Broad functional profiling of fission yeast proteins using phenomics and machine learning

María Rodríguez-López[1]*, Nicola Bordin[2]*, Jon Lees[2], Harry Scholes[2], Shaimaa Hassan[1], Quentin Saintain[1], Stephan Kamrad[1], Christine Orengo[2+], Jürg Bähler[1+]

1 University College London, Institute of Healthy Ageing and Department of Genetics, Evolution & Environment, London, United Kingdom.

2 University College London, Institute of Structural and Molecular Biology, London, United Kingdom.

* Equal contribution

[+] Equal senior/corresponding authors

Keywords: Cell phenotype, Protein function, *Schizosaccharomyces pombe,* Oxidative stress, Computational prediction, Gene Ontology, Functional genomics, Cellular ageing, Unknown protein, Machine learning, CATH FunFams, CATHPredictGO

Abstract

Many proteins remain poorly characterized even in well-studied organisms, presenting a bottleneck for research. We applied phenomics and machine-learning approaches with *Schizosaccharomyces pombe* for broad cues on protein functions. We assayed colony-growth phenotypes to measure the fitness of deletion mutants for all 3509 non-essential genes in 131 conditions with different nutrients, drugs, and stresses. These analyses exposed phenotypes for 3492 mutants, including 124 mutants of 'priority unstudied' proteins conserved in humans, providing varied functional clues. For example, over 900 new proteins affected the resistance to oxidative stress. Phenotype-correlation networks suggested roles for poorly characterized proteins through 'guilt by association' with known proteins. For complementary functional insights, we predicted Gene Ontology (GO) terms using machine learning methods exploiting protein-network and protein-homology data (NET-FF). We obtained 56,594 high-scoring GO predictions, of which 22,060 also featured high information content. Our phenotype-correlation data and NET-FF predictions showed a strong concordance with existing PomBase GO annotations and protein networks, with integrated analyses revealing 1,675 novel GO predictions for 783 genes, including 47 predictions for 23 priority unstudied proteins. Experimental validation identified new proteins involved in cellular ageing, showing that these predictions and phenomics data provide a rich resource to uncover new protein functions.

Introduction

Most biomedical research focuses on genes that are already well-studied, a situation which has changed little in the past 40 years of biochemical, cell biological and genetic investigations (1–9). The cellular functions of many genes thus remain poorly characterized or unknown, even with the availability of whole-genome sequences (10). This disparity is reflected in a strong bias in citations and publications which has remained constant over the past two decades (2,3). A synthetic bacterium with a minimal genome contains 473 genes, defining a bare essential for life (11); remarkably, 149 (~32%) of these most vital genes have unknown cellular roles. Moreover, a recent proteome survey across the evolutionary range finds that 38.4% of the identified proteins are not associated with any biological process, including 22.9% of the 100 most abundant proteins of each species (12). These examples highlight how much of the 'dark proteome' awaits discovery and the need to characterize gene function in an unbiased manner. To fully understand cells, we ought to know what all parts do and how they contribute to biological systems and disease.

The fission yeast, *Schizosaccharomyces pombe*, provides a powerful platform to interrogate eukaryotic gene function owing to its relative simplicity, well-annotated genome, deletion-mutant libraries for high-throughput assays, and genetic tractability under tightly controlled conditions (13). The *S. pombe* genome annotation contains 5134 protein-coding genes, 3624 (70.6%) of which are conserved in metazoa. Unlike budding yeast, fission yeast did not undergo any genome duplication, so there is less gene redundancy and mutations are more likely to result in phenotypes (14). Moreover, unlike in mammalian cells, ~90% of the genes are expressed under standard growth conditions (15), facilitating analyses of their functions at organism level. Fission yeast resembles mammalian cells in many respects (e.g., symmetrical cell division, chromatin and RNAi pathways, centromeres and replication origins), and thus provides a complementary model system to the more widely studied budding yeast (16). As in other model organisms (17), genes are being associated with cellular functions and phenotypes at higher rates in fission yeast than they are in humans. Fission yeast researchers, supported by the model organism database PomBase (18,19), have experimentally characterized 2498 proteins in publications, and the biological functions of 1916 additional proteins can be reliably inferred from orthologs in other organisms (status March 2023). Systematic genetic screens have been published for several, mostly well-studied conditions, including DNA damage (20), caffeine tolerance (21), catalase expression (22), mycelial development (23), aneuploid viability (24), autophagy (25), heavy-metal tolerance (26), cell-cycle progression and cell shape (27), mitotic competence (28), TORC1 inhibition (29–31), respiratory growth (32–34), and chronological lifespan (35,36). Nevertheless, the numbers of entirely unknown proteins are hardly decreasing over time; so without fresh approaches and resources directed towards unknown proteins, it will take dozens of years to uncover their functions (8). Empirical data show that genes that are linked with biological processes and functions in model organisms

become characterized much more readily in other organisms (3,6,8,37,38). Thus, initial dedicated efforts to link unstudied proteins with cellular functions and biological processes are key to trigger deeper follow-on analyses of these proteins.

The concept of protein function is somewhat imprecise and hierarchical, involving several layers of detail (39). PomBase labels *S. pombe* genes as 'unknown' if they cannot be associated with any informative, high-level GO Biological-Process terms (40), applying strict criteria based on functional assays for mutant phenotypes, protein localization and interactions in *S. pombe*, or on corresponding functional assays in other organisms to infer the function (8). *S. pombe* encodes 641 'unknown' genes (PomBase, status March 2023). There are of course many more poorly characterized proteins for which little meaningful information is available because the associated annotations are either incorrect or too general to be informative. Among these 643 unknown proteins, many are apparently found only in the fission yeast clade, but 380 are more widely conserved. Among these conserved proteins, 135 'priority unstudied' proteins show clear human orthologs, and our analysis using Orthologous Matrix (41) reveal broad and strong conservation profiles across 100 metazoans. These priority unstudied proteins have not been directly studied in any organism but can be assumed to have pertinent biological roles conserved over 500 million years of evolution. Notably, 49 (~36%) of the 135 unstudied proteins conserved from fission yeast to humans are not conserved in budding yeast, the best characterized eukaryote.

Initial efforts to learn about unknown genes need to be broad and exploratory, because hypothesis-driven research cannot target proteins for which nothing is known. Many genes may remain unknown because they are not required under benign laboratory conditions. Phenomics seeks to uncover gene function by rigorously identifying all phenotypes associated with gene mutants under many different conditions (13,42). Microbes like yeast are particularly suitable for such high-throughput phenotyping. For example, while only 34% of all budding yeast deletion mutants display growth phenotypes under a standard condition, 97% show phenotypes in specific chemical or environmental perturbations (43). Such broad genetic screens provide a basis to understand the function of unstudied genes. Besides offering direct functional clues about biological processes being affected by specific genes, they also give rich phenotypic signatures reflecting systemic responses to genetic perturbations in different conditions. Clustering and network analyses can point to the functions of unstudied genes that show similar phenotypic signatures to known genes, a principle called 'guilt by association' (44,45).

Besides phenomics, advances in machine learning now provide potent opportunities for unbiased functional predictions based on available data associated with genes (46–48). Over the past decade, various approaches have been independently benchmarked to determine the most useful strategies. Recent evaluations highlight the importance of protein-sequence homology as a key

component of successful methods (49,50). Future trends are likely to see more important contributions from deep-learning feature extraction on protein sequences (51). A particularly successful homology-based strategy in the last round of the Critical Assessment of Functional Annotations, CAFA4, exploits CATH protein functional families (FunFams) and has consistently been highly ranked for precision in Gene Ontology (GO) Molecular Function and Biological Process prediction (CAFA1 (52), CAFA2 (53), CAFA3 (49) and CAFA4 (ISMB2020, personal communication from organizers). However, homology-based data are not the only useful source for functional prediction, and combining other data types with machine learning can help to maximize performance (54). For GO Biological Process prediction, network data can help in achieving high prediction accuracy (e.g., biological pathways, genetic interactions, protein interactions). An ongoing challenge remains how to process this type of data contained in large, complex networks in order to combine it with more structured data (e.g., CATH FunFam inferred GO annotations). One method that shows some success is the Multi-Modal Auto Encoder (MMAE), which processes network data to a lower dimensional latent space that can then be integrated with other data types more effectively (55).

Here we combine complementary phenomics data, predicted annotations from integrated machine-learning methods as well as curated and experimental PomBase annotations to provide diverse clues about gene functions and functional associations in fission yeast. Using these potent wet and dry methods, we obtained 103,520 quantitative phenotype datapoints for 3,492 non-essential genes across 131 diverse conditions. By applying new machine-learning methods, we also report a high-confidence set of 1,675 novel GO term associations for 783 genes, derived by combining our phenotype data and functional predictions with PomBase annotations. This rich and wide-ranging functional information for *S. pombe* proteins, including the priority unstudied proteins and many poorly characterized proteins, provides a rich framework and testable hypotheses for follow-on studies.

Results and Discussion

Overview of the study

We performed colony-based phenotyping of the deletion mutants for all non-essential *S. pombe* genes in response to different nutrient, drug and stress conditions. This phenomics dataset provides both functional clues and insights into functional relationships from phenotype correlations. In parallel, we developed a meta-predictor (NET-FF) that combines established methods for exploiting protein-network and protein-family data to predict new GO associations for *S. pombe* genes. We then integrated our phenomics data and NET-FF predictions with the gold-standard annotations in PomBase, reflecting extensive experimental and curation work (18,19), to validate our data and identify novel protein functions and associations. Figure 1 provides an overview of this study.

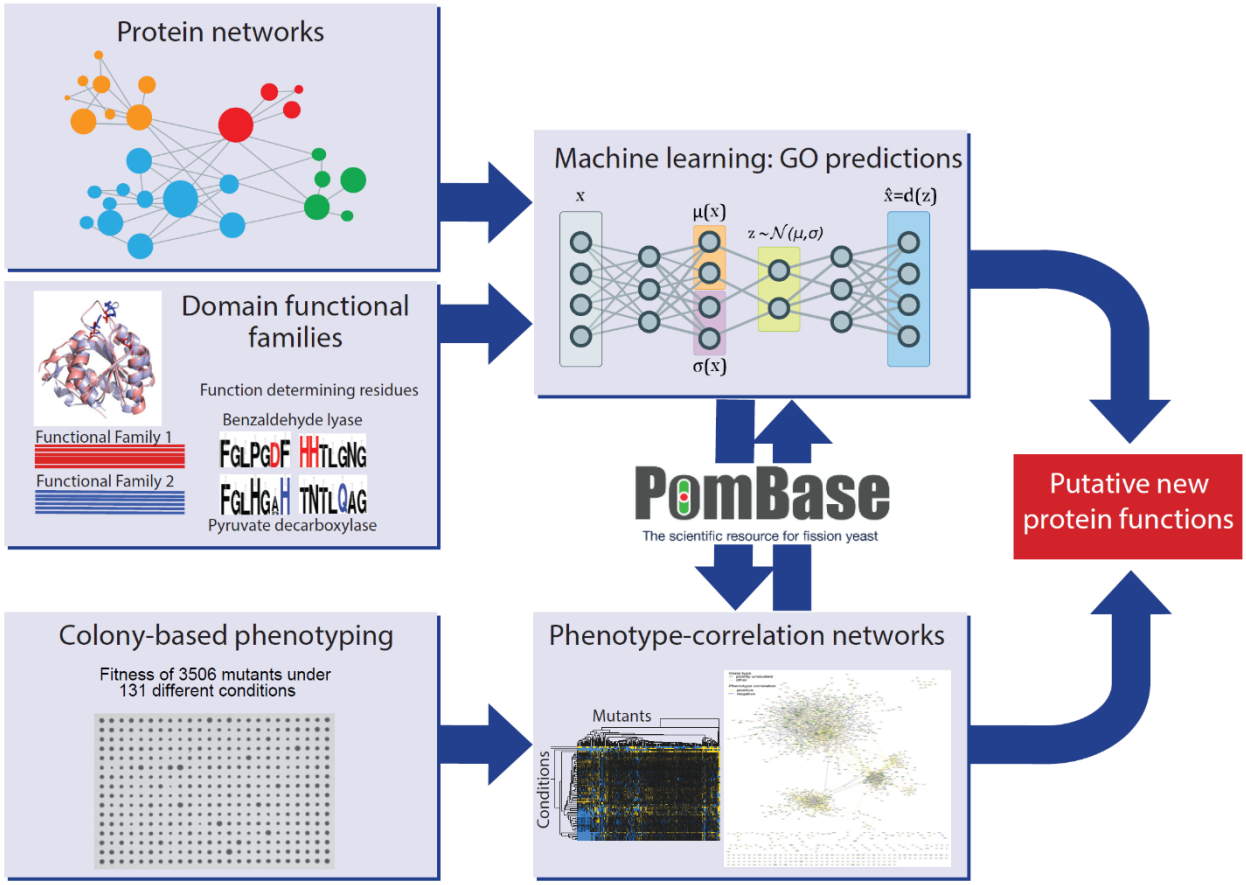

Figure 1. Schematic overview of study design.
Experimental and computational tasks along with the relationships between the different aspects of the work to generate new information on protein function.

Colony-based phenotyping of deletion mutants in multiple conditions

We combined two *S. pombe* deletion-mutant libraries (56,57) and crossed out all auxotrophy mutations, allowing us to screen 3,509 non-essential genes in prototroph strains that only contained the query mutants. We measured colony growth as a fitness proxy under different conditions using *pyphe,* our python package for phenomics analyses (58,59). The mutant strains were arrayed in 1536-colony format, including a 384-colony grid of wild-type strains as control. We assayed the deletion mutants in response to 131 diverse environmental conditions, both benign and stressful, including different nutrients and drugs as well as oxidative, osmotic, heavy-metal, protein-homeostasis and DNA-metabolism stresses. We also assayed some combined conditions which can reveal additional phenotypes through non-additive effects that are not evident from single conditions (60). For some drugs and stresses, we assayed both low and high doses, in which wild-type cell growth is normal or inhibited, respectively, to uncover both sensitive or resistant mutants. In total, these assays produced 2,832,384 data points, providing 103,520 phenotypes across 3,492 mutants (Supplementary file 1). These data exceed all the combined annotated phenotypes available in PomBase, currently encompassing 92,681 gene-phenotype associations of which 59,382 are unique (status March 2023).

Growth phenotypes in standard conditions

We first assayed fitness differences in standard, benign conditions: growth in rich or minimal medium at 32°C. We looked for mutants showing a significant difference in colony growth relative to wild-type control cells, applying a threshold of $p < 0.05$ after Benjamini-Hochberg correction and a colony-size difference of ≥5% compared to wild-type. Among the 3,509 deletion mutants tested, 732 and 760 mutants grew slower than wild-type cells in rich and minimal media, respectively, while 265 and 22 mutants grew faster in the same media (Figure 2A; Figure 2-figure supplement 1A; Supplementary file 1). The two media showed substantial overlap but also distinct responses. As an example, cells deleted for the transcription factor Fil1, which regulates the response to amino acid starvation, showed slower growth in minimal medium, as reported previously (61), but faster growth in rich medium. The differences between the two media were also reflected in distinct functional enrichments (Figure 2-figure supplement 1B). For example, mutants that grew slowly only in minimal medium were enriched in amino-acid biosynthesis and sulphate assimilation, reflecting auxotrophies, while mutants that grew slowly only in rich medium were enriched in the maintenance and fidelity of DNA replication, possibly reflecting an increased need for quality control in rapidly proliferating cells. Among the mutants analysed, we paid particular attention to those deleted for so-called ‘priority unstudied’ genes, which are widely conserved from fission yeast to human but remain entirely uncharacterized (8). We could phenotype 124 mutants of the priority unstudied genes, 13 of which grew slower in either or both media while 11 grew faster

(Figure 2A; Figure 2-figure supplement 1A). This finding shows that even under standard conditions some of these mutants readily reveal fitness phenotypes.

We compared our results with the relevant fission yeast phenotype ontology (FYPO) terms annotated in PomBase (19,62), including ‘increased cell population growth’ (253 mutants), and ‘decreased vegetative cell population growth’ or ‘slow vegetative cell population growth’ (1043 mutants). Overall, our screen confirmed 426 (40.8%) of the previously annotated slow-growth mutants but only 23 (9%) of the annotated fast-growth mutants. Besides technical discrepancies, the mutants showing different growth behaviour in the different assays may reflect differences in the type of mutant, between solid and liquid media, and/or between mutants growing in isolation or in pools given that a previous growth screen involved parallel mutant profiling by Bar-seq (35). Our phenotype data also showed good agreement with our recent study of 238 coding-gene mutants that have been phenotyped along with non-coding RNA mutants (59). Taken together, we identified 940 mutants that grow more slowly and 265 mutants that grow more rapidly in standard media, including 552 and 21 mutants that grow more slowly or rapidly, respectively, in both media.

Phenotypes in diverse nutrient, drug, and stress conditions

We then assayed fitness differences of the deletion mutants in the presence of various stresses or other treatments, relative to the same mutants growing in standard conditions and normalized for wild-type growth. Because only four repeats were measured per condition, fewer than for the benign conditions, we did not use p-values but applied a more stringent threshold to identify sensitive (slowly growing) or resistant (rapidly growing) mutants, using an effect-size difference of ≥10% compared to the standard condition in the same medium. Overall, we could measure colony growth for 3,506 mutants, 3,492 (99.6%) of which showed phenotypes in at least one condition, including 55,577 sensitive and 47,943 resistant phenotypes (Figure 2B; Supplementary file 1). Thirty-eight mutants displayed many different phenotypes in 100 to 118 conditions (Supplementary file 1), pointing to pleiotropic genes that exert multiple functions. All 124 mutants of the priority unstudied genes produced phenotypes in at least one condition, with an average of 19.4 phenotypes per gene compared to an average of 29.9 phenotypes per gene for all other mutants (Figure 2B). Moreover, all 129 conditions tested caused phenotypes in 118 to 2,761 mutants (Figure 2C; Supplementary file 1). Thus, we obtained rich and diverse phenotype data for nearly all mutants tested, including those of priority unstudied genes.

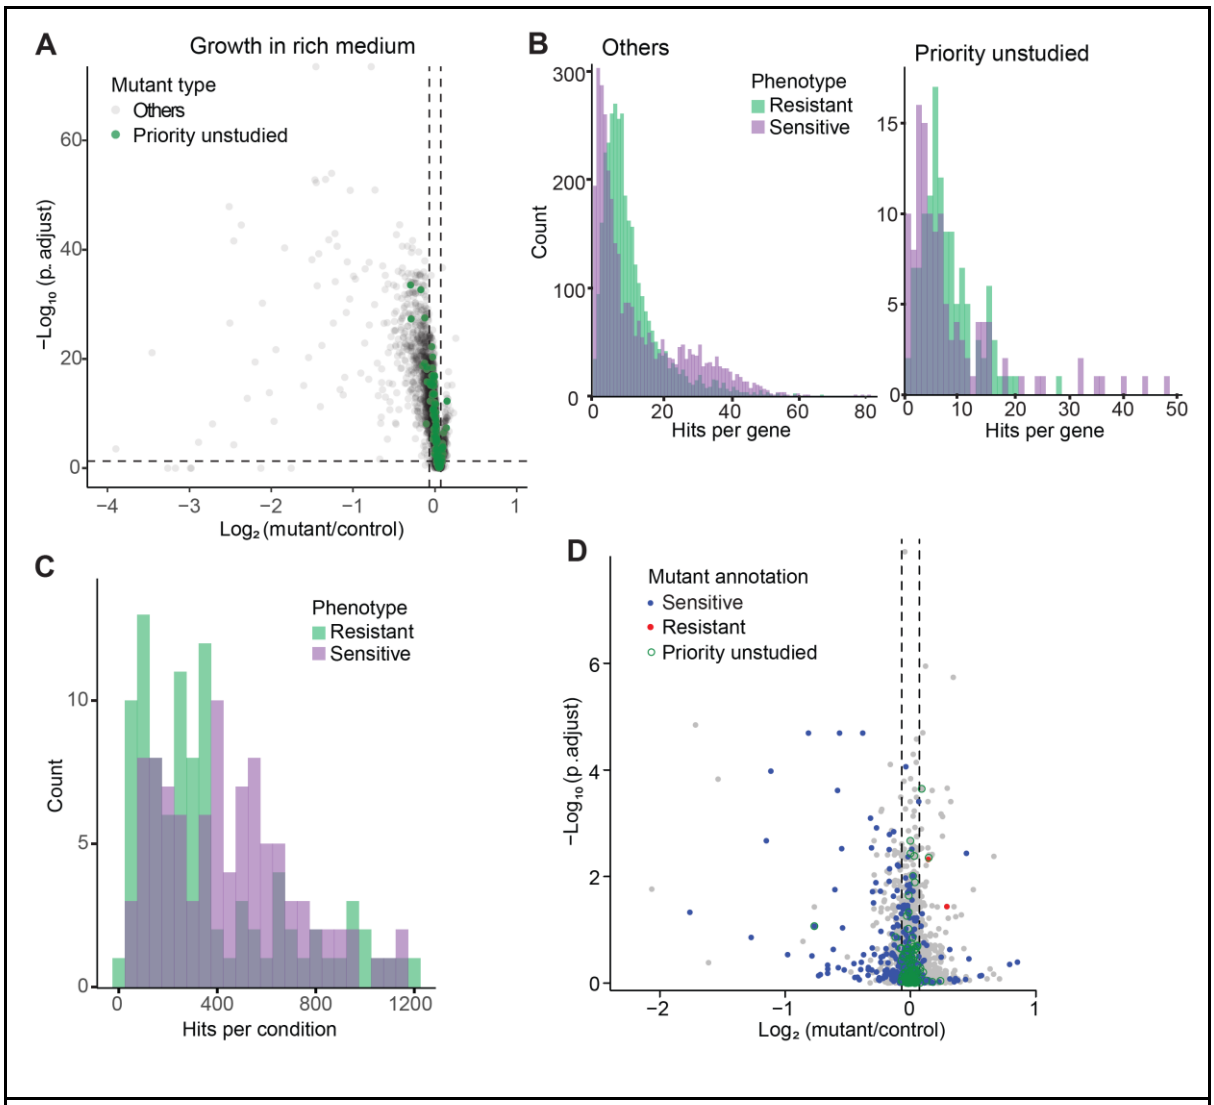

Figure 2. Colony-based phenotyping of deletion mutants.
(A) Volcano plot of mutant colony sizes for priority unstudied genes (green) and all other genes (grey) growing in rich medium. The dashed vertical and horizontal lines show the 5% difference in effect size and significance threshold, respectively. Strains with lower fitness (smaller colonies) are <0 on the x-axis, and those with higher fitness are >0. We applied a significance threshold of 0.05 after Benjamini–Hochberg correction for multiple testing and a difference in fitness of abs(log2(mutant/wild type))> log2(0.05) to call hits based on colony size (n = 60).
(B) Distributions of significant hits per gene for priority unstudied mutants (right) and all other mutants (left) showing faster (green) or slower (purple) colony growth compared to wild-type cells.
(C) Distributions of significant hits per condition for all mutants showing faster (green) or slower (purple) colony growth compared to wild-type cells.
(D) Fitness of mutant cells growing in 2% glycerol+0.01% glucose in this study, plotted as in (A). Blue: mutants identified as slow growing in the same condition in our previous screen (32); red: mutants annotated as 'increased cell population growth on glycerol carbon source' (FYPO:0004167) (62); green rings: mutants in priority unstudied genes. Seventy-nine mutants were slow growing in both screens ($p = 1x10^{-39}$; Fisher's exact test, FDR corrected for multiple testing).

To validate our approach, we compared one condition (growth in medium with glycerol) with previous results from the same condition (32). This analysis revealed a strong overall agreement

between the two datasets (Figure 2D). Differences may reflect the intrinsic variability of high-throughput screens and the more refined normalization procedure applied here (58,59), using wild-type cells as a common reference. We also compared correlations across phenotypes between deleted genes in our phenotype data with protein-protein interaction data in STRING (63). This analysis revealed 425 protein pairs that were supported by phenotype correlations, many more than expected by chance (Figure 2-figure supplement 2). Thus, phenotype correlations between pairs of knockout strains recapitulate known protein interactions assayed independently, demonstrating the validity of this approach.

Mutants showing altered resistance to oxidative stresses

These extensive phenotype data provide rich functional information regarding specific genes and processes that negatively or positively affect cellular fitness in a wide range of environmental or physiological contexts. As an example, we highlight oxidative stress which triggers a widely studied cellular response (64,65). Our data included mutant fitness under different doses or exposure times of three oxidants: hydrogen peroxide (H_2O_2), *t*-butylhydroperoxide (TBH), and diamide. Over 90% of the mutants were sensitive or resistant under at least one of these conditions (Figure 2-figure supplement 3). Sensitivity or resistance was often specific to the oxidant dose. For example, of the 495 and 220 sensitive mutants in the low or high dose of H_2O_2, respectively, only 96 mutants were sensitive in both doses. Even more pronounced, of the 304 and 95 resistant mutants in the low or high dose of H_2O_2, respectively, only 25 mutants were resistant in both doses. This finding is consistent with studies showing that different stress-response pathways and gene-expression programmes are launched in different doses of H_2O_2 (64,66).

We generated conservative lists of mutants that were sensitive or resistant in at least three of the six oxidative-stress conditions tested, resulting in 610 sensitive and 365 resistant mutants (Supplementary file 1). Using AnGeLi (67), we looked for functional enrichments among these lists with respect FYPO (62) and GO terms (40). The 610 genes leading to oxidative-stress sensitivity when deleted showed diverse enrichments, including 434 genes ($p = 8.2x10^{-7}$) involved in metabolic processes, particularly in the negative regulation of metabolism, 30 genes ($p = 5.0x10^{-7}$) involved in endosomal transport, 79 genes ($p = 1.1x10^{-11}$) that regulate transcription, 503 genes ($p = 8.4x10^{-50}$) required for normal growth, 104 genes ($p = 2.0x10^{-20}$) with cytoskeletal functions, and all six genes ($p = 8.9x10^{-5}$) of the adenylate cyclase-activating G-protein coupled receptor signalling pathway. Surprisingly, the 610 genes were not significantly enriched for genes induced during oxidative stress, but they included 42 genes ($p = 3.3x10^{-5}$) encoding ribosomal proteins that are down-regulated during oxidative stress (66). This finding is consistent with studies from budding yeast reporting poor correlation between the genes induced in response to stress and the genes required to survive that stress (68,69). It also raises the interesting possibility that certain ribosomal proteins have specialized roles in stress protection. The 610 sensitive mutants included

only 64 genes ($p = 1.3x10^{-16}$) already annotated with the FYPO term 'sensitive to hydrogen peroxide'. However, 445 of the 610 genes ($p = 3.1x10^{-57}$) are annotated as 'increased sensitivity to chemical', indicating that oxidant-sensitive mutants also tend to be sensitive to other compounds. In contrast to the sensitive mutants, the 265 resistant mutants showed no informative functional enrichments, and only 8 of these genes have been previously annotated to the FYPO term 'resistant to hydrogen peroxide'. Thus, the present study uncovered over 900 proteins not previously linked to oxidative stress in *S. pombe*, including 18 priority unstudied proteins.

Clustering of mutants with similar phenotype profiles

To explore functional signatures of the deletion mutants across all conditions, we applied *k*-medoids clustering of our phenotype data, revealing eight main clusters (Figure 3A). Analyses with Metascape (70) and AnGeLi (67) showed that each of these eight clusters feature distinct functional enrichments in different GO and KEGG pathway terms (Figure 3B; Supplementary file 1) and/or in genes from published expression or phenotyping studies. Examples of significant functional enrichments for these clusters are described below. Cluster 1 (293 genes, 12 priority unstudied) is enriched for genes showing high expression variability across different conditions (71) and for genes induced during meiotic differentiation (72) and in response to TORC1 inhibitors (29). Cluster 2 (570 genes, 20 priority unstudied) is enriched for phenotypes related to cell mating and sporulation, e.g. 'incomplete cell-wall disassembly at cell fusion site' or 'abnormal shmoo morphology' (73). Cluster 3 (806 genes, 29 priority unstudied) is enriched in endosome-vacuole, peroxisome and protein catabolism functions and in core environmental stress response genes (74). Cluster 4 (454 genes, 19 priority unstudied) is enriched in functions related to mitochondrion organisation and oxidoreductase activity as well as in genes affecting the chronological lifespan (36). Cluster 5 (633 genes, 24 priority unstudied) is enriched for citrate-cycle and sulphur-relay pathways as well as phenotypes related to abnormal cell growth and increased sensitivity to chemicals, based on FYPO terms from several studies (62). Cluster 6 (210 genes, 4 priority unstudied) is enriched in many different functions related to actin and microtubule cytoskeletons, cell morphogenesis, chromatin segregation and ribosomes as well as in ribosomal-protein gene clusters that are repressed in oxidative stresses (66). Consistent with these GO enrichments, Cluster 6 is also enriched in phenotypes related to abnormal cell shape and cytoskeleton organisation, based on FYPO terms from numerous studies (62). Cluster 7 (137 genes, 7 priority unstudied) is enriched in autophagy-related functions. Cluster 8 (346 genes, 5 priority unstudied) is enriched in metabolism-related functions like amino-acid biosynthesis and mitochondrion, in genes with slow-growth (35) and caffeine-sensitive (31) phenotypes as well as in genes showing high levels and stability of mRNAs (75,76). These enrichments allow inferring possible functions for the mutants in unknown genes that show similar phenotypic profiles to genes functioning in known processes through the principle of 'guilt by association'.

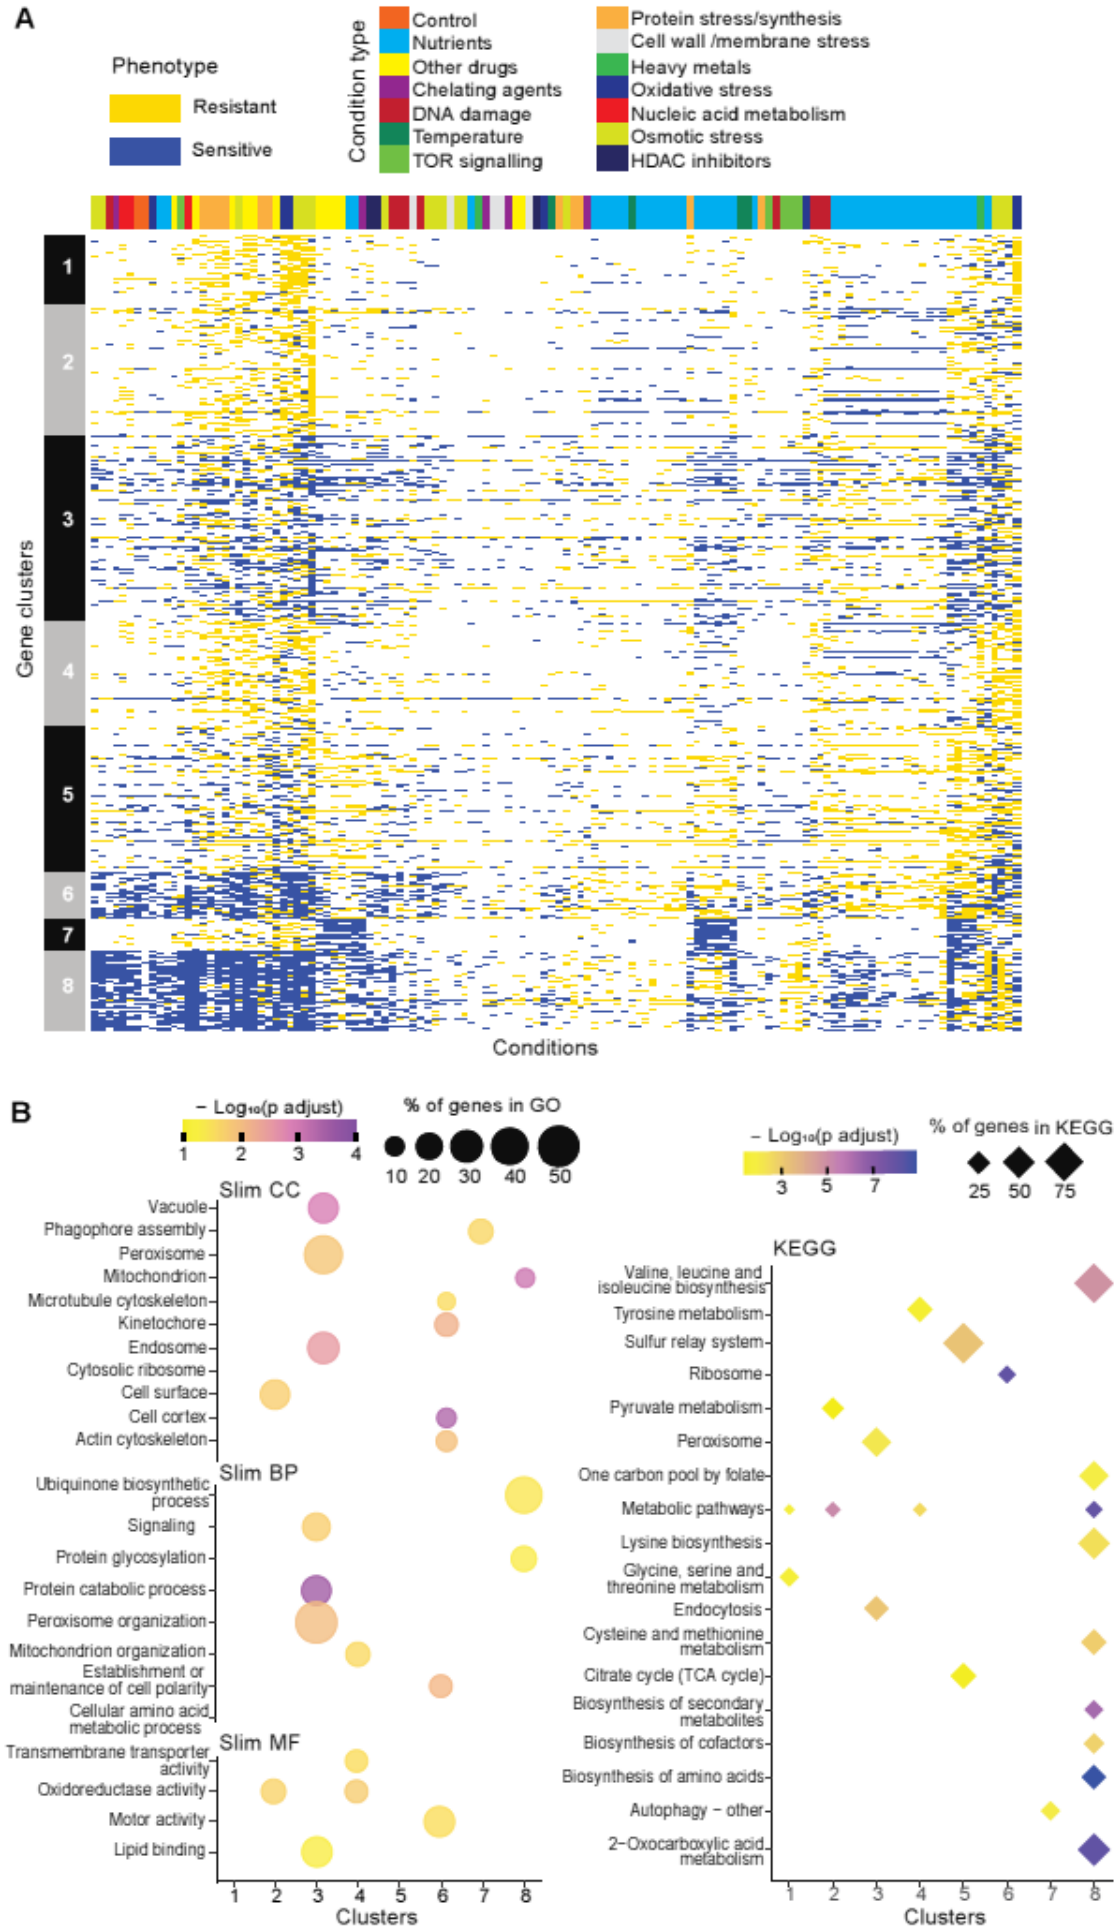
A
Phenotype
Resistant
Sensitive
Condition type
Control
Nutrients
Other drugs
Chelating agents
DNA damage
Temperature
TOR signalling
Protein stress/synthesis
Cell wall /membrane stress
Heavy metals
Oxidative stress
Nucleic acid metabolism
Osmotic stress
HDAC inhibitors
Gene clusters
1
2
3
4
5
6
7
8
Conditions
B
− Log10(p adjust)
1 2 3 4
% of genes in GO
10 20 30 40 50
Slim CC
Vacuole
Phagophore assembly
Peroxisome
Mitochondrion
Microtubule cytoskeleton
Kinetochore
Endosome
Cytosolic ribosome
Cell surface
Cell cortex
Actin cytoskeleton
Slim BP
Ubiquinone biosynthetic process
Signaling
Protein glycosylation
Protein catabolic process
Peroxisome organization
Mitochondrion organization
Establishment or maintenance of cell polarity
Cellular amino acid metabolic process
Slim MF
Transmembrane transporter activity
Oxidoreductase activity
Motor activity
Lipid binding
1 2 3 4 5 6 7 8
Clusters
− Log10(p adjust)
3 5 7
% of genes in KEGG
25 50 75
KEGG
Valine, leucine and isoleucine biosynthesis
Tyrosine metabolism
Sulfur relay system
Ribosome
Pyruvate metabolism
Peroxisome
One carbon pool by folate
Metabolic pathways
Lysine biosynthesis
Glycine, serine and threonine metabolism
Endocytosis
Cysteine and methionine metabolism
Citrate cycle (TCA cycle)
Biosynthesis of secondary metabolites
Biosynthesis of cofactors
Biosynthesis of amino acids
Autophagy − other
2−Oxocarboxylic acid metabolism
1 2 3 4 5 6 7 8
Clusters

Figure 3. Mutant phenotype signatures.
(A) Clustering of discretized data for 3,449 deletion mutants (rows) in 98 conditions (columns). For discretizing the phenotype data, we classed mutants as either resistant (yellow), sensitive (blue), or similar (white) relative to their fitness in the corresponding control condition, applying an effect-size threshold of ≥10% and, for the two benign conditions, also a significance threshold of $p < 0.05$ (Supplementary file 1). To improve the clustering, we removed conditions and mutants that produced <4 phenotypes, the 3 conditions that produced >2400 phenotypes, and genes with >50% of missing values. Different types of conditions are colour-coded on top. Indicated at left are the 8 main clusters, calculated using *k*-medoids in R (pam function) and a distance of 1-Pearson correlation. Using the function fviz_nbclust (R package factoextra), we determined that 8 is an optimum number of clusters.
(B) Functional enrichments of the 8 gene clusters in (A) using GO slim categories (51) for Cellular Component (CC), Biological Process (BP) and Molecular Function (MF) (left, using SystempipeR) as well as KEGG pathways (right, using g:Profiler) (77).

The NET-FF meta-predictor combines protein-network and -family data to associate GO terms with *S. pombe* proteins

We developed an integrated pipeline (NET-FF) to predict GO terms for *S. pombe* proteins by combining two independent established approaches, NetHom and CATHPredictGO. NetHom uses a machine-learning method to combine network features derived using DeepNF (55) with homology features from CATH protein functional families (FunFams) (78). CATHPredictGO also uses homology features but exploits this data with an independent, complementary approach. To obtain our final set of predictions, NET-FF combines the predicted protein functions from the NetHom and CATHPredictGO predictors to increase confidence and coverage of the predictions (see Methods for details of NetHom, CATHPredictGO and their integration in NET-FF).

Using NET-FF, we predicted a total of 2,390,915 GO terms for *S. pombe* proteins. Filtering for high scoring predictions (>0.7 from maximum of 1.0) and applying a taxon constraint resulted in 56,594 GO terms (Figure 4-figure supplement 1). These terms comprised 42,808 Biological Process (BP) and 13,786 Molecular Function (MF) terms, across 2,852 *S. pombe* proteins, including 53 priority unstudied proteins Supplementary File 2). Below, we discuss the predictions from NET-FF and their validation. Subsequently, we combine our NET-FF predictions and phenotype-correlation data with curated annotations in PomBase to derive novel functional associations of *S. pombe* proteins.

Systematic comparison of NET-FF predictions and PomBase annotations

To assess the validity of our 56,594 high scoring NET-FF predictions, we determined the similarity between the highly curated GO terms annotated by PomBase and those predicted by NET-FF. Although only 60% of the PomBase annotations are supported by direct experimental evidence, a further 15,438 (34.5%) are curated from inferred annotations, providing a rich knowledgebase for the *S. pombe* community. To compare the PomBase annotations with NET-FF predictions, we modified a standard approach for comparing two GO terms (GOGO method; 79) to perform a

pairwise comparison using the same gene with different GO terms sets, e.g. GOGO(gene $A_{PomBase}$ - gene $A_{NET\text{-}FF}$). This approach provided a GO semantic similarity between the two datasets for a given gene, with 0 having completely different GO terms and 1 having identical GO terms. Figure 4A shows the distribution of GOGO scores obtained by comparing the two datasets and a randomized control obtained by shuffling GO terms between genes. Although the datasets differ (median difference GOGO-BP = 0.52, median difference GOGO-MF = 0.64), the real data comparison gives a higher GOGO semantic similarity when compared to random. Supplementary file 2 provides the list of predictions from NET-FF above our cut-off, including those for which our predictions agree or are very similar to those annotated independently in PomBase. We conclude that NET-FF predictions are more similar to the available annotations in PomBase than random.

Aside from false predictions in NET-FF, the difference in GO terms predicted by the GOGO analyses may suggest the presence of more specific GO terms in either dataset, i.e. more informative annotations that occur deeper in the GO hierarchy. We calculated the Information Content (IC) for all GO terms in PomBase (Methods) and, for a given gene, extracted the GO with the highest IC. We then calculated for each gene the ΔIC as ($IC_{NET\text{-}FF}$ - $IC_{Pombase}$), with positive or negative results indicating that NET-FF or PomBase, respectively, have the more informative GO for this gene. Figure 4B shows the distribution of ΔIC for the BP and MF gene datasets that are in common between the NET-FF predictions and PomBase annotations. For MF, 30% of the PomBase annotations were more informative than NET-FF predictions, while 28% of the NET-FF predictions were more informative than PomBase annotations (Figure 4B). For BP, the proportion of more informative PomBase annotations (51%) was higher than for more informative NET-FF predictions (33%). Notably, the situation was different for the priority unstudied genes, where PomBase was more informative with respect to MF (42.5% vs 13.5%), while NET-FF was more informative with respect to BP (14.9% vs 8.5%) (Figure 4B). Given that BP terms are key to understand the cellular roles of proteins, this result indicates that the NET-FF predictions provide valuable new clues for the function of priority unstudied proteins.

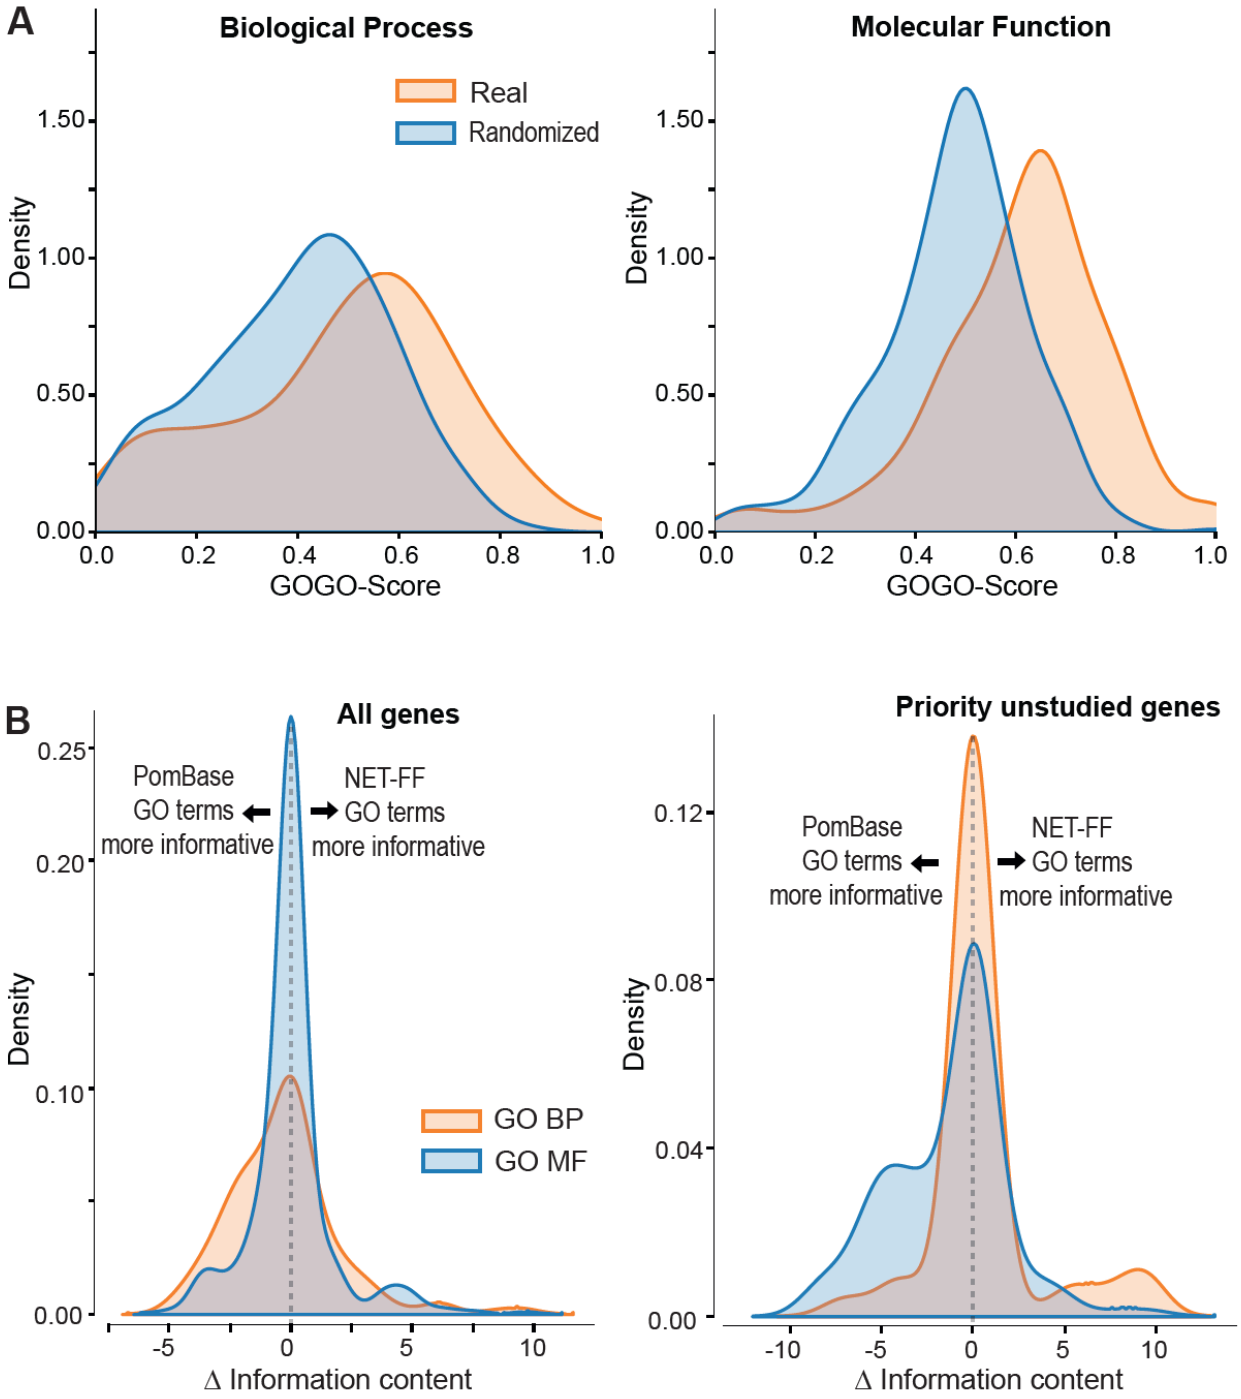

Figure 4. NET-FF predictions of GO terms for *S. pombe* proteins.
(A) GOGO semantic similarity scores for GO terms predicted by NET-FF and GO terms annotated by PomBase, both for real (orange) and randomized (blue) datasets from NET-FF. Left: Biological Process data, right: Molecular function data.
(B) Information content distribution for GO terms predicted by NET-FF and GO terms annotated by PomBase, both for GO Biological Process (orange) and GO Molecular Function (blue) terms. Left: dataset for all genes, right: subset of priority unstudied genes.

PomBase annotations are subject to manual curation including extensive literature searches for available gene information (18,19). Therefore, we naturally expected the proportion of informative PomBase annotations to be higher than for NET-FF predictions. Of the 56,594 NET-FF predictions, 22,060 had a high information content (IC >5), reflecting more specific GO terms (Figure 4-figure supplement 1; Supplemental file 2). Of these, 9,374 terms were identical to PomBase GO terms with experimental annotations for the same genes. This overlap in predictions is not that surprising, because NET-FF was trained on GO terms with experimental annotations but not on terms inferred from electronic annotations. Notably, of the remaining 12,686 NET-FF predictions, 6,052 (48%) were identical to curated or electronically inferred PomBase annotations for the same genes (PomBase GO version 13/01/23). This substantial overlap provides compelling validation of NET-FF.

We further filtered the 6,634 NET-FF predictions that were not identical to PomBase annotations to remove 276 GO terms with more informative (deeper) annotations in PomBase, 498 GO terms that became obsolete, and 4,185 GO terms that were up-propagated and thus redundant with more informative terms. This resulted in 1,675 novel NET-FF GO predictions for 783 *S. pombe* genes, including 47 predictions for 23 priority unstudied genes (Supplementary file 2). Of these novel predictions, 1,481 terms are in parts of the GO tree not previously assigned to these genes and 194 terms provide more informative annotations than PomBase. Some of the novel predictions contain ‘regulation of’, suggesting a role in a regulatory process that could be direct or indirect. The substantial overlap between the NET-FF predictions and the high-confidence annotations available in PomBase support the notion that the 1,675 novel predictions can provide unique functional clues for the associated genes.

Manual comparison of selected NET-FF predictions with GO terms and phenotypes annotated in PomBase

We also manually validated a selection of the NET-FF GO terms using the rich extant information in PomBase, including both GO and FYPO terms (18,19,62). We selected 312 genes with predicted GO terms that relate to broad cellular processes: ageing/cell death, growth/response to nutrients, reproduction, recombination, cold/heat response, histone acetylation/deacetylation, DNA damage/repair, autophagy, and actin (Supplementary file 2). We checked to what extent aspects of our NET-FF predictions for these genes had already been annotated by searching PomBase for GO and/or FYPO terms related to the predicted GO terms. For example, ‘decreased cell population growth at low temperature’ (FYPO:0000080) relates to ‘cold acclimation’ (GO:0009631) and ‘response to cold’ (GO:0009409). This analysis identified 165 of the 312 genes (52.9%) with NET-FF GO terms similar to the PomBase FYPO and/or GO annotations (Supplementary file 2). Moreover, several NET-FF GO terms were supported by our colony-based phenotype data. For

example, mutants in 4 of 13 genes with predictions related to cold/heat response showed phenotypes at low and/or high growth temperatures, mutants in 3 of 6 genes with predictions related to histone acetylation/deacetylation showed phenotypes in conditions affecting those processes, mutants in 31 of 34 genes with predictions related to growth/nutrient response showed altered growth in different nutrients, and mutants in 30 of 31 genes with predictions related to DNA damage/repair showed phenotypes in conditions causing different types of DNA damage. Thus, this independent experimental evidence further supports our NET-FF predictions.

For the remaining 147 of the 312 selected genes (47.1%), the NET-FF GO terms represent novel predictions, revealing putative functions that were not previously known (Supplementary file 2). The proportions of novel predictions varied greatly for different process-related categories. As expected, processes that are well-studied by *S. pombe* researchers were already well-captured in PomBase, with relatively few novel predictions. For example, only 2 of 13 genes (15.4%) in the category 'Histone acetylation/deacetylation' and 7 of 42 genes (16.7%) in the category 'DNA damage/repair' were not annotated as such in PomBase. In contrast, 11 of 12 genes (91.7%) in the category 'ageing/cell death' and 26 of 44 genes (59.1%) in the category 'growth/response to nutrients' were not annotated as such in PomBase. Five of the genes associated with novel GO-term predictions were priority unstudied. The new predictions allow us to make specific functional inferences about these unknown proteins (Supplementary file 2). For example, *SPAC15A10.10* was predicted to function in the regulation of nuclear division and actin-related processes, and *SPAC25B8.08* in actin-dependent ATPase activity and organelle localisation. Notably, deletion of *SPBC2D10.03c,* predicted to function in the response to heat, leads to decreased growth at high temperature (80), thus validating the prediction. Overall, the NET-FF predictions recapitulate GO and/or phenotype predictions in PomBase for half of the genes in the selected broad categories, thus adding confidence to our method.

Experimental validation of genes predicted to function in cellular ageing

The general category 'ageing/cell death' included the largest proportion of genes linked to NET-FF predictions not available in PomBase (11 of 12 genes; 91.7%). This result raises the possibility that these are novel genes affecting the chronological lifespan (CLS) of stationary phase *S. pombe* cells, a model for cellular ageing. We determined the maximal CLS for eight deletion mutants of these genes and control strains (Figure 5A). Seven of these mutants showed subtle to strong effects on the CLS, including one mutant in a priority unstudied gene (*SPAC23C4.09c*) featuring a strong longevity phenotype. In contrast to this result obtained using stationary-phase cells, the *SPAC23C4.09c* mutant has been reported to reduce the CLS in quiescent cells (81). We have shown before that some genes exert opposite effects on the CLS in stationary-phase cells, limited by glucose, and in quiescent cells, limited by nitrogen (35). Besides cell death, *SPAC23C4.09c* was predicted to function in DNA-damage repair (Supplementary file 2), a process important for

ageing (82). Notably, the *SPAC23C4.09c* gene formed a tight phenotype-correlation network with ten other genes, including four with known CLS phenotypes (Figure 5B). Several of these genes are involved in respiration and carbon metabolism, including *otg1* (galactosyltransferase), *cox20* (cytochrome c oxidase assembly protein), *sdh8* (mitochondrial respiratory chain complex II assembly factor), *sdh3* (succinate dehydrogenase cytochrome b subunit), and *emi5* (succinate dehydrogenase complex assembly). Consistent with these interactions, the *SPAC23C4.09c* mutant grows slowly on non-fermentable media (33). Together, these experimental results validate the computational predictions by confirming new genes with roles in cellular ageing and illustrate how our predictions and the phenomics data complement each other to reveal new gene functions.

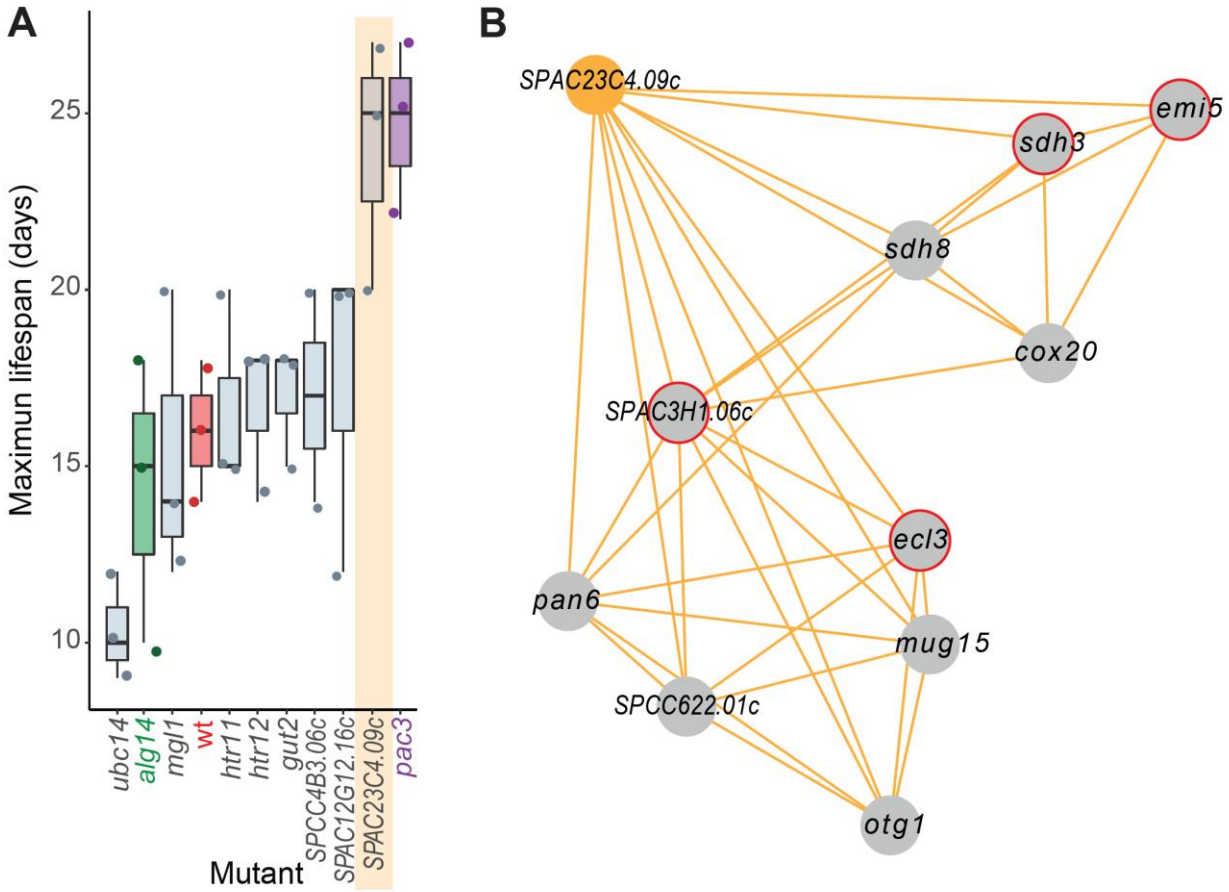

Figure 5. Experimental validation of new ageing-related gene reflecting its NET-FF GO predictions and phenotype-correlation network.
(A) Maximal CLS for eight selected deletion mutants (grey) of genes with NET-FF GO terms related to ageing/cell death and DNA-damage repair, measured using a high-throughput CLS assay (36). Of the eight mutants, we have recently shown that *ubc14* and *htr12* are short- and long-lived, respectively (36). As controls, the maximal CLS of known short-lived (*alg14,* green) and long-lived (*pac3,* purple) mutants (36) and wild-type cells (wt, red) are included. Three independent biological repeats were performed with actual datapoints shown as dots. Orange band: mutant in priority unstudied gene (*SPAC23C4.09c*) showing strong longevity phenotype.
(B) Cluster from Cytoscape network (see Figure 6) representing phenotype correlations between the *SPAC23C4.09c* gene-deletion mutant and ten other mutants. Orange edges show positive phenotype correlations, and red circles indicate genes with previously known CLS phenotypes (36,83). Details discussed in main text.

Combining NET-FF predictions with phenotype correlation data to validate functional associations and derive higher-confidence predictions

To further filter our set of 22,060 NET-FF predictions with high scores and high information content, we identified pairs of genes showing high similarity both in their GO-term predictions (GOGO scores) and in their deletion-mutant phenotypes (PHEPHE scores). This approach should provide a higher confidence subset of gene pairs for which the two independent, orthogonal methods agree. For each gene pair with a predicted GOGO BP score >0.5, the GOGO score was multiplied by the PHEPHE score of the same gene pair (obtained from Pearson correlation of phenotype profiles across conditions). If the combined product of GOGO and PHEPHE scores was >0.5, we added the gene pair link to a network. This analysis resulted in a much larger network for the gene pairs with high GOGO and PHEPHE scores than for random gene pairs (Figure 6-figure supplement 1A). Supplementary file 2 provides a list of the gene pairs with high GOGO and high PHEPHE scores, both for BP terms (553 gene pairs) and MF terms (606 gene pairs). The agreement between these two independent approaches increases the confidence in both our NET-FF predictions and phenotype data.

For comparison, we used an analogous approach to identify sets of gene pairs that showed high GO similarity in their PomBase annotations. Supplementary file 2 provides a list of these gene pairs with high GOGO and high PHEPHE scores, both for BP terms (12,587 gene pairs) and MF terms (20,433 gene pairs). As expected, these are larger gene sets than for the NET-FF predictions, generating a much larger network for these gene pairs than random pairs (Figure 6-figure supplement 1B). This analysis adds further confidence in our phenotype data.

Integrated analysis of phenotype-correlation networks and GOGO similarities

We used our phenotype-correlation data to construct a Cytoscape network consisting of 138 clusters (Supplementary file 2). Given that the phenotype-correlation clusters are enriched in genes with similar functions, we looked for gene pairs with similar GO terms (GOGO score >0.7) based on NET-FF predictions and/or PomBase annotations (Figure 6A-C; hatched blue and purple edges, respectively). As expected, the clusters contained many more high GOGO gene pairs based on PomBase annotations, but the NET-FF predictions added some new information (Supplementary file 2). An example is the unique NET-FF-based link between *rpl1001* and *hse1* in Cluster 31 (Figure 6B). In many cases, high GO similarities linked gene pairs that were not directly linked via high phenotype correlations although they were part of the same cluster, thus independently supporting their cluster membership. Of the nine clusters containing at least 15 genes, six showed more high GOGO gene pairs than expected by chance based on NET-FF predictions (Figure 6-figure supplement 2). Supplementary file 2 contains information on the

clusters whose high GOGO gene pairs are enriched for specific GO terms based on NET-FF, PomBase, and/or the combination of NET-FF and PomBase.

We then examined whether inclusion of the NET-FF predictions led to additional GO term enrichments of the clustered data over existing PomBase annotations. Based on PomBase, 308 GO terms were significantly enriched, while the NET-FF predictions added a further 63 GO-term enrichments, including 32 GO terms directly based on NET-FF and 31 GO terms becoming only enriched when combining PomBase and NET-FF terms (Supplementary file 2). Together, these analyses show that the NET-FF predictions can add new information to the phenotype clusters, and the concordance between high GOGO gene pairs and phenotype correlations increase confidence in both the NET-FF predictions and phenotype-correlation data.

The phenotype-correlation network allows for a fine-grained, nuanced analysis of functional relationships (59). Based on the genes with known roles, the clusters were functionally coherent based on gene enrichment analyses and manual inspection, reflecting protein complexes and/or biological pathways. For example, Cluster 13 consisted of a branch with all members of the PAN complex, a polyA-specific ribonuclease, and another branch containing six genes involved in cAMP signalling, along with two genes that function in RNA interference (Figure 6A). Consistent with the latter branch, published reports suggest a link between cAMP signalling and RNA interference (84–86). Other clusters contained mostly genes involved in vacuolar/endosomal transport and peroxisome function, along with poorly characterized genes (Figure 6B). Some clusters linked members of protein complexes and even reflected different sub-complexes (Figure 6C). Of the 124 priority unstudied genes, 70 showed phenotype correlations with at least one characterized gene (Supplementary file 2). For example, five priority unstudied genes implicated in mitochondrial functions correlated with genes that play known roles in mitochondria. Figure 6D shows two such networks including priority unstudied genes along with their nearest neighbours, placing the unstudied genes in a functional context based on phenotype correlations. Figure 6-figure supplement 3 provides the entire Cytoscape network which can be interactively explored as specified in the legend. Thus, these phenotype-correlation clusters can reveal unexpected relationships between different cellular processes and provide functional clues for poorly characterized or unstudied proteins.

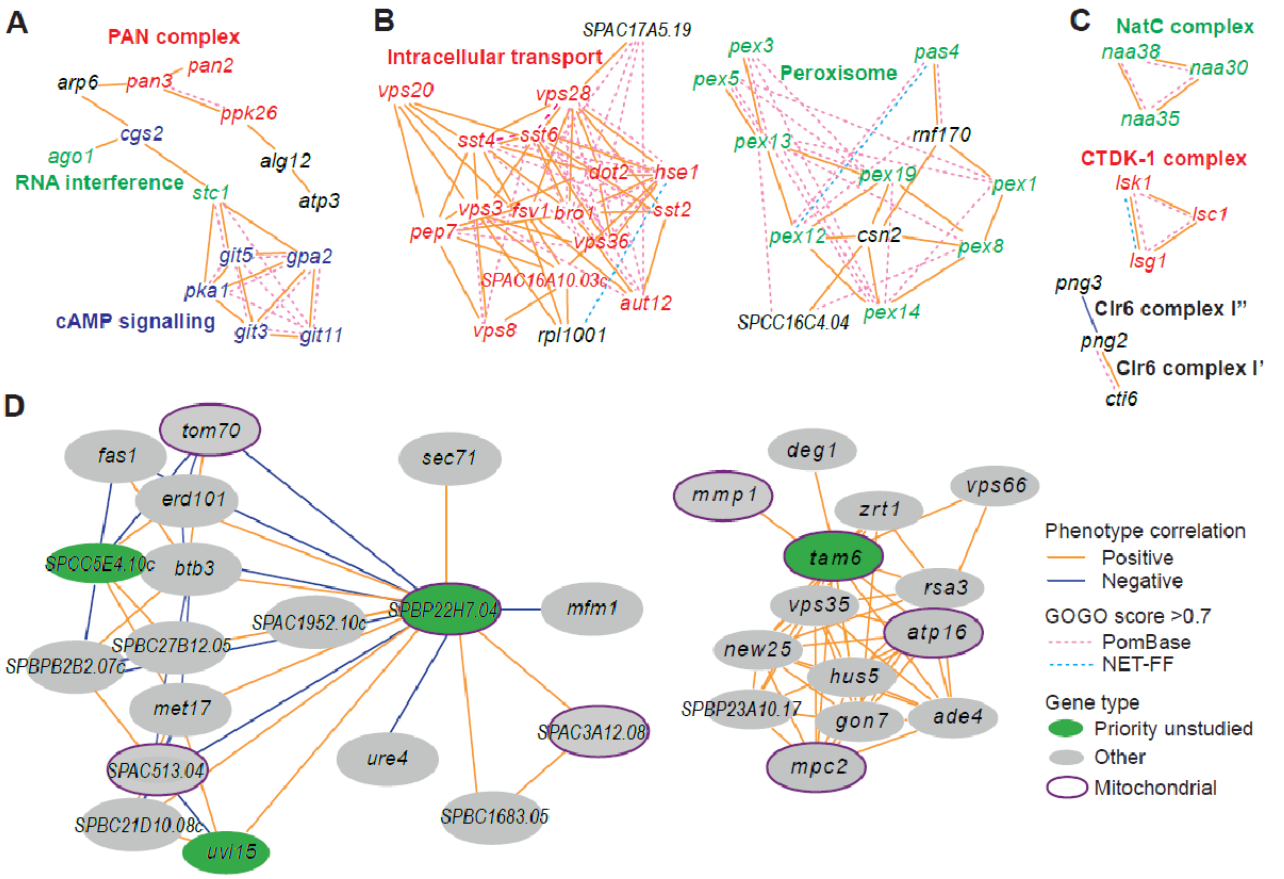

Figure 6. Example of phenotype-correlation clusters from Cytoscape network. Positive and negative phenotype correlations across the 131 conditions are indicated as orange and blue edges, respectively. Gene pairs with high GO similarities based on PomBase annotations or NET-FF predictions are indicated as hatched pink and bright blue edges, respectively.
(A) Cluster 13 consisting of a branch with all members of the PAN complex (red) and a branch containing six genes involved in cAMP signalling (blue), along with two genes involved in RNA interference (green).
(B) Cluster 31 enriched for genes involved in vacuolar and endosomal transport (red) and Cluster 22 enriched for genes with peroxisome functions (green).
(C) Three clusters linking members of the same protein complexes, involved in protein maturation (NatC, green), regulation of transcription elongation (CTDK-1, red), and histone deacetylation (Clr6, black). The data reflect different Clr6 sub-complexes, with *png2* and *cti6* being positively correlated (Clr6 complex I') but *png2* and *png3* being negatively correlated (Clr6 complex I") (87).
(D) Two networks containing priority unstudied genes (green), including those with implicated mitochondrial functions (green with purple ellipses), together with their nearest neighbours based on phenotype correlation, which include genes with known mitochondrial functions (grey with purple ellipses). For simplicity, high GOGO pairs are not indicated in these networks.

Conclusions

We applied a phenomics approach using colony-based assays for 3509 deletion mutants in all non-essential genes of fission yeast under many different conditions, including different nutrients and drugs as well as oxidative, osmotic, heavy-metal, protein-homeostasis and DNA-metabolism stresses. All 131 conditions tested led to growth phenotypes in over 100 mutants. We detected phenotypes in at least one condition for 99.6% of the mutants, including all 124 mutants of priority unstudied genes for which we obtained data. Together, these rich data exceed all the phenotypes currently annotated in PomBase. By combining two orthologous gene-function prediction methods (NET-FF), we predicted 22,060 GO terms with high scores and high information content across 2167 genes. Compared to existing PomBase annotations, these GO predictions are more similar than random and of comparable information content. Moreover, our phenotype-correlation data and NET-FF predictions showed a strong overall concordance with the PomBase GO annotations and provide new functional clues for many genes. The agreement between the independent phenomics, curation, and computational approaches increases the confidence in both our NET-FF predictions and phenotype data. Notably, 15,426 NET-FF predictions are identical to PomBase GO annotations for the same genes, while 1,675 predictions provide novel GO associations for 783 genes, including 47 predictions for 23 priority unstudied proteins. Experimental validation based on selected predictions revealed new proteins involved in cellular ageing and showed that combining the NET-FF predictions and phenomics data can reveal new protein functions. Integrated analyses revealed good agreement between phenotype-correlations and GO similarities for gene pairs, with the NET-FF predictions adding some unique information to the PomBase annotations. These wet and dry approaches provide a rich framework to better understand the functional relationships between proteins and to mechanistically dissect the roles of proteins in physiologically relevant conditions.

Methods

Yeast strains

Yeast strains used in this study are described in Supplementary file 1; most of the strains are derived from a prototroph version of the Bioneer V.5 deletion collection (56) generated as described (32), and the rest of the strains were kindly provided by Kathy Gould (57). Prior to phenotyping, the strains were arranged into a 1536 format with grid of a control wild-type strain (*h-972*) as described before (58). Cells were grown on standard Edinburgh minimal medium (EMM) or rich yeast extract medium with supplements (YES) as specified.

Colony-based phenotyping of deletion-mutant library

The conditions used for phenotyping are provided in Supplementary file 1. For phenotypic screens, plates containing 1536 colonies were pinned onto YES plates using a RoToR arraying robot (Singer Instruments) and incubated for 24 hours at 32°C; these plates served as template. Strains from the template plates were pinned onto the phenotyping plates using 1536-pin pads and applying 4% pressure to minimize the biomass inoculated onto each plate. Each template plate was used to inoculate 4-6 phenotyping plates. After pinning plates were wrapped in plastic film to prevent excessive drying and incubated at 32°C (except for plates used in temperature assay, which were incubated at the appropriate temperature) for 20 hours prior to scanning. For conditions where the growth was highly impaired plates were incubated for 40 hours prior to scanning.

Plates were scanned and processed using *pyphe* (58,88). After processing with grid correction, median fitness and corrected (Benjamini-Hochberg) p-value thresholds (5% median fitness difference with respect to wt and adjusted p-value <0.05) were used to determine phenotypes in benign conditions (YES and EMM at 32°C). Every mutant was measured under these conditions in over 10 repeats. To determine phenotypes in the other conditions, measured in only 4 repeats, we applied a higher threshold (10% median fitness difference with respect to wt, and growth in control media) without using p-values.

Clustering of phenotyping data

To cluster the phenotypic data, we simplified the dataset by converting the data into ternary encoding using a median phenotypic value threshold: phenotypes showing a reduction of ≥10% on the phenotypic score were coded as -1, those showing an increase of ≥10% were coded as +1, and the weaker phenotypes in between were coded as 0. We removed the strains that contained >50% missing values and the conditions with >90% of the scores being 0.

Clusters for Figure 3 were determined using k medoids in R (pam function from package cluster) where the distance between the genes was calculated as 1-Pearson correlation, using 8 as the optimal number of cluster determined by the R package factoextra (fviz_nbclust function).
For the networks in Figure 6, we used Pearson correlations and filtered on absolute r values >0.7 and adjusted p <0.01 (59). The network was visualised using Cytoscape and clustering was done using community clustering (GLay) from the clustermaker extension (89).

NET-FF integrated method for prediction of GO Molecular Function and Biological Process terms

We developed a new computational pipeline (NET-FF) combining two independent approaches (NetHom and CATHPredictGO), for predicting GO terms associated with *S. pombe* proteins.

NetHom is a machine learning method that combines protein family data from CATH functional families with protein network features generated using the DeepNF method (55). CATHPredictGO is a simple decision tree approach that only uses CATH protein family data and has been endorsed for precision in multiple rounds of CAFA independent assessment (49,52,53). It is included to increase confidence in predicted GO terms.

NetHom – a machine learning predictor based on protein network and homology based features: NetHom combines network features derived using the DeepNF approach pioneered by the Bonneau group with homology based features from CATH-FunFams (78). Each type of feature is described below.

Network embeddings: Unsupervised feature learning was used to extract the context *S. pombe* genes within networks from the STRING database (v11.0) (90). We applied a modified version of deepNF (55), that uses a multimodal deep autoencoder, a type of neural network, to embed genes in a low-dimensional space, according to their multi-network context. Individual networks were constructed for the 5,100 *S. pombe* proteins contained in STRING and each of six interaction types: `neighbourhood', `fusion', `cooccurence', `coexpression', `experimental' and `database'. The autoencoder compresses 30,600 dimensions of network data into 256 dimensions that can be used as features to train off-the-shelf machine learning models. A multimodal deep autoencoder was used with hidden layers for each of the six networks, followed by a single, shared embedding layer, followed by hidden layers for each of the six networks, with all layers containing 256 neurons. Sigmoid activations were used on the embedding layer, so that embedding values ranged between -1 and +1, whilst ReLU activations were used on the hidden layers. Models were trained using the Adam optimiser for 500 epochs with batch sizes of 128 examples using data from 90% of proteins. The remaining 10% of proteins were used as a validation set to monitor training using the binary cross-entropy loss function. After training, weights from the epoch with the lowest validation loss were used to generate embeddings for all proteins. Models were implemented using Keras (v2.1.5) and TensorFlow (v1.8.0).

Homology assignments from functional families (FunFams): To predict GO terms using homology data, we scanned the *S. pombe* proteins against Hidden Markov Models for CATH-FunFams (78). FunFams are subsets of protein domains within a CATH Superfamily predicted to share the same function (91,92). FunFams are generated using a purely sequence based approach that segregates sets of homologous domain sequences within a CATH superfamily according to differentially conserved residue positions likely to be associated with functional properties. Our protocol uses agglomerative clustering and exploits HHsuite (93) to iteratively compare HMMs derived for clusters of relatives within a CATH superfamily. The starting clusters are sets of relatives sharing at least 90% sequence identity. At each step the two most similar clusters are

merged to ultimately give a tree of relationships across the superfamily. Subsequently the FunFamer algorithm (94) is used to determine whether nodes within the tree should be merged into the same functional family (FunFam) based on sharing highly conserved residues, differentially conserved in other FunFams. Each FunFam contains at least one experimentally characterised relative with an assigned GO MF or GO BP term. All 5,396 protein sequences (downloaded from PomBase on Feb 21, 2019) were searched against the FunFam hidden Markov model (HMM) library from CATH (v4.2) using hmmsearch from HMMER3 (94,95) and an E-value threshold of E $<10^{-3}$. In total, 3,319 sequences had 149,099 hits to 23,900 FunFams. For machine learning, E-values were $-\log_{10}$-transformed to convert to a linear scale and transformed to a [n proteins x m FunFams] matrix, with 99.8% sparsity. For each gene and GO term pair, the corresponding position in the target matrix was set to 1 if the gene was annotated with the term. Because 23,901 FunFams were hit, the FunFam feature matrix was very wide, from a machine learning perspective. So, when predicting some GO term *g*, models were trained on FunFams that contain at least one protein annotated with *g*. GO annotations (91,96) for all UniProt (92) accessions contained in FunFams were downloaded in February 2020. All annotations were included, except those with NAS, ND, TAS or IEA evidence codes (except UniProtKB-kw IEA curated terms, which were included). GO terms were associated with FunFams by identifying all terms annotated to proteins in each FunFam. Ancestor terms that have `is_a', `has_part', `part_of' and `regulates' relationships were also included. Release ‘2018-11-12' of the GO was used.

Integrating the multiple data types (protein network, protein family data) using machine learning
NetHom: Supervised machine learning was used to predict GO annotations for *S. pombe* proteins using random forests. We used Julia (v1.5) and the DecisionTree.jl package, which implements random forests of classification and regression trees. The cost function employed the following criteria: no maximum depth, so trees could grow arbitrarily deep; a minimum of two samples is needed to split a node, resulting in terminal nodes with single samples in each; and no minimum purity increase, here defined according to minimising entropy. Forests of 500 trees were grown, where the trees were not pruned after they were grown. Overall model performance was estimated using five-fold stratified cross-validation. The data was shuffled before each cross-validation. Five independent trials of cross-validation were performed to estimate the model performance under different train-test splits. Terms were predicted using the one-vs-rest multiclass strategy. Two hyperparameters, the number of features {10, 25, 50, √n} and the partial sampling of examples {0.50, 0.75, 1.00}, were optimised using an exhaustive grid search. Each combination of parameters was assessed using a nested five-fold stratified cross-validation, evaluated using the area under the precision-recall curve (AUPR). Combinations of features—network embeddings, FunFams — were benchmarked for their ability to predict *S. pombe* protein function. We also assessed the benefit of including features on growth phenotypes. Functions were predicted for *S. pombe* proteins for which growth phenotypes were collected. The 53 *S. pombe* GO Slim terms

were used as targets (accessed from PomBase on Nov 14, 2018). GO Slim terms were chosen as targets, as these are sufficiently informative terms in the ontology. Following this, GO term annotations were predicted for *S. pombe* proteins using the combination of features that produced the highest AUPR. We used a five-fold stratified cross-validation strategy for prediction, whereby models predict labels for samples that they have never seen before, which forces the model to not overfit on the training data. Because we did not use growth phenotype features to make final predictions, we did not exclude any genes from our dataset. GO terms were included in the target set if they were annotated to between 50 and 1,000 proteins. GO annotations were propagated to their less specific parent terms in the ontology.

Performance of NetHom predictor: The protein network embeddings and/or CATH homology data were used as input to train a Random Forest method using one-vs-rest classification with 53 GO Slim BP terms for *S. pombe* as classes for the model (18). Different combinations of protein network and homology features were used to predict GO slim terms and their performance was evaluated using five-fold classification. Alone, network embeddings were the best set of features followed by FunFam homology data, with areas under the precision-recall curve (AUPR) of 0.58 and 0.42, respectively (Figure 4-figure supplement 2). We tested three combinations of features. The combination of network embeddings and FunFam data showed a performance that was 3% higher than the network embeddings alone, with an AUPR of 0.60 (Figure 4-figure supplement 2). We also tested for any benefit of using the colony-growth phenotype data described above, by including these data as a simple vector of phenotype presence or absence. Surprisingly, adding the phenotype data to the network data produced a slight drop in performance. Although the phenotype data has predictive value (Figure 2-figure supplement 2), these results suggest that it does not add to the comprehensive information derived from the STRING protein-interaction network. For this reason, we did not use the phenotype data for the prediction model but used it as an orthogonal cross-validation dataset.

CATHPredictGO - protein function prediction by protein family data alone: We also implemented a protein function prediction method based on membership of protein families. The predictor was constructed as a fast and frugal tree, which is a type of decision tree that asks a total of *n* questions and has *n* + 1 exits. First, a protein is scanned against the FunFams from CATH (v4.3) (78), using the FunFam inclusion thresholds. The set of non-overlapping domain hits with the highest bit scores were resolved using CATH-Resolve-Hits (97). If the protein had a significant match to one or more FunFams, the FunFam match with the highest bit score above the inclusion threshold was used. All GO terms annotated to existing members of that FunFam were predicted for the protein, with probability set to the proportion of members annotated with each GO term. Low probability predictions with $p < 0.1$ were removed. If the protein did not have a significant match to at least one FunFam, then the protein was scanned against Pfam-FunFams that we generated

from Pfam v32 (98) using the FunFam generation pipeline (99). GO terms were predicted using Pfam-FunFams in the same way as for CATH-FunFams. GO term annotations for UniProt accessions were downloaded in Feb 2020 and associated to FunFams and Pfam-FunFams. GO terms were up-propagated to their less specific parent terms in the ontology.

NET-FF protein function prediction by combining machine learning (NetHom) and CATHPredictGO predictions: Predicted protein functions from the machine learning and protein family-based methods were combined into a single dataset. If both methods predicted a particular GO term, with probabilities p_1 and p2, the combined probability was reported as min(p_1 + p_2, 1). GO taxon constraints were applied to remove GO terms that never occur in *S. pombe* taxa. GO term information content was added to each predicted GO term using the *S. pombe* GO annotation set (100). New terms that had not previously been annotated to *S. pombe* proteins were assigned the highest information content observed in the PomBase annotation set. Although both NetHom and CATHPredictGO exploit CATH FunFams they do so in different ways. Furthermore, since FunFams have been endorsed for precision in CAFA, we decided to upweight their contribution in this way.

Pairwise GO semantic similarity calculations

Gene Ontology terms from the combined machine-learning and protein family-based methods were filtered by Information Content (IC >5) and confidence score assigned by the predictor (p >0.7), reducing the number of genes with GO terms to 2167. All available genes and their GO terms predictions were collected from PomBase (n genes = 5396). For both datasets, we used GOGO (79) in an all-vs-all fashion to calculate the pairwise GO semantic similarity (GOGO score) for each pair of genes, i.e. genes from the NET-FF dataset (n=2167, nGO=22,060) with a GO probability of 0.7 and information content over 5, and all genes from PomBase with an associated set of GO terms (n=5396, nGO=45,542), resulting in 14 million and 1.5 million GOGO scores for PomBase and NET-FF, respectively. Results were subsequently separated by their ontology (Molecular Function; MFO) and Biological Process; BPO) and deduplicated. All 5396 genes in PomBase had GOGO scores for Biological Process, while 6 genes had missing scores for Molecular Function. The majority of genes in the NET-FF dataset had resulting GOGO scores, with 1674 genes having at least one positive GO:GO score for Biological Process and 1438 genes with a positive score for Molecular Function.

Information content calculation

The Information Content for GO terms was calculated as the negative log probability of the term occurring in all GO terms prediction for PomBase:

$$IC(GOterm) = -log(p(GOterm))$$

.

Where the frequency p(GOterm) is:

$$p(GO) = \frac{n_{GO'}}{N} |GO' \varepsilon \{GO, GO descendants\}|$$

nGO' is the number of annotations with the term GO' and N is the total number of GO terms in PomBase. The code to generate the IC content was adapted from the following Github gist: https://gist.github.com/avrilcoghlan/047a086c3b1b97071e177af6f0f1916d

Chronological lifespan measurements for experimental validation

We used our recently developed high-throughput assay (36) to determine the CLS of the strains in Figure 6. In brief, aliquots of ageing cultures were taken daily and serially diluted using an automated multichannel pipette (Integra Assist; Integra Biosciences Ltd). The serially diluted droplets were pinned in quadruplicate (384-well format) onto YES agar plates using a Singer RoToR HDA pinning robot (Singer Instruments). The plates were incubated at 32°C for 2-4 days until colonies were clearly visible. To collect images of agar plates, *pyphe-scan* was used with a transmission mode Epson V700 scanner (58,59). The R package *DeadOrAlive* was used to analyse the plate pictures and calculate the total number of colony forming units (CFUs) in the ageing cultures (36). The maximum CLS for each mutant was determined as the number of days from reaching stationary phase (100% viability) until the cells appeared to reach 0% viability (CFUs <1). For this experiment, 3 independent biological repeats were measured for each strain, except for the wild-type control and *SPCC4B3.06c* mutant which were measured in 2 independent repeats.

Acknowledgments

We thank Val Wood, Manuel Lera-Ramirez and Melania D'Angiolo for helpful comments on the manuscript, John Shawe-Taylor for advice on machine-learning approaches, and StJohn Townsend for exploring alternate clustering approaches. This work was funded by BBSRC Research Grant BB/R009597/1 to C.O. and J.B.

References

1. Kustatscher G, Collins T, Gingras AC, Guo T, Hermjakob H, Ideker T, et al. Understudied proteins: opportunities and challenges for functional proteomics. Nat Methods. 2022 Jul;19(7):774–9.

2. Edwards AM, Isserlin R, Bader GD, Frye SV, Willson TM, Yu FH. Too many roads not taken. Nature. 2011 Feb;470(7333):163–5.

3. Stoeger T, Gerlach M, Morimoto RI, Nunes Amaral LA. Large-scale investigation of the reasons why potentially important genes are ignored. Freeman T, editor. PLOS Biol. 2018 Sep 18;16(9):e2006643.

4. Pfeiffer T, Hoffmann R. Temporal patterns of genes in scientific publications. Proc Natl Acad Sci. 2007 Jul 17;104(29):12052–6.

5. Su AI, Hogenesch JB. Power-law-like distributions in biomedical publications and research funding. Genome Biol. 2007;8(4):404.

6. Haynes WA, Tomczak A, Khatri P. Gene annotation bias impedes biomedical research. Sci Rep. 2018 Dec;8(1):1362.

7. Oprea TI, Bologa CG, Brunak S, Campbell A, Gan GN, Gaulton A, et al. Unexplored therapeutic opportunities in the human genome. Nat Rev Drug Discov. 2018 May;17(5):317–32.

8. Wood V, Lock A, Harris MA, Rutherford K, Bähler J, Oliver SG. Hidden in plain sight: what remains to be discovered in the eukaryotic proteome? Open Biol. 2019 Feb;9(2):180241.

9. Sinha S, Eisenhaber B, Jensen LJ, Kalbuaji B, Eisenhaber F. Darkness in the Human Gene and Protein Function Space: Widely Modest or Absent Illumination by the Life Science Literature and the Trend for Fewer Protein Function Discoveries Since 2000. Proteomics. 2018 Nov;18(21–22):e1800093.

10. Gates AJ, Gysi DM, Kellis M, Barabási AL. A wealth of discovery built on the Human Genome Project — by the numbers. Nature. 2021 Feb 11;590(7845):212–5.

11. Hutchison CA, Chuang RY, Noskov VN, Assad-Garcia N, Deerinck TJ, Ellisman MH, et al. Design and synthesis of a minimal bacterial genome. Science. 2016 Mar 25;351(6280):aad6253–aad6253.

12. Müller JB, Geyer PE, Colaço AR, Treit PV, Strauss MT, Oroshi M, et al. The proteome landscape of the kingdoms of life. Nature. 2020 Jun;582(7813):592–6.

13. Rallis C, Bähler J. Cell-based screens and phenomics with fission yeast. Crit Rev Biochem Mol Biol. 2016;51(2):86–95.

14. Wood V, Gwilliam R, Rajandream MA, Lyne M, Lyne R, Stewart A, et al. The genome sequence of Schizosaccharomyces pombe. Nature. 2002 Feb 21;415(6874):871–80.

15. Marguerat S, Schmidt A, Codlin S, Chen W, Aebersold R, Bähler J. Quantitative analysis of fission yeast transcriptomes and proteomes in proliferating and quiescent cells. Cell. 2012 Oct 26;151(3):671–83.

16. Hoffman CS, Wood V, Fantes PA. An Ancient Yeast for Young Geneticists: A Primer on the Schizosaccharomyces pombe Model System. Genetics. 2015 Oct;201(2):403–23.

17. McGary KL, Park TJ, Woods JO, Cha HJ, Wallingford JB, Marcotte EM. Systematic discovery of nonobvious human disease models through orthologous phenotypes. Proc Natl Acad Sci U S A. 2010 Apr 6;107(14):6544–9.

18. Harris MA, Rutherford KM, Hayles J, Lock A, Bähler J, Oliver SG, et al. Fission stories: using PomBase to understand Schizosaccharomyces pombe biology. Genetics. 2021 Dec 23;iyab222.

19. Lock A, Rutherford K, Harris MA, Hayles J, Oliver SG, Bähler J, et al. PomBase 2018: user-driven reimplementation of the fission yeast database provides rapid and intuitive access to diverse, interconnected information. Nucleic Acids Res. 2019 Jan 8;47(D1):D821–7.

20. Deshpande GP, Hayles J, Hoe KL, Kim DU, Park HO, Hartsuiker E. Screening a genome-wide S. pombe deletion library identifies novel genes and pathways involved in genome stability maintenance. DNA Repair. 2009 May 1;8(5):672–9.

21. Calvo IA, Gabrielli N, Iglesias-Baena I, García-Santamarina S, Hoe KL, Kim DU, et al. Genome-wide screen of genes required for caffeine tolerance in fission yeast. PloS One. 2009 Aug 12;4(8):e6619.

22. García P, Encinar Del Dedo J, Ayté J, Hidalgo E. Genome-wide Screening of Regulators of Catalase Expression: Role of a Transcription Complex and Histone and tRNA Modification Complexes on Adaptation to Stress. J Biol Chem. 2016 Jan 8;291(2):790–9.

23. Dodgson J, Avula H, Hoe KL, Kim DU, Park HO, Hayles J, et al. Functional genomics of adhesion, invasion, and mycelial formation in Schizosaccharomyces pombe. Eukaryot Cell. 2009 Aug;8(8):1298–306.

24. Tange Y, Kurabayashi A, Goto B, Hoe KL, Kim DU, Park HO, et al. The CCR4-NOT complex is implicated in the viability of aneuploid yeasts. PLoS Genet. 2012;8(6):e1002776.

25. Sun LL, Li M, Suo F, Liu XM, Shen EZ, Yang B, et al. Global analysis of fission yeast mating genes reveals new autophagy factors. PLoS Genet. 2013;9(8):e1003715.

26. Guo L, Ganguly A, Sun L, Suo F, Du LL, Russell P. Global Fitness Profiling Identifies Arsenic and Cadmium Tolerance Mechanisms in Fission Yeast. G3 Bethesda Md. 2016 Oct 13;6(10):3317–33.

27. Hayles J, Wood V, Jeffery L, Hoe KL, Kim DU, Park HO, et al. A genome-wide resource of cell cycle and cell shape genes of fission yeast. Open Biol. 2013 May 22;3(5):130053.

28. Sajiki K, Tahara Y, Uehara L, Sasaki T, Pluskal T, Yanagida M. Genetic regulation of mitotic competence in G_0 quiescent cells. Sci Adv. 2018 Aug;4(8):eaat5685.

29. Rodríguez-López M, Gonzalez S, Hillson O, Tunnacliffe E, Codlin S, Tallada VA, et al. The GATA Transcription Factor Gaf1 Represses tRNAs, Inhibits Growth, and Extends Chronological Lifespan Downstream of Fission Yeast TORC1. Cell Rep. 2020 Mar 10;30(10):3240-3249.e4.

30. Lie S, Banks P, Lawless C, Lydall D, Petersen J. The contribution of non-essential Schizosaccharomyces pombe genes to fitness in response to altered nutrient supply and target of rapamycin activity. Open Biol. 2018 May;8(5).

31. Rallis C, López-Maury L, Georgescu T, Pancaldi V, Bähler J. Systematic screen for mutants resistant to TORC1 inhibition in fission yeast reveals genes involved in cellular ageing and growth. Biol Open. 2014 Feb 15;3(2):161–71.

32. Malecki M, Bähler J. Identifying genes required for respiratory growth of fission yeast. Wellcome Open Res. 2016;1:12.

33. Malecki M, Bitton DA, Rodríguez-López M, Rallis C, Calavia NG, Smith GC, et al. Functional and regulatory profiling of energy metabolism in fission yeast. Genome Biol. 2016 Nov 25;17(1):240.

34. Zuin A, Gabrielli N, Calvo IA, García-Santamarina S, Hoe KL, Kim DU, et al. Mitochondrial Dysfunction Increases Oxidative Stress and Decreases Chronological Life Span in Fission Yeast. Wölfl S, editor. PLoS ONE. 2008 Jul 30;3(7):e2842.

35. Sideri T, Rallis C, Bitton DA, Lages BM, Suo F, Rodríguez-López M, et al. Parallel profiling of fission yeast deletion mutants for proliferation and for lifespan during long-term quiescence. G3 Bethesda Md. 2014 Dec 1;5(1):145–55.

36. Romila CA, Townsend S, Malecki M, Kamrad S, Rodríguez-López M, Hillson O, et al. Barcode sequencing and a high-throughput assay for chronological lifespan uncover ageing-associated genes in fission yeast. Microb Cell Graz Austria. 2021 Jul 5;8(7):146–60.

37. Hu Y, Comjean A, Mohr SE, The FlyBase Consortium, Perrimon N. Gene2Function: An Integrated Online Resource for Gene Function Discovery. G3 GenesGenomesGenetics. 2017 Aug 1;7(8):2855–8.

38. Mungall CJ, McMurry JA, Köhler S, Balhoff JP, Borromeo C, Brush M, et al. The Monarch Initiative: an integrative data and analytic platform connecting phenotypes to genotypes across species. Nucleic Acids Res. 2017 Jan 4;45(D1):D712–22.

39. Szklarczyk D, Franceschini A, Wyder S, Forslund K, Heller D, Huerta-Cepas J, et al. STRING v10: protein-protein interaction networks, integrated over the tree of life. Nucleic Acids Res. 2015 Jan;43(Database issue):D447-452.

40. Gene Ontology Consortium. The Gene Ontology resource: enriching a GOld mine. Nucleic Acids Res. 2021 Jan 8;49(D1):D325–34.

41. Altenhoff AM, Levy J, Zarowiecki M, Tomiczek B, Warwick Vesztrocy A, Dalquen DA, et al. OMA standalone: orthology inference among public and custom genomes and transcriptomes. Genome Res. 2019 Jul;29(7):1152–63.

42. Brochado AR, Typas A. High-throughput approaches to understanding gene function and mapping network architecture in bacteria. Curr Opin Microbiol. 2013 Apr;16(2):199–206.

43. Hillenmeyer ME, Fung E, Wildenhain J, Pierce SE, Hoon S, Lee W, et al. The chemical genomic portrait of yeast: uncovering a phenotype for all genes. Science. 2008 Apr 18;320(5874):362–5.

44. Ryan CJ, Cimermančič P, Szpiech ZA, Sali A, Hernandez RD, Krogan NJ. High-resolution network biology: connecting sequence with function. Nat Rev Genet. 2013 Dec;14(12):865–79.

45. Wiwie C, Baumbach J, Röttger R. Comparing the performance of biomedical clustering methods. Nat Methods. 2015 Nov;12(11):1033–8.

46. Huttenhower C, Hibbs MA, Myers CL, Caudy AA, Hess DC, Troyanskaya OG. The impact of incomplete knowledge on evaluation: an experimental benchmark for protein function prediction. Bioinforma Oxf Engl. 2009 Sep 15;25(18):2404–10.

47. Libbrecht MW, Noble WS. Machine learning applications in genetics and genomics. Nat Rev Genet. 2015 Jun;16(6):321–32.

48. Angermueller C, Pärnamaa T, Parts L, Stegle O. Deep learning for computational biology. Mol Syst Biol. 2016 Jul 29;12(7):878.

49. Zhou N, Jiang Y, Bergquist TR, Lee AJ, Kacsoh BZ, Crocker AW, et al. The CAFA challenge reports improved protein function prediction and new functional annotations for hundreds of genes through experimental screens. Genome Biol. 2019 Nov 19;20(1):244.

50. You R, Zhang Z, Xiong Y, Sun F, Mamitsuka H, Zhu S. GOLabeler: improving sequence-based large-scale protein function prediction by learning to rank. Bioinforma Oxf Engl. 2018 Jul 15;34(14):2465–73.

51. Sanderson T, Bileschi ML, Belanger D, Colwell LJ. ProteInfer: deep networks for protein functional inference [Internet]. Bioinformatics; 2021 Sep [cited 2021 Oct 28]. Available from: http://biorxiv.org/lookup/doi/10.1101/2021.09.20.461077

52. Radivojac P, Clark WT, Oron TR, Schnoes AM, Wittkop T, Sokolov A, et al. A large-scale evaluation of computational protein function prediction. Nat Methods. 2013 Mar;10(3):221–7.

53. Jiang Y, Oron TR, Clark WT, Bankapur AR, D'Andrea D, Lepore R, et al. An expanded evaluation of protein function prediction methods shows an improvement in accuracy. Genome Biol. 2016 Dec;17(1):184.

54. Yao S, You R, Wang S, Xiong Y, Huang X, Zhu S. NetGO 2.0: improving large-scale protein function prediction with massive sequence, text, domain, family and network information. Nucleic Acids Res. 2021 Jul 2;49(W1):W469–75.

55. Gligorijević V, Barot M, Bonneau R. deepNF: deep network fusion for protein function prediction. Wren J, editor. Bioinformatics. 2018 Nov 15;34(22):3873–81.

56. Kim DU, Hayles J, Kim D, Wood V, Park HO, Won M, et al. Analysis of a genome-wide set of gene deletions in the fission yeast Schizosaccharomyces pombe. Nat Biotechnol. 2010 Jun;28(6):617–23.

57. Chen JS, Beckley JR, McDonald NA, Ren L, Mangione M, Jang SJ, et al. Identification of new players in cell division, DNA damage response, and morphogenesis through construction of Schizosaccharomyces pombe deletion strains. G3 Bethesda Md. 2014 Dec 31;5(3):361–70.

58. Kamrad S, Rodríguez-López M, Cotobal C, Correia-Melo C, Ralser M, Bähler J. Pyphe, a python toolbox for assessing microbial growth and cell viability in high-throughput colony screens. eLife. 2020 Jun 16;9:e55160.

59. Rodríguez-López M, Anver S, Cotobal C, Kamrad S, Malecki M, Correia-Melo C, et al. Functional profiling of long intergenic non-coding RNAs in fission yeast. eLife. 2022 Jan 5;11:e76000.

60. Rallis C, Codlin S, Bähler J. TORC1 signaling inhibition by rapamycin and caffeine affect lifespan, global gene expression, and cell proliferation of fission yeast. Aging Cell. 2013 Aug;12(4):563–73.

61. Duncan CDS, Rodríguez-López M, Ruis P, Bähler J, Mata J. General amino acid control in fission yeast is regulated by a nonconserved transcription factor, with functions analogous to Gcn4/Atf4. Proc Natl Acad Sci U S A. 2018 Feb 20;115(8):E1829–38.

62. Harris MA, Lock A, Bähler J, Oliver SG, Wood V. FYPO: the fission yeast phenotype ontology. Bioinforma Oxf Engl. 2013 Jul 1;29(13):1671–8.

63. Szklarczyk D, Gable AL, Nastou KC, Lyon D, Kirsch R, Pyysalo S, et al. The STRING database in 2021: customizable protein-protein networks, and functional characterization of user-uploaded gene/measurement sets. Nucleic Acids Res. 2021 Jan 8;49(D1):D605–12.

64. Vivancos AP, Jara M, Zuin A, Sansó M, Hidalgo E. Oxidative stress in Schizosaccharomyces pombe: different H2O2 levels, different response pathways. Mol Genet Genomics MGG. 2006 Dec;276(6):495–502.

65. Ikner A, Shiozaki K. Yeast signaling pathways in the oxidative stress response. Mutat Res. 2005 Jan 6;569(1–2):13–27.

66. Chen D, Wilkinson CRM, Watt S, Penkett CJ, Toone WM, Jones N, et al. Multiple pathways differentially regulate global oxidative stress responses in fission yeast. Mol Biol Cell. 2008 Jan;19(1):308–17.

67. Bitton DA, Schubert F, Dey S, Okoniewski M, Smith GC, Khadayate S, et al. AnGeLi: A Tool for the Analysis of Gene Lists from Fission Yeast. Front Genet [Internet]. 2015 Nov 16 [cited 2021 Oct 28];6. Available from: http://journal.frontiersin.org/Article/10.3389/fgene.2015.00330/abstract

68. Giaever G, Chu AM, Ni L, Connelly C, Riles L, Véronneau S, et al. Functional profiling of the Saccharomyces cerevisiae genome. Nature. 2002 Jul 25;418(6896):387–91.

69. Berry DB, Gasch AP. Stress-activated genomic expression changes serve a preparative role for impending stress in yeast. Mol Biol Cell. 2008 Nov;19(11):4580–7.

70. Zhou Y, Zhou B, Pache L, Chang M, Khodabakhshi AH, Tanaseichuk O, et al. Metascape provides a biologist-oriented resource for the analysis of systems-level datasets. Nat Commun. 2019 Apr 3;10(1):1523.

71. Pancaldi V, Schubert F, Bähler J. Meta-analysis of genome regulation and expression variability across hundreds of environmental and genetic perturbations in fission yeast. Mol Biosyst. 2010 Mar;6(3):543–52.

72. Mata J, Lyne R, Burns G, Bähler J. The transcriptional program of meiosis and sporulation in fission yeast. Nat Genet. 2002 Sep;32(1):143–7.

73. Dudin O, Merlini L, Bendezú FO, Groux R, Vincenzetti V, Martin SG. A systematic screen for morphological abnormalities during fission yeast sexual reproduction identifies a mechanism of actin aster formation for cell fusion. PLoS Genet. 2017 Apr;13(4):e1006721.

74. Chen D, Toone WM, Mata J, Lyne R, Burns G, Kivinen K, et al. Global Transcriptional Responses of Fission Yeast to Environmental Stress. Botstein D, editor. Mol Biol Cell. 2003 Jan;14(1):214–29.

75. Lackner DH, Beilharz TH, Marguerat S, Mata J, Watt S, Schubert F, et al. A network of multiple regulatory layers shapes gene expression in fission yeast. Mol Cell. 2007 Apr 13;26(1):145–55.

76. Hasan A, Cotobal C, Duncan CDS, Mata J. Systematic analysis of the role of RNA-binding proteins in the regulation of RNA stability. PLoS Genet. 2014 Nov;10(11):e1004684.

77. Raudvere U, Kolberg L, Kuzmin I, Arak T, Adler P, Peterson H, et al. g:Profiler: a web server for functional enrichment analysis and conversions of gene lists (2019 update). Nucleic Acids Res. 2019 Jul 2;47(W1):W191–8.

78. Sillitoe I, Bordin N, Dawson N, Waman VP, Ashford P, Scholes HM, et al. CATH: increased structural coverage of functional space. Nucleic Acids Res. 2021 Jan 8;49(D1):D266–73.

79. Zhao C, Wang Z. GOGO: An improved algorithm to measure the semantic similarity between gene ontology terms. Sci Rep. 2018 Dec;8(1):15107.

80. A Rahaman SN, Mat Yusop J, Mohamed-Hussein ZA, Aizat WM, Ho KL, Teh AH, et al. Crystal structure and functional analysis of human C1ORF123. PeerJ. 2018;6:e5377.

81. Zahedi Y, Durand-Dubief M, Ekwall K. High-Throughput Flow Cytometry Combined with Genetic Analysis Brings New Insights into the Understanding of Chromatin Regulation of Cellular Quiescence. Int J Mol Sci. 2020 Nov 27;21(23):9022.

82. Schumacher B, Pothof J, Vijg J, Hoeijmakers JHJ. The central role of DNA damage in the ageing process. Nature. 2021 Apr;592(7856):695–703.

83. Ohtsuka H, Azuma K, Murakami H, Aiba H. hsf1 (+) extends chronological lifespan through Ecl1 family genes in fission yeast. Mol Genet Genomics MGG. 2011 Jan;285(1):67–77.

84. Piazzon N, Maisonneuve C, Guilleret I, Rotman S, Constam DB. Bicc1 links the regulation of cAMP signaling in polycystic kidneys to microRNA-induced gene silencing. J Mol Cell Biol. 2012 Dec;4(6):398–408.

85. Liu X, Zhang H, Zhou P, Yu Y, Zhang H, Chen L, et al. CREB1 acts via the miR-922/ARID2 axis to enhance malignant behavior of liver cancer cells. Oncol Rep. 2021 May;45(5):79.

86. Mortensen RD, Serra M, Steitz JA, Vasudevan S. Posttranscriptional activation of gene expression in Xenopus laevis oocytes by microRNA-protein complexes (microRNPs). Proc Natl Acad Sci U S A. 2011 May 17;108(20):8281–6.

87. Zilio N, Codlin S, Vashisht AA, Bitton DA, Head SR, Wohlschlegel JA, et al. A novel histone deacetylase complex in the control of transcription and genome stability. Mol Cell Biol. 2014 Sep 15;34(18):3500–14.

88. Kamrad S, Bähler J, Ralser M. High-Throughput, High-Precision Colony Phenotyping with Pyphe. In: Devaux F, editor. Yeast Functional Genomics [Internet]. New York, NY: Springer US; 2022 [cited 2023 Mar 27]. p. 381–97. (Methods in Molecular Biology; vol. 2477). Available from: https://link.springer.com/10.1007/978-1-0716-2257-5_21

89. Morris JH, Apeltsin L, Newman AM, Baumbach J, Wittkop T, Su G, et al. clusterMaker: a multi-algorithm clustering plugin for Cytoscape. BMC Bioinformatics. 2011 Dec;12(1):436.

90. Szklarczyk D, Gable AL, Lyon D, Junge A, Wyder S, Huerta-Cepas J, et al. STRING v11: protein–protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets. Nucleic Acids Res. 2019 Jan 8;47(D1):D607–13.

91. The Gene Ontology Consortium. The Gene Ontology Resource: 20 years and still GOing strong. Nucleic Acids Res. 2019 Jan 8;47(D1):D330–8.

92. The UniProt Consortium. UniProt: a worldwide hub of protein knowledge. Nucleic Acids Res. 2019 Jan 8;47(D1):D506–15.

93. Steinegger M, Meier M, Mirdita M, Vöhringer H, Haunsberger SJ, Söding J. HH-suite3 for fast remote homology detection and deep protein annotation. BMC Bioinformatics. 2019 Dec;20(1):473.

94. Das S, Lee D, Sillitoe I, Dawson NL, Lees JG, Orengo CA. Functional classification of CATH superfamilies: a domain-based approach for protein function annotation. Bioinforma Oxf Engl. 2015 Nov 1;31(21):3460–7.

95. Mistry J, Finn RD, Eddy SR, Bateman A, Punta M. Challenges in homology search: HMMER3 and convergent evolution of coiled-coil regions. Nucleic Acids Res. 2013 Jul 1;41(12):e121–e121.

96. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, et al. Gene Ontology: tool for the unification of biology. Nat Genet. 2000 May;25(1):25–9.

97. Lewis TE, Sillitoe I, Lees JG. cath-resolve-hits: a new tool that resolves domain matches suspiciously quickly. Hancock J, editor. Bioinformatics. 2019 May 15;35(10):1766–7.

98. El-Gebali S, Mistry J, Bateman A, Eddy SR, Luciani A, Potter SC, et al. The Pfam protein families database in 2019. Nucleic Acids Res. 2019 Jan 8;47(D1):D427–32.

99. Adeyelu T, Bordin N, Waman VP, Sadlej M, Sillitoe I, Moya-Garcia AA, et al. KinFams: De-Novo Classification of Protein Kinases Using CATH Functional Units. Biomolecules. 2023 Feb 2;13(2):277.

100. Mistry M, Pavlidis P. Gene Ontology term overlap as a measure of gene functional similarity. BMC Bioinformatics. 2008 Aug 4;9:327.

Legends for figures, figure supplements and supplementary files

Figure 1. Schematic overview of study design.
Experimental and computational tasks along with the relationships between the different aspects of the work to generate new information on protein function.

Figure 2. Colony-based phenotyping of deletion mutants.
(A) Volcano plot of mutant colony sizes for priority unstudied genes (green) and all other genes (grey) growing in rich medium. The dashed vertical and horizontal lines show the 5% difference in effect size and significance threshold, respectively. Strains with lower fitness (smaller colonies) are <0 on the x-axis, and those with higher fitness are >0. We applied a significance threshold of 0.05 after Benjamini–Hochberg correction for multiple testing and a difference in fitness of abs(log2(mutant/wild type))> log2(0.05) to call hits based on colony size (n = 60).
(B) Distributions of significant hits per gene for priority unstudied mutants (right) and all other mutants (left) showing faster (green) or slower (purple) colony growth compared to wild-type cells.
(C) Distributions of significant hits per condition for all mutants showing faster (green) or slower (purple) colony growth compared to wild-type cells.
(D) Fitness of mutant cells growing in 2% glycerol+0.01% glucose in this study, plotted as in (A). Blue: mutants identified as slow growing in the same condition in our previous screen (32); red: mutants annotated as ‘increased cell population growth on glycerol carbon source’ (FYPO:0004167) in PomBase (62); green rings: mutants in priority unstudied genes. Seventy-nine mutants were slow growing in both screens ($p = 1x10^{-39}$; Fisher's exact test, FDR corrected for multiple testing).

Figure 2-figure supplement 1: Colony-based phenotyping of mutants in benign conditions.
(A) Left: Volcano plot of mutant colony sizes for priority unstudied genes (green) and all other genes (grey) growing in minimal medium. The dashed vertical and horizontal lines show the 5% difference in effect size and significance threshold, respectively (statistics as in Fig. 2A; n = 8). Right: Visualisation of overlapping genes among different sets as indicated using Metascape (70). We classified the mutants as fast- or slow-growing relative to their growth compared to the wild-type strain in minimal (EMM) and rich (YES) media and separately show the mutants in priority unstudied genes.
(B) Top GO and Reactome terms enriched for fast- and slow-growth mutants in minimal (EMM) and rich (YES) medium. The different mutant classes showing phenotypes were subjected to GO enrichment analysis using Metascape.

Figure 2-figure supplement 2: Our phenotype association matrix has much higher concordance with STRING than expected by chance.

Number of edges obtained from STRING protein interactions that are concordant with phenotype correlations (red vertical line) compared to 100 randomisations of the phenotype data (blue curve). For each STRING protein pair with a STRING-SCORE ≥750 the association was kept and the STRING-SCORE was multiplied by the phenotype-correlation score for the corresponding protein pair (obtained by measuring Pearson correlation of phenotype profiles for protein pair). In total, 425 protein pairs (edges) had a combined product of STRING-SCORE and phenotype-correlation score >0.5 (red dashed line). This number was never obtained by repeating the process 100 times after randomising the genes in the phenotype matrix. A maximum of only 79 protein pairs (edges) was obtained for the random process (blue curve). This result indicates that the STRING data show above random concordance with the phenotype data.

Figure 2-figure supplement 3: Clusters of mutant phenotype profiles for oxidative-stress conditions.

Hierarchical clustering of discretized data for 3449 deletion mutants (rows) in 6 conditions (columns). For discretizing the phenotype data, we classed mutants as either resistant (yellow), sensitive (blue), or similar (white) relative to their fitness in the corresponding control condition (as in Figure 3A). Basic clustering was performed in R using heatmap.2, with euclidean measure to obtain distance matrix and complete agglomeration method.

Figure 3. Mutant phenotype profiles.

(A) Clustering of discretized data for 3,449 deletion mutants (rows) in 98 conditions (columns). For discretizing the phenotype data, we classed mutants as either resistant (yellow), sensitive (blue), or similar (white) relative to their fitness in the corresponding control condition, applying an effect-size threshold of ≥10% and, for the two benign conditions, also a significance threshold of $p < 0.05$ (Supplementary file 1). To improve the clustering, we removed conditions and mutants that produced <4 phenotypes, the 3 conditions that produced >2400 phenotypes, and genes with >50% of missing values. Different types of conditions are colour-coded on top. Indicated at left are the 8 main clusters, calculated using *k*-medoids in R (pam function) and a distance of 1-Pearson correlation. Using the function fviz_nbclust (R package factoextra), we determined that 8 is an optimum number of clusters.

(B) Functional enrichments of the 8 gene clusters in (A) using GO slim categories (40) for Cellular Component (CC), Biological Process (BP) and Molecular Function (MF) (left, using SystempipeR) as well as KEGG pathways (right, using g:Profiler) (77).

Figure 4. NET-FF predictions of GO terms for *S. pombe* proteins.

(A) GOGO semantic similarity scores for GO terms predicted by NET-FF and GO terms annotated by PomBase, both for real (orange) and randomized (blue) datasets from NET-FF. Left: Biological Process data, right: Molecular function data.

(B) Information content distribution for GO terms predicted by NET-FF and GO terms annotated by PomBase, both for GO Biological Process (orange) and GO Molecular Function (blue) terms. Left: dataset for all genes, right: subset for priority unstudied genes.

Figure 4-figure supplement 1: Filtering of predictions by information content and probability value.

Illustration of the criteria for filtering the predictions (green area) by information content >5 (left) and by probability value >0.7 (right).

Figure 4-figure supplement 2: Performance of NET-FF predictor by source data.

Precision-recall curves plotted for each repeat (translucent curves) and micro-averaged data (solid curves). FF: FunFam homology data; NET: network data; NET-FF: combined FF and NET information; numbers indicate corresponding micro-averaged AUPR.

Figure 5. Experimental validation of new ageing-related gene reflecting its NET-FF GO predictions and phenotype-correlation network.

(A) Maximal CLS for eight selected deletion mutants of genes with NET-FF GO terms related to ageing/cell death and DNA-damage repair, measured using a high-throughput CLS assay (36). Known short- and long-lived mutants are highlighted in green and purple, respectively. Of the eight mutants, *ubc14* and *htr12* were known to be short- and long-lived, respectively (36). As controls, the maximal CLS of known short-lived (*alg14*) and long-lived (*pac3*) mutants (36) and wild-type cells (wt, red) are included. Three independent biological repeats were performed with actual datapoints shown as dots. Orange band: mutant in priority unstudied gene (*SPAC23C4.09c*) showing strong longevity phenotype.

(B) Cluster from Cytoscape network (see Figure 6) representing phenotype correlations between the *SPAC23C4.09c* gene-deletion mutant and ten other mutants. Orange edges show positive phenotype correlations, and red circles indicate genes with previously known CLS phenotypes (36,83). Details discussed in main text.

Figure 6. Example phenotype-correlation clusters from Cytoscape network. Positive and negative phenotype correlations across the 131 conditions are indicated as orange and blue edges, respectively. Gene pairs with high GO similarities based on PomBase annotations or NET-FF predictions are indicated as hatched pink and bright blue edges, respectively.

(A) Cluster 13 consisting of a branch with all members of the PAN complex (red) and a branch containing six genes involved in cAMP signalling (blue), along with two genes involved in RNA interference (green).

(B) Cluster 31 enriched for genes involved in vacuolar and endosomal transport (red) and Cluster 22 enriched for genes with peroxisome functions (green).

(C) Three clusters linking members of the same protein complexes, involved in protein maturation (NatC, green), regulation of transcription elongation (CTDK-1, red), and histone deacetylation (Clr6, black). The data reflect different Clr6 sub-complexes, with *png2* and *cti6* being positively correlated (Clr6 complex I') but *png2* and *png3* being negatively correlated (Clr6 complex I") (87).

(D) Two networks containing priority unstudied genes (green), including those with implicated mitochondrial functions (green with purple ellipses), together with their nearest neighbours based on phenotype correlation, which include genes with known mitochondrial functions (grey with purple ellipses). For simplicity, high GOGO pairs are not indicated in these networks.

Figure 6-figure supplement 1: GO association networks show higher concordance with our phenotype data than expected by chance.

(A) For each gene pair with a predicted GOGO BP (score >0.5) the GOGO score was multiplied by the PHEPHE score of the same gene pair (obtained from Pearson correlation of phenotype profiles across conditions for each gene). If the resulting combined product of GOGO and PHEPHE scores was >0.5, we added the gene pair link to a graph (network). For our dataset, this resulted in a graph of 31 edges (red line above). Graphs of this size were never obtained by repeating the process 100 times through randomising the PHEPHE matrix column gene labels before the lookup on the PHEPHE score (blue data). A maximum graph size of 17 edges was found through 100 randomisations. This result indicates that the GO predictions show above random concordance with the phenotype data.

(B) For each gene pair with a PomBase GOGO MF or BP score >0.7, the GOGO score was multiplied by the PHEPHE score of the same gene pair (obtained from Pearson correlation of phenotype profiles across conditions). If the resulting combined product of PomBase GOGO and PHEPHE scores was >0.5, we added the gene pair link to a graph (network). For our dataset, this resulted in a graph of over 2100 edges (red line above). Graphs of this size were never obtained by repeating the process 100 times through randomising the PHEPHE matrix column gene labels before the lookup on the PHEPHE score (blue data). This result indicates that the PomBase GO associations show above random concordance with the phenotype data.

Figure 6-figure supplement 2: Enrichment of high GOGO gene pair edges in nine clusters from the Cytoscape clustering of phenotype correlations, containing at least 15 genes. Six showed more high GOGO gene pairs based on NET-FF predictions than expected by chance.

Figure 6-figure supplement 3: Phenotype-correlation clusters from entire Cytoscape network. Examples are shown in Figure 6. The data can be interactively explored here: https://www.ndexbio.org/#/network/d4e109d2-05e9-11ed-ac45-0ac135e8bacf?accesskey=46cbcedb7360ef936f5cd68fb551ef8f9cd7c6bf068f279bc2c61e3e94e4a490

Supplementary file 1: Data for phenomics assays

Supplementary file 2: Data for NET-FF predictions and integrated analyses

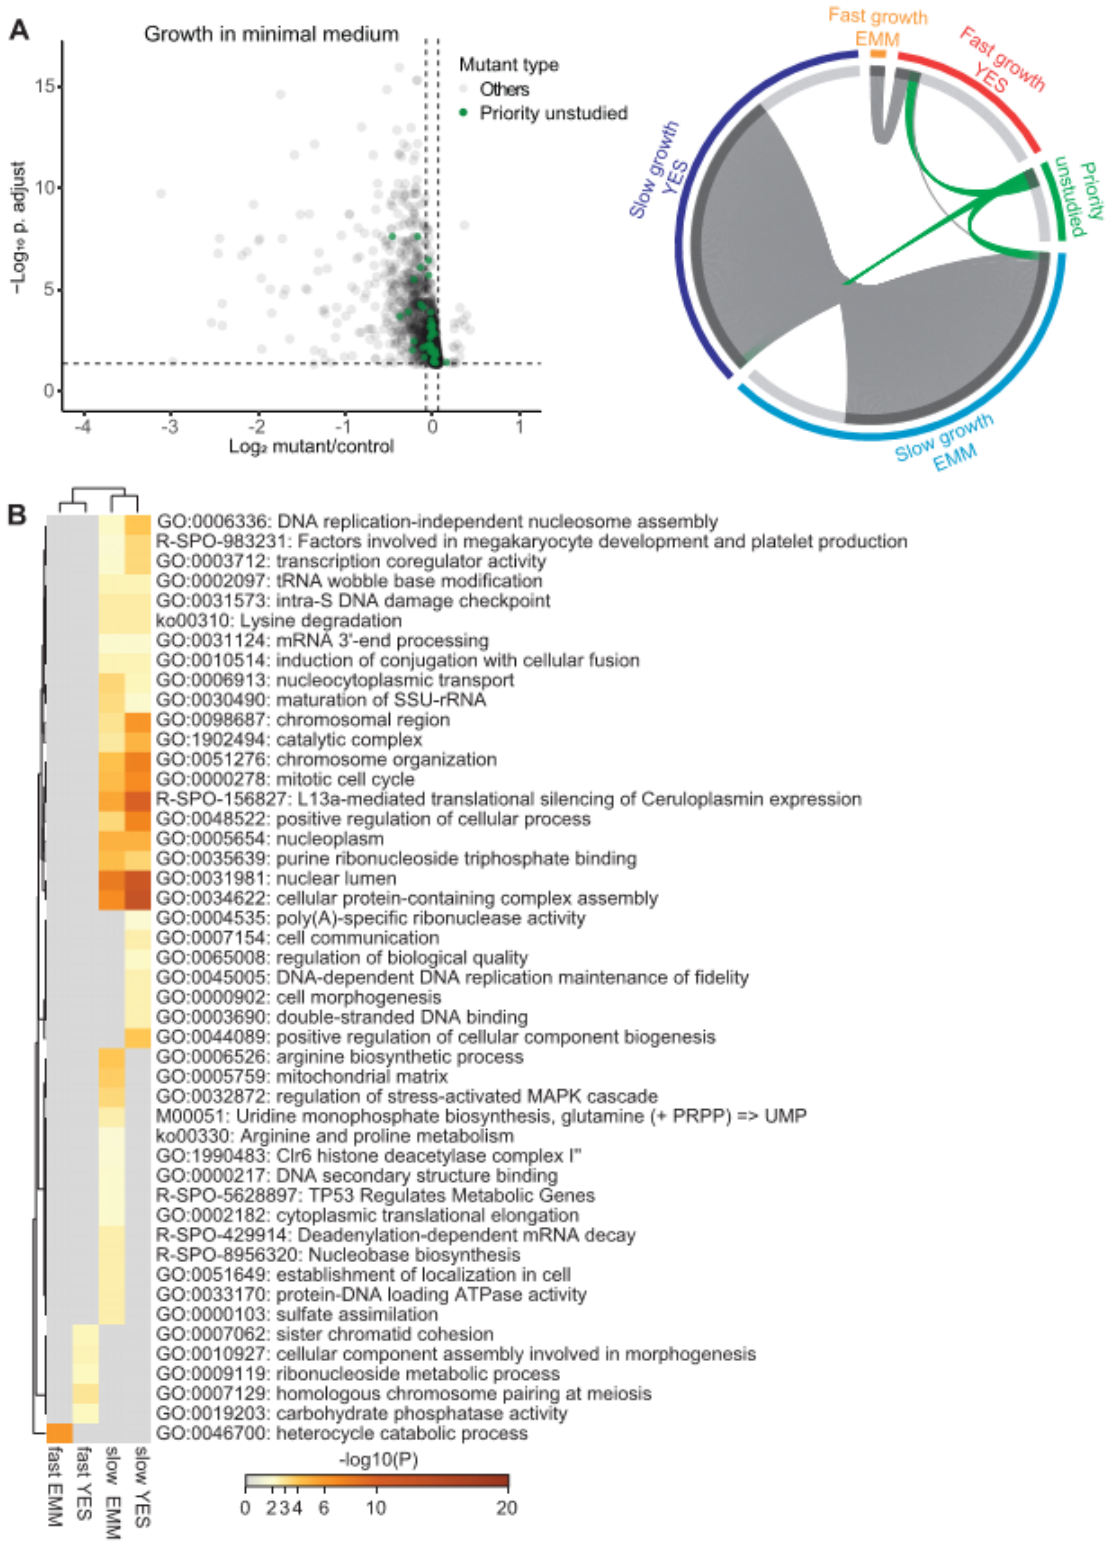

Figure 2-figure supplement 1: Colony-based phenotyping of mutants in benign conditions.
(A) Left: Volcano plot of mutant colony sizes for priority unstudied genes (green) and all other genes (grey) growing in rich medium. The dashed vertical and horizontal lines show the 5% difference in effect size and significance threshold, respectively (statistics as in Fig. 2A; n = 8). Right: Visualisation of overlapping genes among different sets as indicated using Metascape (1). We classified the mutants as fast- or slow-growing relative to their growth compared to the wild-type strain in minimal (EMM) and rich (YES) media and separately show the mutants in priority unstudied genes.
(B) Top GO and Reactome terms enriched for fast- and slow-growth mutants in minimal (EMM) and rich (YES) medium. The different mutant classes showing phenotypes were subjected to GO enrichment analysis using Metascape.

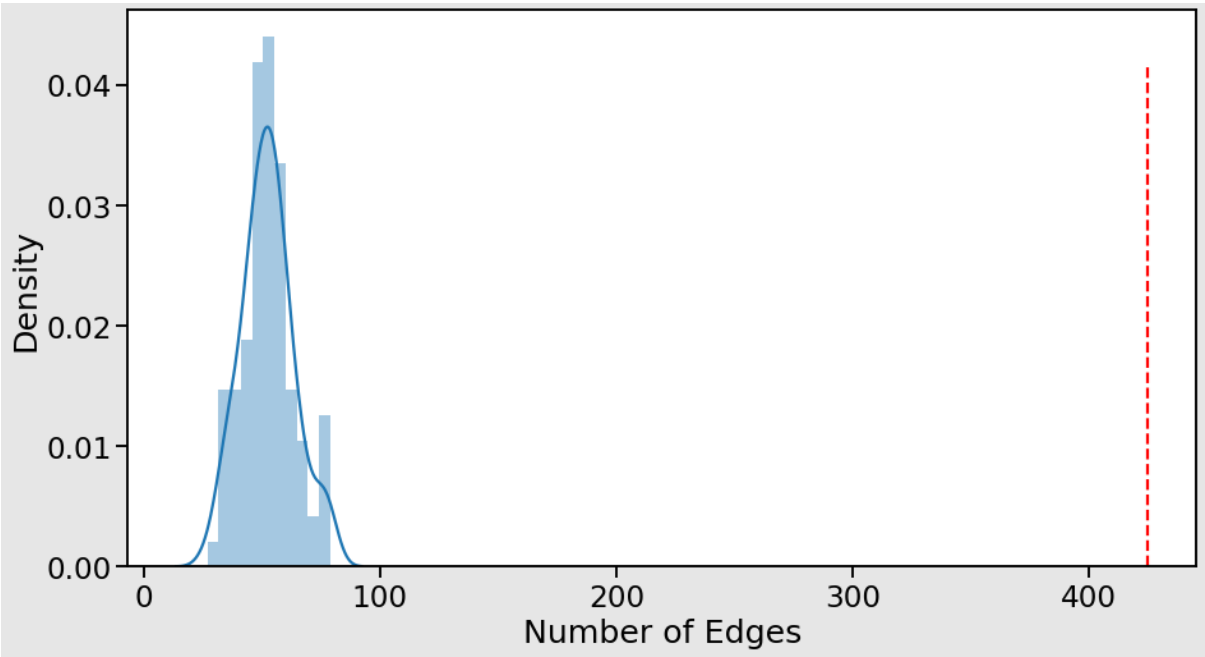

Figure 2-figure supplement 2: Our phenotype association matrix has much higher concordance with STRING than expected by chance.
Number of edges obtained from STRING protein interactions that are concordant with phenotype correlations (red vertical line) compared to 100 randomisations of the phenotype data (blue curve). For each STRING protein pair with a STRING-SCORE ≥750 the association was kept and the STRING-SCORE was multiplied by the phenotype-correlation score for the corresponding protein pair (obtained by measuring Pearson correlation of phenotype profiles for protein pair). In total, 425 protein pairs (edges) had a combined product of STRING-SCORE and phenotype-correlation score >0.5 (red hatched line). This number was never obtained by repeating the process 100 times after randomising the genes in the phenotype matrix. A maximum of only 79 protein pairs (edges) was obtained for the random process (blue curve). This result indicates that the STRING data show above random concordance with the phenotype data.

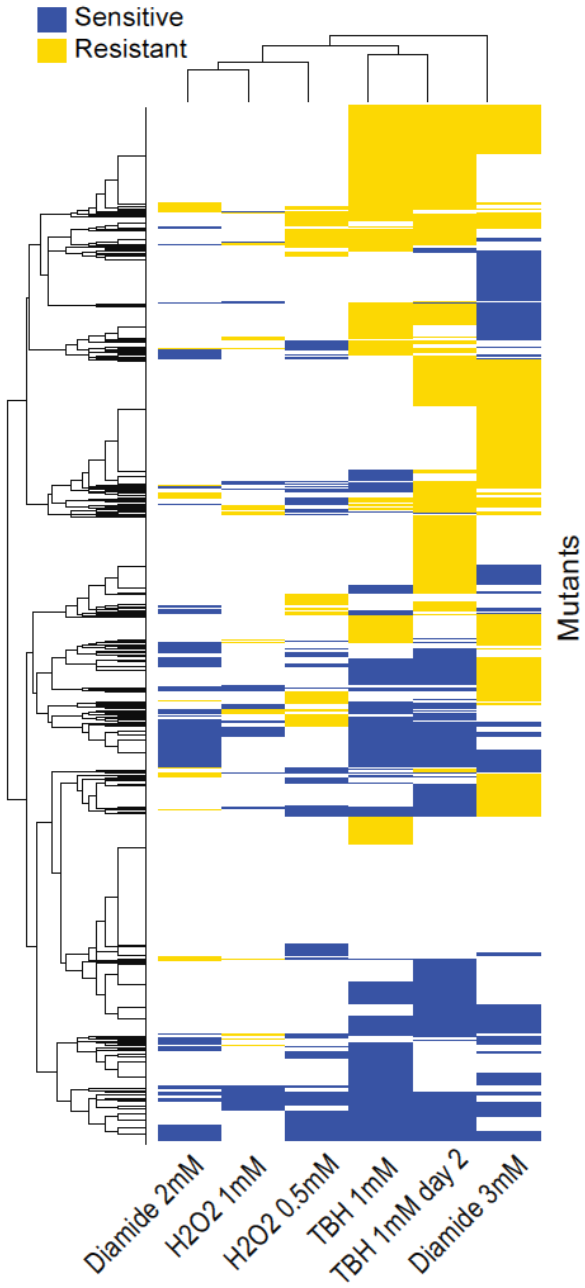

Figure 2-figure supplement 3: Clusters of mutant phenotype profiles for oxidative-stress conditions.
Hierarchical clustering of discretized data for 3449 deletion mutants (rows) in 6 conditions (columns). For discretizing the phenotype data, we classed mutants as either resistant (yellow), sensitive (blue), or similar (white) relative to their fitness in the corresponding control condition (as in Figure 3A). Basic clustering was performed in R using heatmap.2, with euclidean measure to obtain distance matrix and complete agglomeration method.

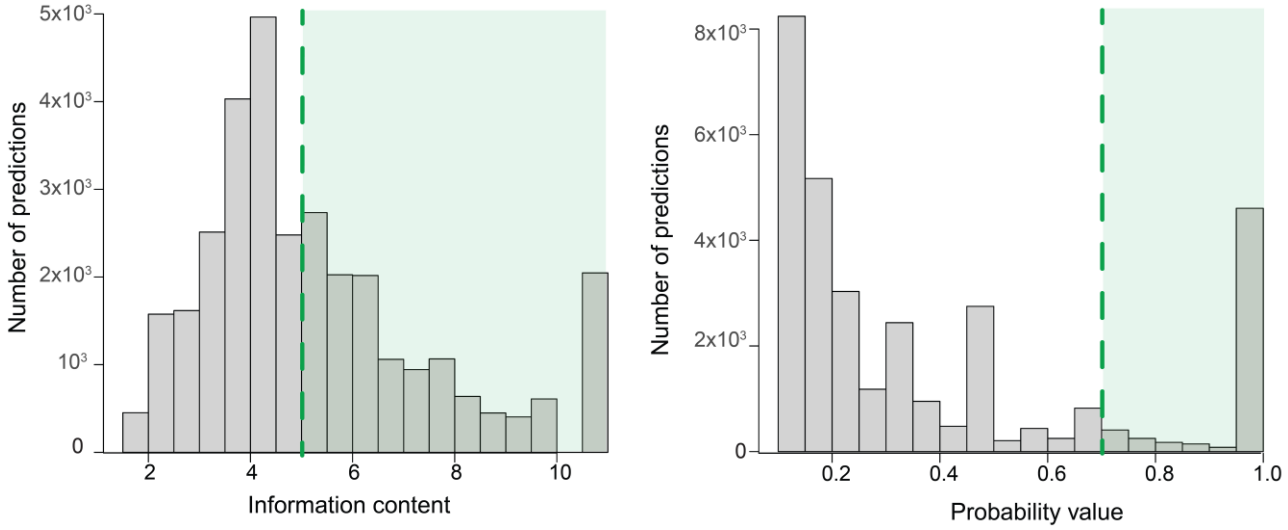

Figure 4-figure supplement 1: Filtering of predictions by information content and probability value.
Illustration of the criteria for filtering the predictions (green area) by information content >5 (left) and by probability value >0.7 (right).

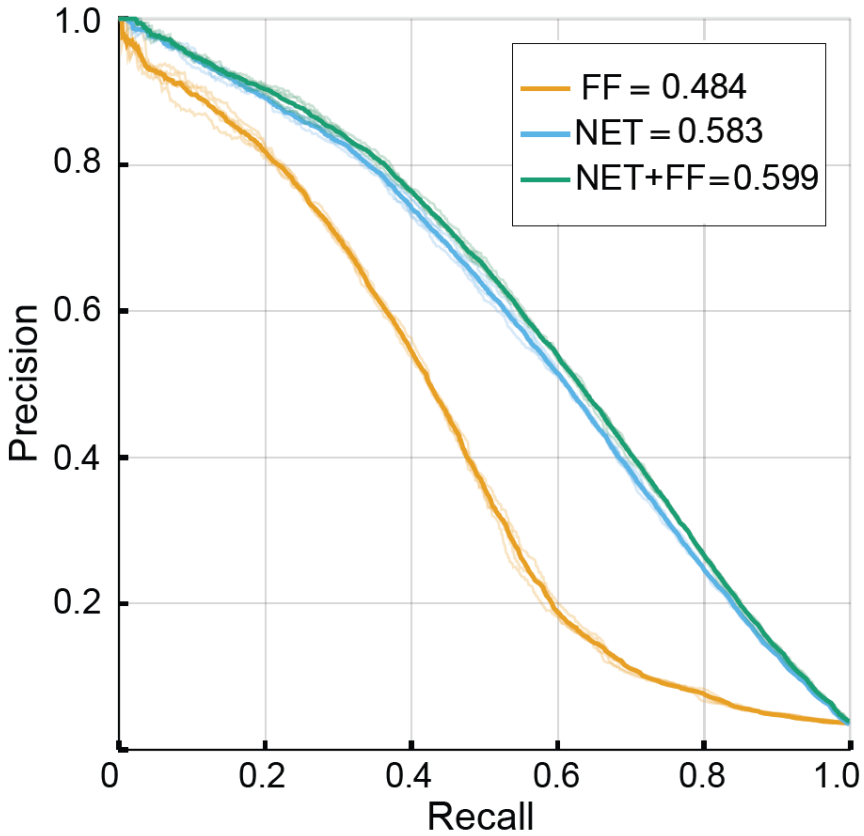

Figure 4-figure supplement 2: Performance of NET-FF predictor by source data.
Precision-recall curves plotted for each repeat (translucent curves) and micro-averaged data (solid curves). FF: FunFam homology data; NET: network data; NET-FF: combined FF and NET information; numbers indicate corresponding micro-averaged AUPR.

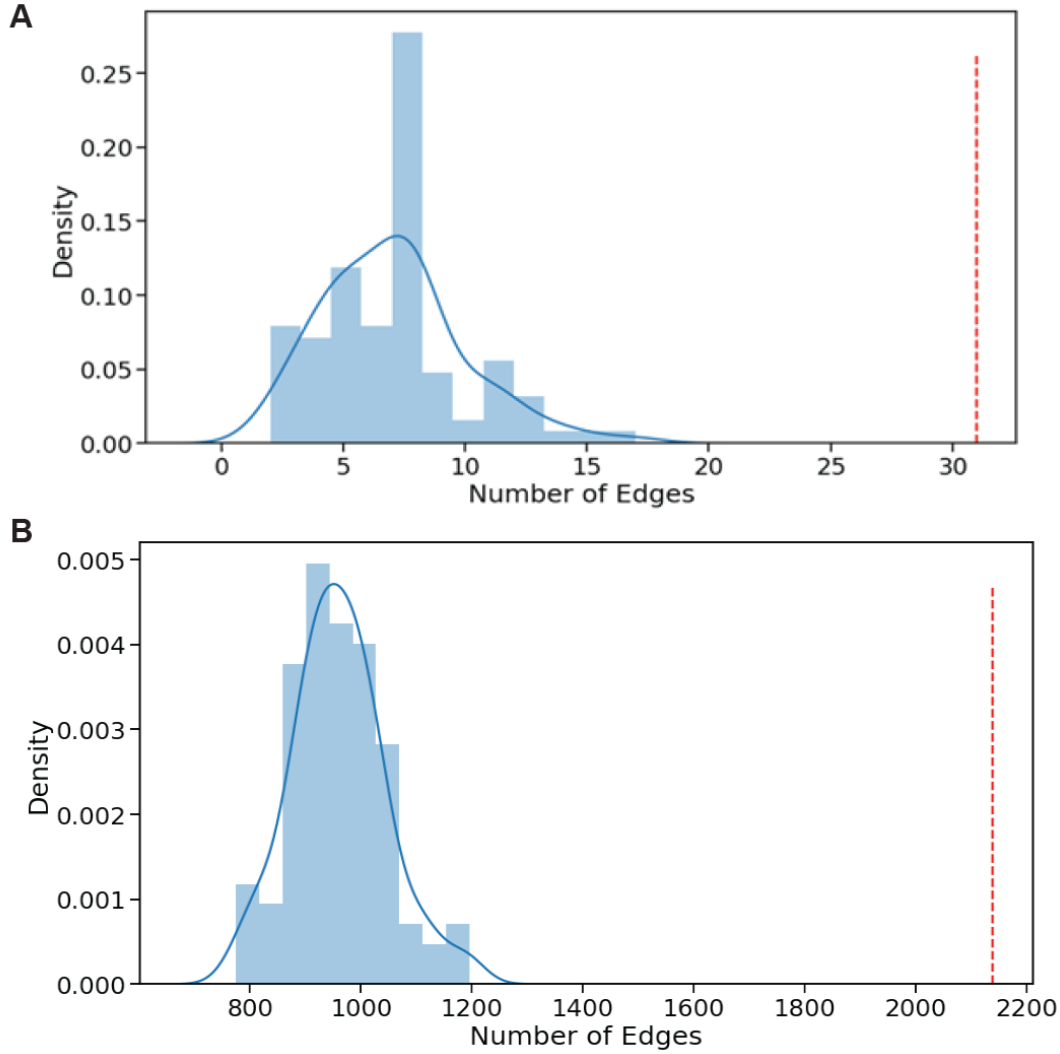

Figure 6-figure supplement 1: GO association networks show higher concordance with our phenotype data than expected by chance.
(A) For each gene pair with a predicted GOGO BP (score >0.5) the GOGO score was multiplied by the PHEPHE score of the same gene pair (obtained from Pearson correlation of phenotype profiles across conditions). If the resulting combined product of GOGO and PHEPHE scores was >0.5, we added the gene pair link to a graph (network). For our dataset, this resulted in a graph of 31 edges (red line above). Graphs of this size were never obtained by repeating the process 100 times through randomising the PHEPHE matrix column gene labels before the lookup on the PHEPHE score (blue data). A maximum graph size of 17 edges was found through 100 randomisations. This result indicates that the GO predictions show above random concordance with the phenotype data.
(B) For each gene pair with a PomBase GOGO MF or BP score >0.7, the GOGO score was multiplied by the PHEPHE score of the same gene pair (obtained from Pearson correlation of phenotype profiles across conditions). If the resulting combined product of PomBase GOGO and PHEPHE scores was >0.5, we added the gene pair link to a graph (network). For our dataset, this resulted in a graph of over 2100 edges (red line above). Graphs of this size were never obtained by repeating the process 100 times through randomising the PHEPHE matrix column gene labels before the lookup on the PHEPHE score (blue data). This result indicates that the PomBase GO associations show above random concordance with the phenotype data.

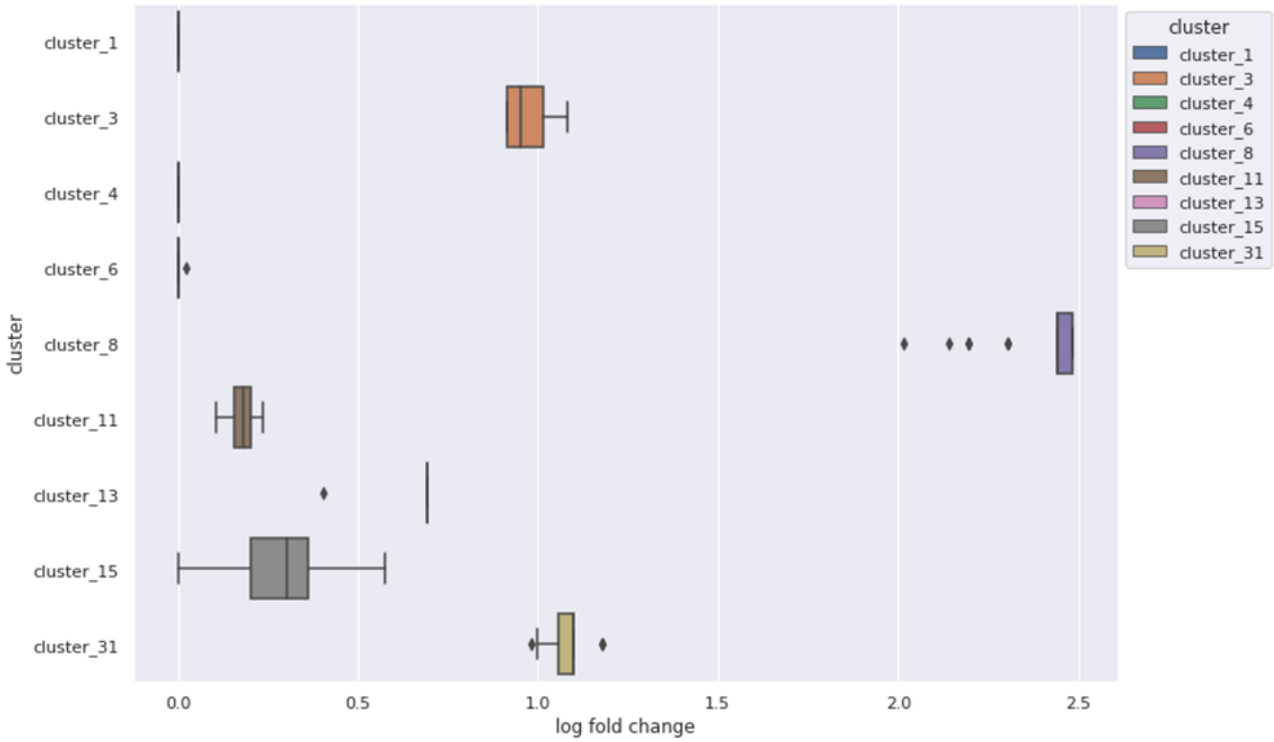

Figure 6-figure supplement 2: Enrichment of high GOGO gene pair edges in nine clusters from the Cytoscape clustering of phenotype correlations, containing at least 15 genes. Six showed more high GOGO gene pairs based on NET-FF predictions than expected by chance.

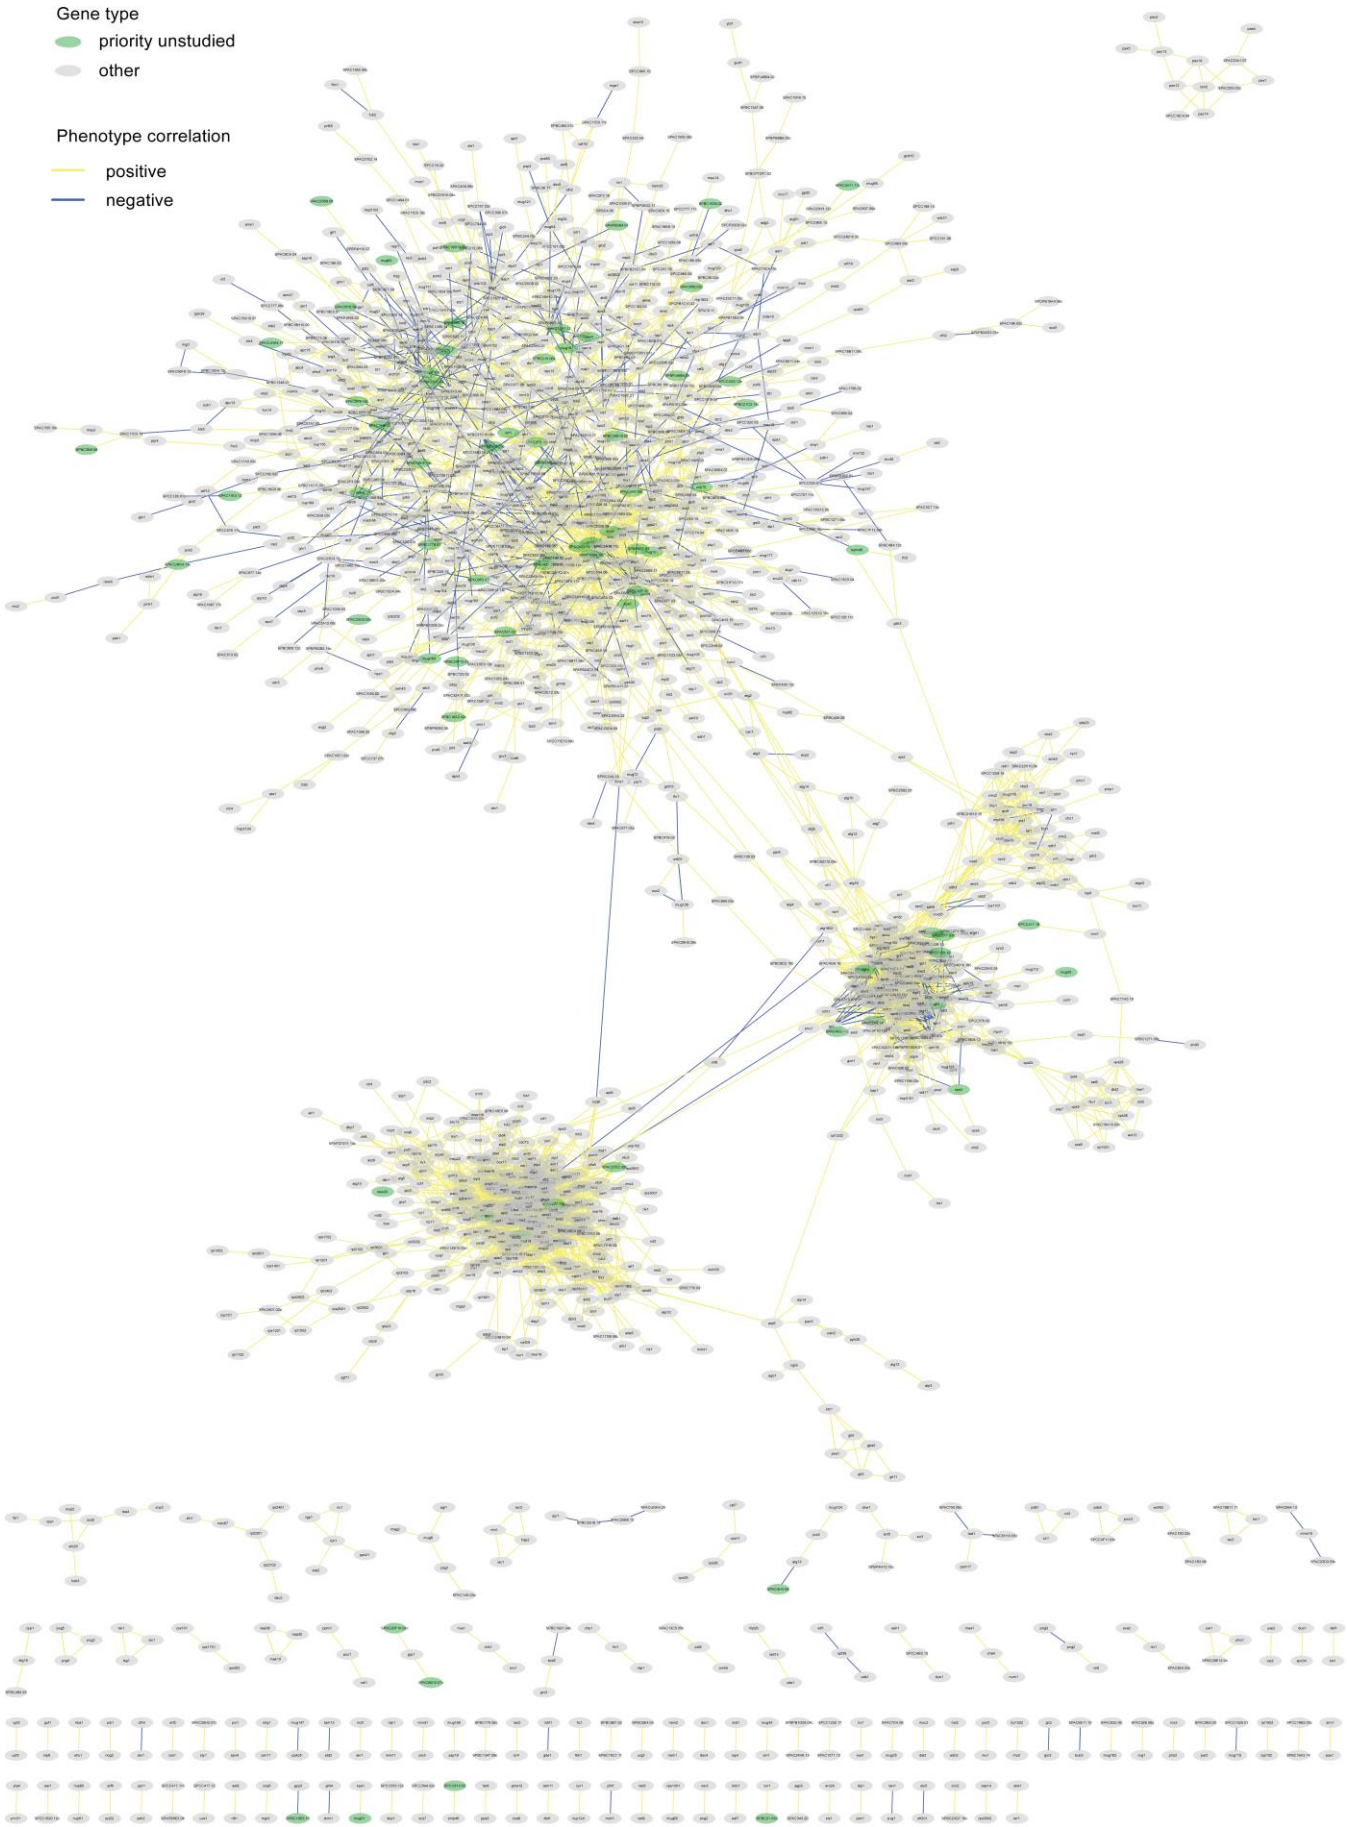

Figure 6-figure supplement 3: Phenotype-correlation clusters from entire Cytoscape network. Examples are shown in Fig. 6. The data can be interactively explored here: https://www.ndexbio.org/#/network/d4e109d2-05e9-11ed-ac45-0ac135e8bacf?accesskey=46cbcedb7360ef936f5cd68fb551ef8f9cd7c6bf068f279bc2c61e3e94e4a490