (E) complexes. In 5 cases (CTNNBIP1, VHL, REA, WDR8, E), the Integrated model shows higher median semantic similarity than the Matrix model, and in 4 of these cases (CTNNBIP1, VHL, WDR8, E), the difference between Integrated and Matrix models is significant (Wilcoxon Test p -values < 0.1).